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

Data-driven and Model-based Verification: a Bayesian Identification Approach

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

Abstract— **This work develops a measurement-driven and model-based formal verification approach, applicable to systems with partly unknown dynamics. We provide a principled method, grounded on reachability analysis and on Bayesian inference, to compute the confidence that a physical system driven by external inputs and accessed under noisy measurements verifies a temporal logic property. A case study is discussed, where we investigate the bounded- and unbounded-time safety of a partly unknown linear time invariant system.**

I. Introduction

The design of complex, high-tech, safety-critical systems such as autonomous vehicles, intelligent robots, and cyber-physical infrastructures, demands guarantees on their correct and reliable behaviour. Correct functioning and reliability over models of systems can be attained by the use of formal methods. Within the computer sciences, the formal verification of software and hardware has successfully led to industrially relevant and impactful applications [8].

The strength of formal techniques, such as model checking, is bound to a fundamental requirement: having access to a model of the system of interest. In practice, for most physical systems their dynamical behaviour is known only in part. This holds in particular for biological systems [1] or for classes of engineered systems where, as a consequence, the use of uncertain dynamical models built from data is common practice [17].

Only limited work within the formal methods community deals with the verification of models with partly unknown dynamics. Classical results [3], [14] consider the verification problem for non-stochastic models described by differential equations and with bounded parametric uncertainty. Similarly, but for continuous-time probabilistic models, [5], [6] explore the parameter space with the objective of model verification. Whenever full state measurements of the system are available, Statistical Model Checking (SMC) [18], [24] replaces model(-based) verification procedures with empirical (statistical) testing. However SMC is limited to stochastic systems with little or no non-determinism, and may require gathering a large set of measurements. Extensions towards the inclusion of non-determinism have been studied in [13], [19], with preliminary steps towards Markov decision processes. Related to SMC techniques but bound to finite state models, [7], [21], [23] assume that the system is encompassed by a finite-state Markov chain and efficiently use data to learn a model and to consequently verify it.

An alternative approach put forward in this work, allowing both partly unknown dynamics over uncountable (continuous) variables and noisy output measurements, is the usage of a Bayesian framework relating the confidence in the asserted validity of a formal property to the uncertainty built over a model from data. This approach provides a new integration of model-based verification methods and data-driven inference techniques. When applied on nonlinearly parameterised linear time invariant (LTI) models this approach introduces computational issues, which as proposed in [10] can be mitigated by statistical methods. Instead, in order to obtain reliable numerical solutions, we propose the use of linearly parameterised model sets defined through orthonormal basis functions to represent partially unknown systems of interest. This is a broadly used framework in system identification [16], [17]: it allows for the incorporation of prior knowledge, while maintaining the benefits (computational aspects) of linear parameterisations. Practically, it has been widely used for the modelling of physical systems, such as the thermal dynamics of buildings [25]. In this contribution we extend the results in [12], obtained for a time-bounded subset of temporal logic properties, to unbounded-time temporal logic properties, and analyse their robustness.

The framework and problem statement are given in Section II. The main results both for bounded and unbounded-time temporal logic properties are discussed in Section III. A case study for a bounded-time safety property (Sec.III-B) and for its unbounded-time extension (Sec. III-D) is given to complement the theory. The proofs of the statements can be found in [11].

II. General Framework and Problem Statement

Denote with $\mathbf{S}$ a physical system, whose input signal $u(t) \in \mathbb{U}, t \in \mathbb{N}$, captures how the environment acts on the system, and whose output signal $y_0(t) \in \mathbb{Y}$ indicates how the system interacts with the environment, or how the system can be measured. The knowledge of the behaviour of the system is often limited or uncertain, and as such there is no "true" mathematical model available to analyse its behaviour. A-priori available knowledge allows us to construct a model set $\mathcal{G}$ with elements $\mathbf{M} \in \mathcal{G}$: this set supports the uncertainty about the "true" model as a distribution over a parameterisation $\theta \in \Theta$, $\mathcal{G} = \{\mathbf{M}(\theta)|\theta \in \Theta\}$. The unknown "true" model $\mathbf{M}(\theta^0)$ representing $\mathbf{S}$, is assumed to be an element of $\mathcal{G}$, namely $\theta^0 \in \Theta$.

Samples can be drawn from the underlying physical system via a measurement set-up, as depicted in Figure 1. An experiment consists of a finite number (N_s) of input-output samples drawn from the system, and is denoted by $Z^{N_s} = \{u(t)_{ex}, \tilde{y}(t)_{ex}\}_{t=1}^{N_s}$, where $u(t)_{ex} \in \mathbb{U}$ is the input for the experiment and $\tilde{y}(t)_{ex}$ is a (possibly noisy)

S. Haesaert and P.M.J. Van den Hof are with the Department of Electrical Engineering, Eindhoven University of Technology, NL. A. Abate is with the Department of Computer Science, University of Oxford, UK. Work 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 Oxford University Press (OUP) Research Fund.

measurement[1] of $y_0(t)_{ex}$. We assume that at the beginning of the measurement procedure (say at $t = 0$), the initial condition of the system, encompassed by the initial state $\mathbf{x}(0)_{ex}$ of models in $\mathbf{M}$, is either known, or, when not known, has a structured uncertainty distribution based on the knowledge of past inputs and/or outputs.

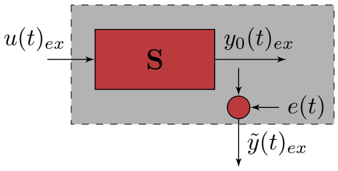

Fig. 1: System and measurement setup. In the measurement setup (grey box) the measured output $\tilde{y}(t)_{ex}$ includes the system output $y_0(t)_{ex}$ and the measurement noise $e(t)$. Data collected from experiments comprises the input $u(t)_{ex}$ and the measured output $\tilde{y}(t)_{ex}$ signals.

The end objective is to analyse the behaviour of system $\mathbf{S}$. We consider properties encoded as specifications ψ and expressed in a temporal logic of choice (to be detailed shortly). Let us remark that the behaviour of $\mathbf{S}$ to be analysed is bound to a set of operating conditions that are pertinent to the verification problem and that will be indexed with ver: this comprises the set of possible input signals $u(t)_{ver}$ (e.g., a white or coloured noise signal, or a non-deterministic signal $u(t)_{ver} \in \mathbb{U}_{ver} \subseteq \mathbb{U}$), and of the set of initial states $\mathbf{x}(0)_{ver} \in \mathbb{X}_{ver}$. The system satisfies a property if the "true" model representing it satisfies it, namely $\mathbf{S} \vDash \psi$ if and only if $\mathbf{M}(\theta^0) \vDash \psi$.

In this work we consider the satisfaction of a property $\mathbf{M}(\theta) \vDash \psi$ as a *binary-valued mapping* from the parameter space Θ. More generally, when in addition to the measurements of the system also its transitions are disturbed by noise, then property satisfaction is a mapping from the parameter space Θ to the interval $[0, 1]$, and quantifies the probability that the model $\mathbf{M}(\theta)$ satisfies the property.

Definition 1 (Satisfaction Function, [5]): Let $\mathcal{G}$ be a set of models $\mathbf{M}$ that is indexed by a parameter $\theta \in \Theta$, and let ψ be a formula in a suitable temporal logic. The satisfaction function $f_\psi : \Theta \to [0, 1]$ associated with ψ is

$$f_\psi(\theta) = \mathbf{P}\left(\mathbf{M}(\theta) \vDash \psi\right). \qquad \square$$

Let us assume that the satisfaction function f_ψ is measurable and entails a decidable verification problem for all $\theta \in \Theta$. The computation of the satisfaction function, or equivalently the exploration of a parameter set over a formal property, has been studied exclusively for continuous time, autonomous models in [3], [9], [14].

Problem 1: For a partly unknown physical system $\mathbf{S}$*, under prior knowledge on the system given as a parameterised model class* $\mathcal{G}$ *supporting an uncertainty distribution over the parameterisation, draw data from the measurement setup and verify properties on* $\mathbf{S}$ *expressed in a temporal logic of choice, while quantifying the confidence of the assertion.* □

A. A Bayesian Framework for Data-driven Verification

Denote loosely with $\mathbf{P}(\cdot)$ and $p(\cdot)$ respectively a probability measure and a probability density function, both defined over a continuous domain. We employ Bayesian probability calculus [20] to express the confidence in a property as a measure of the uncertainty distribution defined over the set $\mathcal{G}$. Within the Bayesian framework, uncertainty distributions are handled as probability distributions of random variables. Therefore the confidence in a property is computed as a probability measure $\mathbf{P}(\cdot)$, integrating the density $p(\cdot)$ of the uncertainty variable over $\mathcal{G}$.

Proposition 2 (Bayesian Confidence): Given a specification ψ and a data set Z^{N_s}, the confidence that $\mathbf{S} \vDash \psi$ can be quantified via inference as

$$\mathbf{P}\left(\mathbf{S} \vDash \psi \mid Z^{N_s}\right) = \int_\Theta f_\psi(\theta) p\left(\theta | Z^{N_s}\right) d\theta\ . \tag{1}$$

The *a-posteriori* uncertainty distribution $p\left(\theta|Z^{N_s}\right)$, given the data set Z^{N_s}; is based on parametric inference over θ as

$$p\left(\theta|Z^{N_s}\right) = \frac{p\left(Z^{N_s}|\theta\right)p(\theta)}{\int_\Theta p(Z^{N_s}|\theta)p(\theta)d\theta}\ , \tag{2}$$

which employs an uncertainty distribution $p(\theta)$ over the parameter set Θ, representing the prior knowledge. □

The statement can be formally derived based on standard Bayesian calculus [20]. In general (1)-(2) in Proposition 2 lack analytical solutions, and the assessment of the binary satisfaction function (1) may be computationally intensive. On the other hand, statistical methods, such as [10] based on Bayesian theory, lead to involved sampling and introduce additional uncertainty from Monte Carlo techniques.

On the contrary, in the next section, we propose a novel numerical approach over a class of discrete-time linear time-invariant systems. Exploiting linear parameterisations and considering properties expressed within a fragment of linear-time temporal logic (LTL) leads to computational procedures for the feasible sets and the confidence on the properties.

III. LTL Verification of LTI systems

Consider a system $\mathbf{S}$ that can be represented by a class of finite-dimensional dynamical models that evolve in discrete-time, and are linear, time-invariant (LTI), and not probabilistic. These models depend on input and output signals ranging over $\mathbb{R}$. The experimental measurement setup, as depicted in Figure 1, consists of the signals $u(t)_{ex}$ and $\tilde{y}(t)_{ex} = y_0(t)_{ex} + e(t)$, representing the inputs and the measured outputs, respectively, and where $e(t)$ is an additive zero-mean, white, Gaussian-distributed measurement noise with variance σ_e^2 that is uncorrelated from the inputs. N_s samples are collected within a data set $Z^{N_s} = \{u(t)_{ex}, \tilde{y}(t)_{ex}\}_{t=1}^{N_s}$.

System properties are expressed, over a finite set of atomic propositions $p_i \in AP$, $i = 1, \ldots, |AP|$, in Linear-time Temporal Logic [2]. LTL formulae are built recursively via the syntax $\psi ::= \text{true} \mid p \mid \neg\psi \mid \psi \wedge \psi \mid \psi \vee \psi \mid \bigcirc\psi \mid \psi \,\mathsf{U}\, \psi$. Let $\pi = \pi(0), \pi(1), \pi(2), \ldots \in \Sigma^{\mathbb{N}^+}$ be a string composed of letters from the alphabet $\Sigma = 2^{AP}$, then the satisfaction relation between π and ψ is denoted as $\pi \vDash \psi$. Denote the k-bounded and unbounded invariance operator as $\square^k\psi = \bigwedge_{i=0}^{k} \bigcirc^i \psi$ and $\square\psi = \neg(\mathtt{true}\,\mathsf{U}\,\neg\psi)$, respectively.

Of interest are formal properties encoded on the input-output behaviour of the system, and over a time horizon $t \geq 0$. The output $y_0(t)_{ver} \in \mathbb{Y}$ is labeled by a map $L : \mathbb{Y} \to \Sigma$,

[1] Both the operating conditions of the experiment, that is the input signal $u(t)_{ex}$ and the initial state $\mathbf{x}(0)_{ex}$, and the measurements have been indexed with ex to distinguish them from the operating conditions of interest for verification (ver), to be discussed shortly.

which assigns letters α in the alphabet Σ via half spaces on the output, as

$$L(y_0(t)_{ver}) = \alpha \in \Sigma \;\Leftrightarrow\; \bigwedge_{p_i \in \alpha} A_{p_i} y_0(t)_{ver} \leq b_{p_i}, \quad (3)$$

for given $A_{p_i}, b_{p_i} \in \mathbb{R}$ that is, sets of atomic propositions are associated to intervals over $\mathbb{Y} \subset \mathbb{R}$. Let us underline that properties are defined over the behaviour $y_0(t)_{ver}$ of the system, and not over the noisy measurements $\tilde{y}(t)_{ex}$ of the system in the measurement setup. Additionally, for the verification problem the input signal is modelled as a bounded signal $u(t) \in \mathbb{U}_{ver}$, and represents possible external non-determinism of the environment acting on the system. Introduce Θ_ψ to be the feasible set of parameters, such that for every parameter the property ψ holds iff $\theta \in \Theta_\psi$, i.e., $\forall \theta \in \Theta : \theta \in \Theta_\psi \Leftrightarrow \mathbf{M}(\theta) \vDash \psi$. As such Θ_ψ can be characterised as a level set of the satisfaction function, namely f_ψ, $\Theta_\psi = \{\theta \in \Theta : f_\psi(\theta) = 1\}$.

A. Verification of Bounded-Time Properties

Models $\mathbf{M}$ in the set $\mathcal{G}$ have the following representation:

$$\mathbf{M}(\theta): \quad \begin{cases} \mathbf{x}(t+1) &= A\mathbf{x}(t) + Bu(t), \\ \hat{y}(t,\theta) &= \theta^T \mathbf{x}(t) + Du(t), \end{cases} \quad (4)$$

and are parameterised by $\theta = \begin{bmatrix}\theta_1 & \dots & \theta_n\end{bmatrix}^T \in \Theta \subset \mathbb{R}^n$, with a prior probability distribution $p(\theta)$. For a given initial condition $\mathbf{x}(0)$ and input sequence, the output of the "true" model $\hat{y}(t, \theta^0)$ is equal to the system output $y_0(t)$.

Given operating conditions for the experiment set-up, the measured signal $\tilde{y}(t)_{ex}$ can be fully characterised: its probability density, conditional on the parameters θ, is

$$\begin{aligned} p\left(Z^{N_s}|\theta\right) &= \textstyle\prod_{t=1}^{N_s} p\left(\tilde{y}(t)_{ex}|\theta\right) \\ &= \frac{1}{\sqrt{\sigma_e^{2N_s}(2\pi)^{N_s}}} \exp\left[-\frac{\sum_{t=1}^{N_s}(\hat{y}(t,\theta)-\tilde{y}(t)_{ex})^2}{2\sigma_e^2}\right], \end{aligned}$$

and can be directly used in Proposition 2. This conditional density $p\left(Z^{N_s}|\theta\right)$ depends implicitly on the given initial state $\mathbf{x}(0)_{ex}$ and, for the case of a given uncertainty distribution for $\mathbf{x}(0)_{ex}$, $p\left(Z^{N_s}|\theta\right)$ should be marginalised over $\mathbf{x}(0)_{ex}$ [22]. The a-posteriori uncertainty distribution is obtained as the analytical solution of the parametric inference step in (2) [22].

Recall now that for a given specification ψ, we seek to determine the feasible set of parameters Θ_ψ, such that the corresponding models admit property ψ, namely $\mathbf{M}(\theta) \vDash \psi$, $\forall \theta \in \Theta_\psi$. Since models $\mathbf{M}(\theta)$ have a linearly-parameterised state space realisation as per (4), it follows that when the set of initial states and inputs $\mathbb{X}_{ver}$ and $\mathbb{U}_{ver}$ are bounded polyhedra, the verification of a class of safety properties expressed by formulae with labels as in (3) leads to a set of feasible parameters Θ_ψ that is a polyhedron, which can be easily computed. More precisely, following theorem can be derived.

Theorem 3 ([12]): Given a bounded polyhedral set (or equivalently a polytope) of initial states $\mathbf{x}(0) \in \mathbb{X}_{ver}$ and of inputs $u(t) \in \mathbb{U}_{ver}$ for $t \geq 0$, and considering a labelling map as in (3), then the feasible set Θ_ψ of the parameterised model set (4) is a polyhedron for properties ψ composed of the LTL fragment $\psi ::= \alpha|\bigcirc\psi|\psi_1 \wedge \psi_2$, with $\alpha \in \Sigma$. □

In the computation of the feasible set (see [11]), the faces of the polyhedron Θ_ψ are shown to be a function of the vertices[2] of the bounded set of initial states $\mathbb{X}_{ver}$ and of the set of inputs $\mathbb{U}_{ver}$, and are expected to grow in number as a function of the time horizon of the property.

The result in Theorem 3 is valid for any finite syntactical composition within the LTL fragment $\psi ::= \alpha|\bigcirc\psi|\psi_1 \wedge \psi_2$, as such it only holds for finite horizon properties. Properties defined over the infinite horizon will be the objective of Section III-C.

B. Case Study: Bounded-Time Safety Verification

Consider a system $\mathbf{S}$ and verify whether its output $y_0(t)_{ver}$ remains within the interval $\mathcal{I} = \begin{bmatrix}-0.5, \ 0.5\end{bmatrix}$ (which we label as ι), for the next 5 time steps, under $u(t)_{ver} \in \mathbb{U}_{ver} = [-0.2, \ 0.2]$ and $\mathbf{x}(0)_{ver} \in \{0_2\} = \mathbb{X}_{ver}$. To formalise the statement, introduce the alphabet $\Sigma = \{\iota, \tau\}$ and the labelling map $L : L(y) = \iota, \forall y \in \mathcal{I}$, $L(y) = \tau, \forall y \in \mathbb{Y} \setminus \mathcal{I}$. Under the stated conditions, check whether the following LTL property holds: $\mathbf{S} \vDash \bigwedge_{i=1}^{5} (\bigcirc)^i \iota$. We assume that system $\mathbf{S}$ can be represented as an element of a model set $\mathcal{G}$ encompassing transfer functions over a second-order Laguerre basis [15] (this is a special case of orthonormal basis functions), which translates to the following parameterised state-space representation:

$$\begin{aligned} \mathbf{x}(t+1) &= \left[\begin{smallmatrix} a & 0 \\ 1-a^2 & a \end{smallmatrix}\right]\mathbf{x}(t) + \left[\begin{smallmatrix} \sqrt{1-a^2} \\ (-a)\sqrt{1-a^2} \end{smallmatrix}\right]u(t), \\ \hat{y}(t,\theta) &= \theta^T \mathbf{x}(t) \,. \end{aligned} \quad (5)$$

The parameter set is selected as $\theta \in \Theta = [-10, 10]^2$, whereas the coefficient a is chosen to be equal to 0.4. We select, as prior available knowledge on the system, a uniform distribution $p(\theta)$ on the model class, and pick a known variance $\sigma_e^2 = 0.5$ for the white additive noise on the measurement. The set of feasible parameters $\Theta_\psi \subset \Theta$ is represented in Figure 2 and is computed according to Theorem 3. Based exclusively on prior available knowledge, the confidence associated to $\mathbf{S} \vDash \bigwedge_{i=1}^{5} (\bigcirc)^i \iota$ amounts to[3] 0.0165. Let us now set up an experiment on the system with "true parameter" $\theta_0 = [1 \ 0]^T$ (Figure 2): select an input signal $u(t)_{ex}$ as a realisation of white noise with a uniform distribution over $[-0.2, 0.2]$, and measure $\tilde{y}(t)_{ex}$ for 200 consecutive time instances. Note that in this case $\mathbb{U}_{ver}$ and $\mathbb{U}_{ex}$ are the same. The uncertainty distribution is then refined as $p\left(\theta|Z^{N_s}\right)$, and results in a confidence (1) in the property equal to 0.779.

We have repeated this test 100 times, and extended it to several instances of the parameter θ^0 characterising the underlying system $\mathbf{S}$. In all instances, after obtaining 200 measurements the a-posteriori probability leads to the confidence in the safety of the system, and is displayed in Table I via its mean and variance terms.

C. Verification of Unbounded-Time Properties

In this section we extend the approach unfolded in Section III-A, to hold on the LTL fragment $\psi ::= \alpha|\bigcirc\psi|\psi_1 \wedge \psi_2$ with additionally the *unbounded* invariance (safety) operator. Recall the form of the k-bounded and of the unbounded

[2] A polytope can be written as the convex hull of a finite set of vertices.

[3] This is obtained by computation of (1) with probability distribution $p(\theta)$: integrals are numerically solved in `Matlab`.

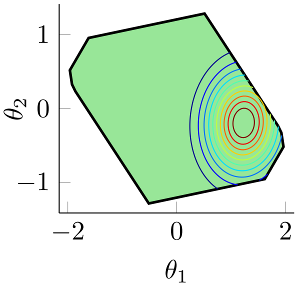

Fig. 2: Feasible set of parameters in Θ, and contour lines of the quantity $p\left(\theta|Z^{N_s}\right)$, obtained for $\theta^0 = [1\ 0]^T$.

invariance operators, namely $\square^k\psi = \bigwedge_{i=0}^{k} \bigcirc^i\psi$ and $\square\psi = \neg(\mathtt{true}\,\mathsf{U}\neg\psi)$ respectively. The extension from a k-bounded operator, covered by the result in Theorem 3, to the unbounded invariance one, is based on the concept of robust positive invariance [4, Def. 4.3], recalled next.

Definition 4: For the system $\mathbf{x}(t+1) = A\mathbf{x}(t) + Bu(t)$, the set $\mathcal{S} \subseteq \mathbb{X}$ is said to be robustly positively invariant if, for all $\mathbf{x}(0) \in \mathcal{S}$ and $u(t) \in \mathbb{U}_{ver}$, the condition $\mathbf{x}(t) \in \mathcal{S}$ holds for all $t \geq 0$. □

Recall that the feasible set Θ_ψ is defined as the set of parameters for which property ψ holds, namely $\forall\theta \in \Theta_\psi : \mathbf{M}(\theta) \vDash \psi$. The satisfaction relation $\mathbf{M}(\theta) \vDash \psi$ depends implicitly on the set of initial states $\mathbf{x}(0) \in \mathbb{X}_{ver}$ and on the set of inputs $\mathbb{U}_{ver}$. Let us extend the definition of the feasible set to explicitly account for its dependence on the set of initial conditions: given a bounded and convex set $\mathcal{S} \subset \mathbb{X}$, let $\Theta_\psi(\mathcal{S})$ be defined as the set of parameters in Θ for which the parameterised models $\mathbf{M}(\theta)$, initialised with $\mathbf{x}(0) \in \mathcal{S}$, satisfy ψ over input signals $u(t) \in \mathbb{U}_{ver}$, $t \geq 0$. Hence the feasible set Θ_ψ can be written as a function of the set of initial states $\mathbb{X}_{ver}$, that is $\Theta_\psi\left(\mathbb{X}_{ver}\right)$. Thus the extended map $\Theta_\psi\left(\cdot\right)$ takes subsets of the state space into subsets of the parameter space. Note that if $\mathcal{S}$ is a robustly positively invariant set that includes the set of initial states $\mathbb{X}_{ver} \subseteq \mathcal{S}$, then for all $\theta \in \Theta_\psi(\mathcal{S})$ the models $\mathbf{M}(\theta)$ satisfy ψ over all infinite-time model traces $\mathbf{x}(t)$: this allows to state that $\mathbf{M}(\theta) \vDash \square\psi$. We can show that the following holds.

Lemma 5: The function $\Theta_\psi(\cdot) : 2^{\mathbb{X}} \to 2^{\Theta}$, for specifications obtained as $\psi ::= \alpha \mid \bigcirc\psi \mid \psi_1 \wedge \psi_2$, is monotonically decreasing: that is if $\mathcal{S}_1 \subseteq \mathcal{S}_2$, then $\Theta_\psi(\mathcal{S}_2) \subseteq \Theta_\psi(\mathcal{S}_1)$. □

Based on the result in Lemma 5, we conclude that the maximal feasible set $\Theta_{\square\psi}$ is obtained as a mapping from the minimal robustly positively invariant set $\mathcal{S}$ that includes $\mathbb{X}_{ver}$: $\Theta_{\square\psi} = \Theta_\psi(\mathcal{S})$. This leads next to consider under which conditions such minimal robustly positively invariant set $\mathcal{S}$ can be exactly computed or approximated.

Feasible set for invariance properties with $\mathbb{X}_{ver} = \{0_n\}$: For $\mathbb{X}_{ver} = \{0_n\}$, assuming a bounded interval $\mathbb{U}_{ver}$ with the origin in its interior, and under some basic assumptions on the dynamics (to be shortly discussed), the minimal robustly positively invariant set can be shown to be a bounded and convex set that includes the origin [4]. Retaining the condition of $\mathbb{U}_{ver}$ being bounded and having the origin in its interior, we first consider the case that $\mathbb{X}_{ver} = \{0_n\}$ and characterise $\mathcal{S}$ via tools available from set theory in systems and control; thereafter we look at extensions to more general sets of initial states $\mathbb{X}_{ver}$.

TABLE I: Mean (μ) and variance (σ^2) of the confidence obtained from 100 experiments with 200 measurements each.

θ^0	μ	σ^2	θ^0	μ	σ^2
[-1 -1]^T	0.348	0.073	[1 -1]^T	0.491	0.085
[-1 0]^T	0.705	0.060	[1 0]^T	0.730	0.056
[-1 1]^T	0.492	0.086	[1 1]^T	0.339	0.065

Assume that $\mathbb{U}_{ver}$ includes the origin, and denote the forward reachability mappings initialised with $\mathcal{R}^{(0)} := \{0_n\} \subset \mathbb{X}$ as

$$\mathcal{R}^{(i)} := \mathrm{Post}(\mathcal{R}^{(i-1)}), \tag{6}$$

with set operation $\mathrm{Post}(X) := \{\mathbf{x}' = A\mathbf{x} + Bu, \mathbf{x} \in X, u \in_{ver}\}$. Denote the limit reachable set as $\mathcal{R}^\infty = \lim_{i\to\infty}\mathcal{R}^{(i)}$. From literature [4] we recall that properties of these i-step reachable sets include the following ones. For a reachable pair (A, B) and an asymptotically stable matrix A, the ∞-reachable set $\mathcal{R}^\infty$ is bounded and convex [4, Proposition 6.9]. The k-step reachable set converges to the ∞-reachable set via (6), since it is monotonically increasing $\mathcal{R}^{(i)} \subseteq \mathcal{R}^{(i+1)}$. Moreover, $\mathcal{R}^\infty$ is the minimal robustly positively invariant set for the system, so that any positively invariant set must include $\mathcal{R}^\infty$ [4, Proposition 6.13]. Thus, starting from $\mathbf{x}(0) = 0_n$, all $\mathbf{x}(t) \in \mathcal{R}^\infty$. Furthermore - of interest to this work - we conclude that $\Theta_{\square^k\psi} = \Theta_\psi\left(\mathcal{R}^{(k)}\right)$ and that $\Theta_{\square\psi} = \Theta_\psi\left(\mathcal{R}^\infty\right)$. For finite iterations the reachable sets $\mathcal{R}^{(i)}$ are polytopes, and if $\mathcal{R}^{(i)} = \mathcal{R}^{(i+1)}$, then $\mathcal{R}^{(i)} = \mathcal{R}^\infty$. Though the iterations can stop in finite time, in general the number of iterations to obtain $\mathcal{R}^\infty$ can be infinite. Whilst the minimal robustly positively invariant set is not necessarily closed or a polytope, there exist methods to approximate $\mathcal{R}^\infty$ as detailed in [4]. For instance, for stable systems, $\mathcal{R}^{(k)}$ is shown to converge to $\mathcal{R}^\infty$, in the sense that for all $\epsilon > 0$ there exists $\bar{k}$ such that for $k \geq \bar{k}$, $\mathcal{R}^{(k)} \subseteq \mathcal{R}^\infty \subseteq (1+\epsilon)\mathcal{R}^{(k)}$ [4, Proposition 6.9].

Feasible set under general polytopic initial states: Recall that the maximal feasible set $\Theta_{\square\psi}$ is obtained as a mapping from the minimal robustly positively invariant set $\mathcal{S}$ including $\mathbb{X}_{ver}$, that is $\Theta_{\square\psi} = \Theta_\psi(\mathcal{S})$. Under the condition $\mathbb{X}_{ver} \subseteq \mathcal{R}^\infty$ and ceteris paribus, then $\mathcal{R}^\infty$ is the minimal robustly positively invariant set that includes $\mathbb{X}_{ver}$, and $\Theta_\psi(\mathcal{R}^\infty) = \Theta_{\square\psi}$.

Let us now extend the study to the case where the conditions $\mathbb{X}_{ver} = \{0_n\}$ or its extension $\mathbb{X}_{ver} \subseteq \mathcal{R}^\infty$ do not apply, while the condition on the bounded set $\mathbb{U}_{ver}$ is maintained, that is $0 \in \mathbb{U}_{ver}$. Consider the more general case where the set of initial states is a polytope but not necessarily a subset of $\mathcal{R}^\infty$. Denote the union of the forward reachability mappings initialised with $\mathcal{R}^{(0)}_{\mathbb{X}_{ver}} := \mathbb{X}_{ver} \subseteq \mathbb{X}$ as

$$\mathcal{R}^{(i)}_{\mathbb{X}_{ver}} := \mathcal{R}^{(i-1)}_{\mathbb{X}_{ver}} \cup \mathrm{Post}(\mathcal{R}^{(i-1)}_{\mathbb{X}_{ver}})\ . \tag{7}$$

This set is also known in the literature as the *reach tube*. The corresponding set over the infinite time horizon is denoted as $\mathcal{R}^\infty_{\mathbb{X}_{ver}} = \lim_{i\to\infty}\mathcal{R}^{(i)}_{\mathbb{X}_{ver}}$. Notice that if $\mathbb{X}_{ver} \subseteq \mathcal{R}^\infty$, then $\mathcal{R}^\infty = \mathcal{R}^\infty_{\mathbb{X}_{ver}}$, as discussed above. The iteration is monotonically increasing $\mathcal{R}^{(i)}_{\mathbb{X}_{ver}} \subseteq \mathcal{R}^{(i+1)}_{\mathbb{X}_{ver}}$, and whenever $\mathcal{R}^{(i)}_{\mathbb{X}_{ver}} = \mathcal{R}^{(i+1)}_{\mathbb{X}_{ver}}$ it stops after a finite number of iterations with $\mathcal{R}^\infty_{\mathbb{X}_{ver}} = \mathcal{R}^{(i)}_{\mathbb{X}_{ver}}$. Of course, also in this more general case, the number of iterations can be unbounded, however the convergence properties of $\mathcal{R}^{(i)}$ extend seamlessly to the case of sets $\mathcal{R}^{(i)}_{\mathbb{X}_{ver}}$. Since $\mathcal{R}^{(i)}_{\mathbb{X}_{ver}}$ is a union of polytopes, it is not guaranteed to be a convex set. Still, it can be shown via the proof of Theorem 3 that the computation of the feasible set $\Theta_\psi(\mathcal{S})$ boils down to that of $\Theta_\psi\big(\operatorname{conv}(\mathcal{S})\big)$.

Robust approximations of the feasible set via $\Theta_\psi(\cdot)$: We exploit the convergence in the computation of the feasible set for invariance properties, in order to bound the error incurred in the approximations of the sets $\mathcal{R}^{\infty}_{\mathbb{X}_{ver}}$ and $\mathcal{R}^{\infty}$. Let $\mathcal{B}$ denote a unit ball centred at the origin and let the Hausdorff distance between sets $\mathcal{R}_1$ and $\mathcal{R}_2$ be defined as

$$\delta_H(\mathcal{R}_1,\mathcal{R}_2)=\inf\{\epsilon\geq 0|\mathcal{R}_1\subseteq\mathcal{R}_2+\epsilon\mathcal{B},\mathcal{R}_2\subseteq\mathcal{R}_1+\epsilon\mathcal{B}\}.$$

Lemma 6: Consider a polytope $\mathcal{R}$, and a property ψ comprised of $\psi ::= \alpha|\bigcirc\psi|\psi_1\wedge\psi_2$, with $\alpha\in\Sigma$, for which $\Theta_\psi(\mathcal{R})$ is a non-empty polytope with vertices v_i and the origin in its interior. Let A be bounded as $\|A\|_2\leq 1$. Then for any $\epsilon_x\geq 0$,

$$\Theta_\psi(\mathcal{R}+\epsilon_x\mathcal{B})\subseteq\Theta_\psi(\mathcal{R})\subseteq\Theta_\psi(\mathcal{R}+\epsilon_x\mathcal{B})+\epsilon_\theta\mathcal{B} \qquad (8)$$

$$\text{if } \epsilon_\theta\geq\frac{\epsilon_x\epsilon_p\max_i(\|v_i\|)^2}{1+\epsilon_x\epsilon_p\max_i(\|v_i\|)},\ \text{ for } \epsilon_p:=\max_{p\in AP}\frac{|A_p|}{|b_p|}. \qquad \square$$

Let us briefly discuss the conditions under which Lemma 6 is applicable. The condition that $\Theta_\psi(\mathcal{R})$ is not empty is raised to avoid the trivial case where $\Theta_\psi(\mathcal{R})=\emptyset$, and (8) holds for all ϵ_θ. The condition that $\Theta_\psi(\mathcal{R})$ is a polytope and hence bounded is necessary to obtain a bounded Hausdorff distance. This distance quantifies the difference between two sets, and is a necessary step to bound the approximation error. The requirement that $\Theta_\psi(\mathcal{R})$ includes the origin is a sufficient condition and relates to well-posedness for bounded input sets including the origin. When considering invariance properties defined for $0\in\mathbb{U}_{ver}$ and for any polytope $\mathbb{X}_{ver}$, the requirement that $0_n\in\Theta_\psi(\cdot)$ is necessary for $\Theta_{\Box\psi}$ to be non-empty: this can be intuitively illustrated by noting that under an assumption of asymptotic stability for A, for any θ and for $u(\cdot)=0$ the output $\hat{y}(t,\theta)$ of the model in (4) converges to 0. Hence for a property to be satisfied under these conditions it should at least hold for the zero output, which is equivalent to demanding that it holds for $\theta=0_n$. For any atomic proposition $p_i\in AP$ (see Eq. (3)) it can be shown that there is an invertible mapping between the row vectors, proportional to the normals of the faces of the polyhedral set $\Theta_{p_i}(\mathbf{x}(0))$, and the initial state $\mathbf{x}(0)$. Therefore, if $\mathcal{R}^{(k)}$ has the origin in its interior, then $\Theta_{p_i}(\mathcal{R}^{(k)})$ has to be bounded, and as a consequence so has any feasible set comprising this atomic proposition. This holds for $k\geq n$ if (A,B) is a reachable pair and if $\mathbb{U}_{ver}$ has 0 in its interior. Under the same conditions there exists a k such that $\mathcal{R}^{(k)}_{\mathbb{X}_{ver}}$ has 0_n in its interior. The generalisation to the case dealing with an Hausdorff distance of the feasible set for invariance properties with a set of inputs $0\notin\mathbb{U}_{ver}$ is outside of the scope of this work.

Convergence properties: We can employ Lemma 6 to bound the Hausdorff distance between $\Theta_\psi(\mathcal{R}^{(k)}_{\mathbb{X}_{ver}})$ and $\Theta_{\Box\psi}$. If $\mathbb{X}_{ver}=\{0_n\}$ and the spectral radius of A is strictly less than 1 (that is $\rho(A)<1$), then the Hausdorff distance can be bounded as

$$\delta_H(\mathcal{R}^{(k)},\mathcal{R}^{\infty})\leq\epsilon(k):=\|A^k\|_2\max_{u\in\mathbb{U}}(|u|)c_1, \qquad (9)$$

with c_1 a bound on $\sum_{i=0}^{\infty}\|A^iB\|$, which is the peak-to-peak performance of the dynamical system obtained from matrices (A,B). In the case that $\mathbb{X}_{ver}\not\subseteq\mathcal{R}^{\infty}$ then the forward reachable iteration can be rewritten as $\mathcal{R}^{(k)}_{\mathbb{X}_{ver}}=\big(\bigcup_{i=0}^{k}A^i\mathbb{X}_{ver}\big)+\mathcal{R}^{(k)}$. The Hausdorff norm can be bounded as $\delta_H(\mathcal{R}^{(k)}_{\mathbb{X}_{ver}},\mathcal{R}^{\infty}_{\mathbb{X}_{ver}})\leq\epsilon(k)+\|A^{k+1}\|_2\delta_H\left(\mathbb{X}_{ver},\{0_n\}\right)$. Note that for $\rho(A)<1$ the norm $\|A^k\|_2\to 0$ for $k\to\infty$. In case the conditions of Lemma 6 on $\mathcal{R}^{(k)}_{\mathbb{X}_{ver}}\subseteq\mathbb{X}$ and $\Theta_\psi\big(\mathcal{R}^{(k)}_{\mathbb{X}_{ver}}\big)$ hold, the Hausdorff distance $\delta_H(\Theta_{\Box^k\psi},\Theta_{\Box\psi})$ can be bounded by

$$\|A^k\|_2\max_i(\|v_i\|)^2\epsilon_p\big(\max_{u\in\mathbb{U}}(|u|)c_1+\|A\|\delta_H(\mathbb{X}_{ver},\{0_n\})\big). \qquad (10)$$

Use in the verification of unbounded-time properties: Based on the convergence properties of the feasible set, the asymptotic behaviour of the confidence computed in Proposition 2 can be stated as follows.

Corollary 7 (Convergence): Under the conditions of Lemma 6, for a Gaussian distribution $p\left(\theta\right)\sim\mathcal{N}\left(\mu_\theta,R_\theta\right)$ with a covariance $R_\theta\succ 0$, $\mathbf{P}\left(\theta\in\Theta_{\Box^k\psi}\right)\to\mathbf{P}\left(\theta\in\Theta_{\Box\psi}\right)$ for $k\to\infty$. $\square$

Theorem 3 can now be generalised to include unbounded-time invariance properties as follows.

Theorem 8: Consider a polytopic set of initial states $x(0)\in\mathbb{X}_{ver}$, inputs $u(t)\in\mathbb{U}_{ver}$ for $t\geq 0$, and a labelling map as in (3). Let $\hat{\mathcal{R}}^{\infty}_{\mathbb{X}_{ver}}$ be a polytopic superset of the minimal robustly positively invariant set that includes $\mathbb{X}_{ver}$, denoted as $\mathcal{R}^{\infty}_{\mathbb{X}_{ver}}$. Then the feasible set admits a polyhedral subset $\hat{\Theta}_\psi\subset\Theta_\psi$ for every specification ψ expressed within the LTL fragment $\psi:=\alpha|\bigcirc\psi|\psi_1\wedge\psi_2|\Box\psi$, and if $\hat{\mathcal{R}}^{\infty}_{\mathbb{X}_{ver}}=\mathcal{R}^{\infty}_{\mathbb{X}_{ver}}$ then $\hat{\Theta}_\psi=\Theta_\psi$. $\square$

D. Case Study (cont.): Unbounded-Time Safety Verification

We study convergence properties for the safety specification ι considered in the case study of Section III-B, maintaining the same operating conditions as before for the safety verification and the experimental setup. In Figure 3a the forward reachability sets $\mathcal{R}^{(k)}$ with $k=1,\ldots,20$ are obtained for the model dynamics in (5). Figure 4 (upper plot) displays bounds $\epsilon(k)$ on the Hausdorff distances $\delta_H(\mathcal{R}^{(k)},\mathcal{R}^{\infty})$ computed with (9): starting from $\mathcal{R}^{(1)}$ (slanted polygon) as in Figure 3a, it can be observed that the

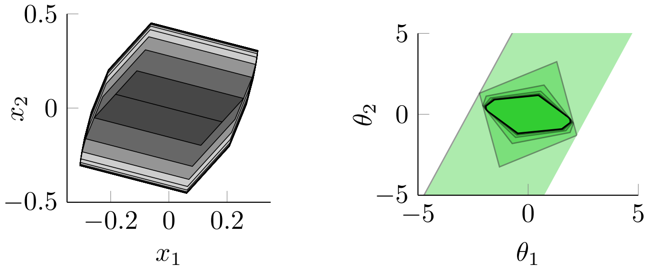

(a) The first 20 iterations of the forward reachable set $\mathcal{R}^{(k)}$, $k=1,\ldots,20$ for the case study. The reachable sets grow in size from dark grey ($k=1$) to light grey ($k=20$), so that $\mathcal{R}^{(k-1)}\subseteq\mathcal{R}^{(k)}$.

(b) The feasible sets for the k-bounded invariance property $\Box^k\iota$, with $k=1,\ldots,20$ (from lighter to darker color), obtained for the case study.

Fig. 3: Reachable and feasible sets for unbounded-time verification problem.

forward reachable sets $\mathcal{R}^{(k)}$ converge rapidly, as confirmed with the error bound displayed in Figure 4 (upper plot).

Based on $\mathcal{R}^{(k)}$, the feasible set for the k-bounded invariance $\square^k\iota$ can be computed as $\Theta_{\square^k\iota} = \Theta_\iota\big(\mathcal{R}^{(k)}\big)$. The feasible sets $\Theta_{\square^k\iota}$ with $k = 1, \ldots, 20$ are plotted in Figure 3b. Observe that the feasible set $\Theta_{\square^1\iota}$ is not bounded, but for $k \geq 2$ the feasible sets are bounded and, as expected, decrease in size with time. In Figure 4 (middle plot) bounds on the Hausdorff distances $\delta_H(\Theta_{\square\iota}, \Theta_{\square^k\iota})$ are given for $k = 2, \ldots, 20$ (no finite bound is computed for the index $k = 1$, since for that instance the feasible set is not bounded). Let us conclude this case study looking at confidence quantification, as a function of the time horizon. Figure 4 (lower plot) represents the confidence over the property $\mathbf{P}\left(\theta \in \Theta_{\square^k\iota} \mid Z^{N_s}\right)$, for indices $k = 1, \ldots, 20$. Unlike the case discussed in Section III-B, which focused on looking at statistics of the confidence via mean and variance drawn over multiple experiments, we zoom in on asymptotic properties by considering a data set Z^{N_s} comprising a single trace made up of 200 measurements, simulated under the same conditions as in Section III-B, and with $\theta_0 = [1\ 0]^T$. From the resulting probability density distribution $p\left(\theta \mid Z^{N_s}\right)$, it is observed that the confidence converges rapidly to the displayed nonzero values.

IV. Discussion and Generalisations

The computational bottleneck of the discussed approach resides on the characterisation and computation of the feasible set. The characterisation based on polytopes allows for an analytical expression, which however may not scale to models with very large dimension (the number of half-planes characterising the feasible set may increase with the time bound of the LTL formula ψ and with the cardinality of the atomic propositions in the alphabet Σ). Further its numerical computation (necessary for infinite-horizon properties) incurs similar limitations.

Note that these computations are essentially similar to known reachability operations, therefore the method is extendable well beyond the 2-dimensional case study when applying sophisticated reachability analysis tools in the literature.

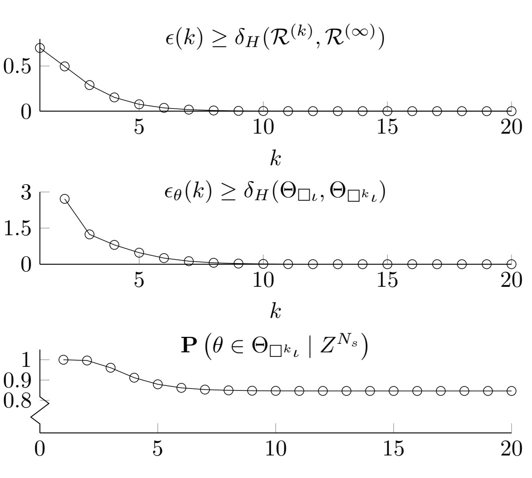

Fig. 4: (Upper plot) Error bound on the approximation level of the k-th forward reachable sets, which is such that $\mathcal{R}^{(\infty)} \subseteq \mathcal{R}^{(k)} + \epsilon(k)$ for $k = 1, \ldots, 20$. (Middle plot) The Hausdorff distance $\epsilon_\theta(k)$ between $\Theta_{\square^k\psi}$ and $\Theta_{\square\psi}$ with $k = 2, \ldots, 20$, obtained for the case study.(Lower plot) Confidence that $\mathbf{S} \vDash \square^k\iota$ for $k = 1, \ldots, 20$ for the case in Section III-B, with a new experiment consisting of 200 samples collected as Z^{N_s}.

References

[1] A. Abate, R. C. Hillen, and S. A. Wahl. Piecewise affine approximation of fluxes and enzyme kinetics from in-vivo ^{13}C labeling experiments. *International Journal of Robust and Nonlinear Control*, pages 1120–1139, 2012. Special Issue on System Identification for Biological Systems.

[2] C. Baier and J.-P. Katoen. Principles of model checking. *MIT Press*, 2008.

[3] G. Batt, C. Belta, and R. Weiss. Model checking genetic regulatory networks with parameter uncertainty. In *HSCC*, pages 61–75. Springer, 2007.

[4] F. Blanchini and S. Miani. *Set-Theoretic Methods in Control*. Birkhäuser Basel, 1st edition, 2007.

[5] L. Bortolussi and G. Sanguinetti. Smoothed model checking for uncertain continuous time Markov chains. *CoRR*, abs/1402.1450, 2014.

[6] L. Brim, M. Češka, S. Dražan, and D. Šafránek. Exploring parameter space of stochastic biochemical systems using quantitative model checking. In N. Sharygina and H. Veith, editors, *CAV*, volume 8044 of *LNCS*, pages 1–17. Springer, 2013.

[7] Y. Chen and T. D. Nielsen. Active learning of Markov decision processes for system verification. In *Conf. on Machine Learning and Applications*, pages 289–294, 2012.

[8] E. M. Clarke. The birth of model checking. In *25 Years of Model Checking*, pages 1–26, 2008.

[9] G. Frehse, S. K. Jha, and B. H. Krogh. A counterexample-guided approach to parameter synthesis for linear hybrid automata. In *HSCC*, pages 187–200. Springer Berlin Heidelberg, 2008.

[10] B. M. Gyori, D. Paulin, and S. K. Palaniappan. Probabilistic verification of partially observable dynamical systems. *CoRR*, abs/1411.0976, 2014.

[11] S. Haesaert, P. M. J. Van den Hof, and A. Abate. Data-driven and Model-based Verification: a Bayesian Identification Approach. *ArXiv e-prints*, Sept. 2015. 1509.03347.

[12] S. Haesaert, P. M. J. Van den Hof, and A. Abate. Data-driven property verification of grey-box systems by Bayesian experiment design. In *American Control Conference*, pages 1800–1805, 2015.

[13] D. Henriques, J. G. Martins, P. Zuliani, A. Platzer, and E. M. Clarke. Statistical model checking for Markov decision processes. In *QEST*, pages 84–93, 2012.

[14] T. Henzinger and H. Wong-Toi. Using hytech to synthesize control parameters for a steam boiler. In *Formal Methods for Industrial Applications*, pages 265–282. Springer Berlin Heidelberg, 1996.

[15] P. S. C. Heuberger, P. M. J. Van den Hof, and O. H. Bosgra. A generalized orthonormal basis for linear dynamical systems. *Automatic Control, IEEE Transactions on*, 40(3):451–465, 1995.

[16] P. S. C. Heuberger, P. M. J. Van den Hof, and B. Wahlberg. *Modelling and identification with rational orthogonal basis functions*. Springer London, 2005.

[17] H. Hjalmarsson. From experiment design to closed-loop control. *Automatica*, pages 393–438, 2005.

[18] A. Legay, B. Delahaye, and S. Bensalem. Statistical model checking: An overview. In H. Barringer, Y. Falcone, B. Finkbeiner, K. Havelund, I. Lee, G. Pace, G. Roşu, O. Sokolsky, and N. Tillmann, editors, *Runtime Verification*, volume 6418 of *LNCS*, pages 122–135. Springer Berlin Heidelberg, 2010.

[19] A. Legay and S. Sedwards. Lightweight Monte Carlo algorithm for Markov decision processes. *CoRR*, abs/1310.3609, 2013.

[20] D. V. Lindley. The philosophy of statistics. *Journal of the Royal Statistical Society: Series D (The Statistician)*, pages 293–337, 2000.

[21] H. Mao and M. Jaeger. Learning and model-checking networks of I/O automata. In *Proc. of Asian Conference on Machine Learning*, 2012.

[22] V. Peterka. Bayesian Approach to System Identification. *Trends Prog. Syst. Identif.*, 1981.

[23] K. Sen, M. Viswanathan, and G. Agha. Learning continuous time Markov chains from sample executions. In *QEST*, pages 146–155, 2004.

[24] K. Sen, M. Viswanathan, and G. Agha. Statistical model checking of black-box probabilistic systems. In R. Alur and D. Peled, editors, *CAV*, volume 3114 of *LNCS*, pages 202–215. Springer, 2004.

[25] G. S. Virk and D. L. Loveday. Model-based control for HVAC applications. In *Conf. on Control Applications*, pages 1861–1866. IEEE, 1994.