

A Notion of Approximation for Systems over Finite Alphabets

Danielle C. Tarraf

Abstract— We consider the problem of approximating plants with discrete sensors and actuators (termed ‘systems over finite alphabets’) by deterministic finite memory systems for the purpose of certified-by-design controller synthesis. We propose a new, control-oriented notion of input/output approximation for these systems, that builds on ideas from robust control theory and behavioral systems theory. We conclude with a brief discussion of the key features of the proposed notion of approximation relative to those of two existing notions of finite state approximation and abstraction.

I. INTRODUCTION

High fidelity models that accurately describe a dynamical system are often too complex to use for controller design. The problem of finding a lower complexity model that approximates a given plant has thus been extensively studied and continues to receive much deserved attention (see [3], [12] and the references therein). In order to be truly useful in this setting, a model complexity reduction approach should provide both a lower complexity model and a rigorous assessment of the quality of approximation that would simultaneously allow one to systematically and efficiently design a controller based on the lower complexity model as well as quantify its performance when implemented in the actual system (assumed to be faithfully captured by the high complexity model).

One of the big success stories has been for the class of linear time-invariant systems. In this setting, model complexity is captured by the order of the system, and the standard model complexity reduction problem (equivalently, the standard LTI model order reduction problem) can be stated as follows: Given a stable LTI plant P of order n , find a stable LTI plant \tilde{P} of order $m < n$ such that the \mathcal{H}_∞ norm of the transfer function of the difference system $P - \tilde{P}$ is minimized. Intuitively, \tilde{P} is considered to be a good approximation of P if the difference between the outputs of the two systems, when initialized to zero and driven side by side using the same input, is not too large relative to the size of the input, in the worst case scenario. Several suboptimal solutions have been developed for this problem, each associated with provable a priori error bounds that are useful in the context of robust control [1], [7], [13], [11]. Three points are key to the utility of this formulation of the model order reduction problem in the control setting, and should be emphasized here: First, the ability of \mathcal{L}_2 gain conditions to capture many performance objectives of interest. Second, the existence of systematic methods for designing \mathcal{H}_∞ controllers, whose complexity scales with the order of

the nominal models. Third, the existence of \mathcal{L}_2 based tools for characterizing the robustness of a system to perturbations, such as those due to approximation uncertainty.

More recently, the problem of approximating hybrid systems, that is systems involving interacting analog and discrete dynamics, by simpler systems has been considered [2], [19], [4]. In particular, the problem of deriving *finite* state approximations of hybrid systems has been the object of intense study over the past two decades. Two main notions of approximation have been systematically explored in this setting: First, ‘qualitative models’, essentially non-deterministic finite state machines whose output or input/output behavior contains that of the original model [14], [10]. The problem of controller synthesis in this case reduces to a standard supervisory control problem [15]. Second, finite simulation or bisimulation ‘abstractions’ of the plant, generally in some approximate sense [6], [16]. Effectively, the set of all output sequences of the original model is contained (exactly or to within some distance) in the set of all output sequences of the lower complexity model. The problem of controller synthesis in this case is a two step procedure whereby a supervisory controller is first designed for the finite abstraction and subsequently refined to a hybrid controller that can be implemented in the original system [17].

In [20], [18], we demonstrated the use of specially constructed finite state approximations (which are neither qualitative models nor simulation/bisimulation abstractions as described above) for a class of systems, namely switched second order homogeneous systems with binary sensors. These finite approximations were constructed to approximately match the input/output behavior of the original model in a specific sense, quantified in terms of a ‘gain’ bound on an appropriately defined approximation error. The finite models and the corresponding error bounds were then used to systematically and efficiently synthesize certified-by-design stabilizing controllers, obtained by solving discrete min-max problems.

In this paper, we take a step back to basics, by proposing a control oriented *notion of input/output finite state approximation* for discrete-time plants that are allowed to interact with their controllers via fixed finite alphabets, of which the demonstration in [18] is a particular instance. While the proposed notion is inspired from ideas in robust control theory and behavioral systems theory, the class of problems considered here poses unique challenges for two main reasons:

- 1) The lack of algebraic structure: In general, the input and output signals of the system are assumed to take their values in arbitrary sets of symbols.

D. C. Tarraf is with the Electrical and Computer Engineering Department, Johns Hopkins University, Baltimore, MD (dtarraf@jhu.edu).

- 2) The need to approximate both the model and the performance objective, while simultaneously quantifying the approximation error in a useful manner.

The notion of finite state approximation proposed in this paper is distinct from both the notion of approximation by qualitative models and the notion of finite abstraction (simulation/bisimulation). It is potentially better suited to the task of control synthesis for certain problems as it leads to a streamlined controller design procedure, though further investigation is necessary to clearly denunciate these problems, as will be discussed in the paper. Another encouraging development is that it does not preclude unstable systems, as demonstrated in [18].

This paper is organized as follows: We begin in Section II by reviewing some basic concepts that will be used in our development. We present the proposed notion of approximation in Section III and explain its significance in the context of control synthesis. We briefly compare and contrast the proposed notion to two existing notions, namely qualitative models and finite abstractions, in Section IV. We conclude in Section V with directions for future work.

A word on notation: \mathbb{Z}_+ and \mathbb{R}_+ denote the set of non-negative integers and non-negative reals, respectively. Given a set \mathcal{A} , $\mathcal{A}^{\mathbb{Z}_+}$ is the set of all infinite sequences over \mathcal{A} . An element of \mathcal{A} is denoted by a while an element of $\mathcal{A}^{\mathbb{Z}_+}$ is denoted by boldface \mathbf{a} . For $\mathbf{a} \in \mathcal{A}^{\mathbb{Z}_+}$, $a(i)$ denotes the i^{th} component of \mathbf{a} .

II. PRELIMINARIES

We begin by reviewing basic concepts that will be useful in our development: Readers are referred to [21] for a more detailed treatment of some of the definitions given in this section.

A discrete-time signal is understood to be an infinite sequence over some prescribed set. A discrete-time system is a process characterized by its feasible signals set, which is simply the set of ordered pairs of all the signals that can be applied as an input to this process, and all the output signals that can be potentially exhibited by the process in response to each of the input signals.

Definition 1: A discrete-time system S is a set of pairs of signals, $S \subset \mathcal{U}^{\mathbb{Z}_+} \times \mathcal{Y}^{\mathbb{Z}_+}$, where \mathcal{U} and \mathcal{Y} are given alphabet sets.

In this setting, system properties of interest are described as ‘integral’ constraints on the feasible signals set.

Definition 2: [21] Consider a system $S \subset \mathcal{U}^{\mathbb{Z}_+} \times \mathcal{Y}^{\mathbb{Z}_+}$ and let $\rho : \mathcal{U} \rightarrow \mathbb{R}$ and $\mu : \mathcal{Y} \rightarrow \mathbb{R}$ be given functions. S is ρ/μ gain stable if there exists a finite non-negative constant γ such that the following inequality is satisfied for all (\mathbf{u}, \mathbf{y}) in S :

$$\inf_{T \geq 0} \sum_{t=0}^T \gamma \rho(u(t)) - \mu(y(t)) > -\infty. \quad (1)$$

In particular, when ρ and μ are non-negative (and not identically zero), a notion of gain can be defined. Let $\rho :$

$\mathcal{U} \rightarrow \mathbb{R}_+$ and $\mu : \mathcal{Y} \rightarrow \mathbb{R}_+$ be given non-negative functions. The ρ/μ gain of S is the infimum of γ such that (1) is satisfied.

In this paper, we are specifically interested in ‘systems over finite alphabets’, essentially discrete-time systems (plants) that are allowed to interact with other systems (controllers) through fixed discrete alphabets.

Definition 3: A system over finite alphabets S is a discrete-time system $S \subset (\mathcal{U} \times \mathcal{R})^{\mathbb{Z}_+} \times (\mathcal{Y} \times \mathcal{V})^{\mathbb{Z}_+}$ whose alphabet sets \mathcal{U} and \mathcal{Y} are finite.

In this setting, S is understood to represent a plant, $\mathbf{r} \in \mathcal{R}^{\mathbb{Z}_+}$ represents an exogenous input to the plant, $\mathbf{u} \in \mathcal{U}^{\mathbb{Z}_+}$ represents the control input to the plant, $\mathbf{v} \in \mathcal{V}^{\mathbb{Z}_+}$ represents the performance output of the plant, and $\mathbf{y} \in \mathcal{Y}^{\mathbb{Z}_+}$ represents the sensor output of the plant. Systems over finite alphabets can thus be simply thought of as plants with finite-valued sensors and actuators. The internal dynamics of the plant may be analog, discrete or hybrid.

Deterministic finite state machines will be used to approximate systems over finite alphabets.

Definition 4: A deterministic finite state machine is a discrete-time system $S \subset \mathcal{U}^{\mathbb{Z}_+} \times \mathcal{Y}^{\mathbb{Z}_+}$ with finite alphabet sets \mathcal{U} and \mathcal{Y} whose feasible input and output signals are related by

$$\begin{aligned} q(t+1) &= f(q(t), u(t)) \\ y(t) &= g(q(t), u(t)) \end{aligned}$$

where $t \in \mathbb{Z}_+$, $q(t) \in \mathcal{Q}$, $u(t) \in \mathcal{U}$, $y(t) \in \mathcal{Y}$, for some finite set \mathcal{Q} , functions $f : \mathcal{Q} \times \mathcal{U} \rightarrow \mathcal{Q}$ and $g : \mathcal{Q} \times \mathcal{U} \rightarrow \mathcal{Y}$.

In this setting, it is understood that finite set \mathcal{Q} represents the set of states of the system, function f is the state transition function, and function g is the output function in the traditional state-space sense.

III. INPUT/OUTPUT APPROXIMATION

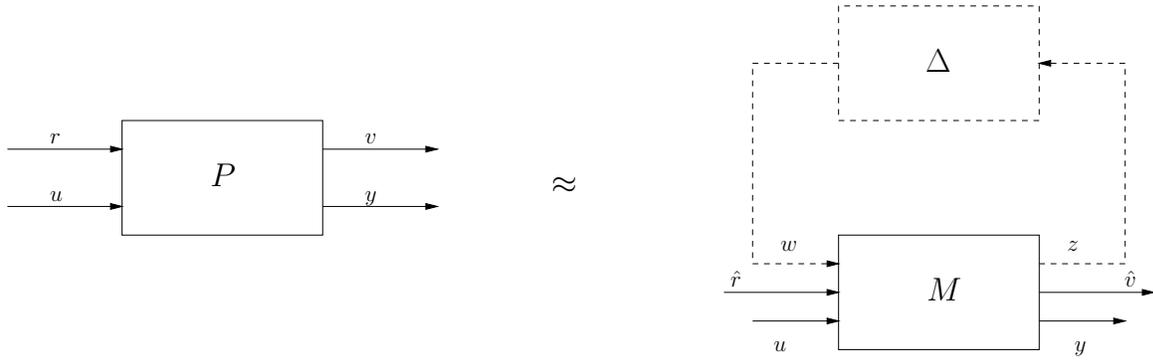
In this section, we propose a control-oriented notion of finite state approximation for systems over finite alphabets whose performance objective is described by a ρ/μ gain stability condition. We formally define this notion in Section III-A, and explain its relevance to the problem of control synthesis in Section III-B. In Section III-C, we revisit the traditional LTI model order reduction problem in the language of the proposed notion of approximation defined in Section III-A for additional insight.

A. Proposed Notion of Approximation

Consider a system over finite alphabets

$$P \subset (\mathcal{U} \times \mathcal{R})^{\mathbb{Z}_+} \times (\mathcal{Y} \times \mathcal{V})^{\mathbb{Z}_+}$$

with exogenous and control inputs r and u respectively, and performance and measurement outputs v and y , respectively, as in Figure 1. It is understood here that alphabet sets \mathcal{U} and \mathcal{Y} are finite, while alphabet sets \mathcal{R} and \mathcal{V} are arbitrary,


 Fig. 1. A finite state approximation of P

meaning they can be finite, countably infinite or simply infinite.

Assume that the purpose of deriving a finite state approximation M of P is to simplify the process of synthesizing a controller $K \subset \mathcal{Y}^{\mathbb{Z}^+} \times \mathcal{U}^{\mathbb{Z}^+}$ such that the closed loop system $(P, K) \subset \mathcal{R}^{\mathbb{Z}^+} \times \mathcal{V}^{\mathbb{Z}^+}$ satisfies a performance objective described by a ρ/μ gain stability condition for some given functions $\rho : \mathcal{R} \rightarrow \mathbb{R}$ and $\mu : \mathcal{V} \rightarrow \mathbb{R}$ (and for $\gamma = 1$),

$$\inf_{T \geq 0} \sum_{t=0}^T \rho(r(t)) - \mu(v(t)) > -\infty \quad (2)$$

Definition 5: A deterministic finite state machine

$$M \subset (\mathcal{U} \times \hat{\mathcal{R}} \times \mathcal{W})^{\mathbb{Z}^+} \times (\mathcal{Y} \times \hat{\mathcal{V}} \times \mathcal{Z})^{\mathbb{Z}^+}$$

with $\hat{\mathcal{R}} \subset \mathcal{R}$ and $\hat{\mathcal{V}} \subset \mathcal{V}$ is a ρ/μ **approximation** of P (plant as described above) if there exists a system $\Delta \subset \mathcal{Z}^{\mathbb{Z}^+} \times \mathcal{W}^{\mathbb{Z}^+}$ and functions $\rho_o : \hat{\mathcal{R}} \rightarrow \mathbb{R}$ and $\mu_o : \hat{\mathcal{V}} \rightarrow \mathbb{R}$ such that

1) There exists a surjective map $\psi : P \rightarrow S$, where

$$S \subset (\mathcal{U} \times \hat{\mathcal{R}})^{\mathbb{Z}^+} \times (\mathcal{Y} \times \hat{\mathcal{V}})^{\mathbb{Z}^+}$$

is the feedback interconnection of M and Δ , as shown in Figure 1, satisfying:

a) $\psi^{-1}(S_{\mathbf{u}, \mathbf{y}}) \supseteq P_{\mathbf{u}, \mathbf{y}}$, where:

$$S_{\mathbf{u}, \mathbf{y}} = \{((\mathbf{u}, \hat{\mathbf{r}}), (\mathbf{y}, \hat{\mathbf{v}})) \in S\}$$

$$P_{\mathbf{u}, \mathbf{y}} = \{((\mathbf{u}, \mathbf{r}), (\mathbf{y}, \mathbf{v})) \in P\}$$

b) If $((\mathbf{u}, \hat{\mathbf{r}}), (\mathbf{y}, \hat{\mathbf{v}})) \in S$ satisfies

$$\inf_{T \geq 0} \sum_{t=0}^T \rho_o(\hat{r}(t)) - \mu_o(\hat{v}(t)) > -\infty \quad (3)$$

then every $((\mathbf{u}, \mathbf{r}), (\mathbf{y}, \mathbf{v}))$ in $\psi^{-1}(\{((\mathbf{u}, \hat{\mathbf{r}}), (\mathbf{y}, \hat{\mathbf{v}}))\})$ satisfies:

$$\inf_{T \geq 0} \sum_{t=0}^T \rho(r(t)) - \mu(v(t)) > -\infty \quad (2)$$

2) Δ is ρ_Δ/μ_Δ gain stable for some non-zero functions $\rho_\Delta : \mathcal{Z}^{\mathbb{Z}^+} \rightarrow \mathbb{R}_+$ and $\mu_\Delta : \mathcal{W}^{\mathbb{Z}^+} \rightarrow \mathbb{R}_+$.

B. Significance of the Proposed Notion

Ultimately, the goal of the approximation is to enable systematic synthesis of a controller $K \subset \mathcal{Y}^{\mathbb{Z}^+} \times \mathcal{U}^{\mathbb{Z}^+}$ such that the closed loop system (P, K) satisfies the performance objective in (2). Defining

$$\bar{P} = \{((\mathbf{u}, \mathbf{r}), (\mathbf{y}, \mathbf{v})) \in P | (\mathbf{y}, \mathbf{u}) \in K\},$$

note that the closed loop system (P, K) is simply the projection of \bar{P} along the second and fourth components. Thus, this controller design problem can be thought of as finding a set $K \subset \mathcal{Y}^{\mathbb{Z}^+} \times \mathcal{U}^{\mathbb{Z}^+}$ such that all the elements of the corresponding set \bar{P} satisfy (2).

The definition of system approximation by deterministic finite state machines proposed in Section III-A is compatible with the goal of robust controller synthesis. Indeed, let M be a ρ/μ approximation of P as in Definition 5. If a K is found such that every element of \bar{S} , defined as:

$$\bar{S} = \{((\mathbf{u}, \hat{\mathbf{r}}), (\mathbf{y}, \hat{\mathbf{v}})) \in S | (\mathbf{y}, \mathbf{u}) \in K\}$$

satisfies the auxiliary performance objective (3), then it follows from condition 1(b) in the definition that every element of $\psi^{-1}(\bar{S})$ satisfies performance objective (2). Thus, in order to ensure that the closed loop system (P, K) also satisfies this performance objective, it is sufficient to ensure that $\bar{P} \subseteq \psi^{-1}(\bar{S})$, which is guaranteed by condition 1(a) in the definition (see Figure 2 for a pictorial illustration of this argument).

Finding a controller K such that all the elements of \bar{S} satisfy (3) is a difficult problem in general, since Δ is an arbitrarily complex system in general. A simpler problem can be posed by characterizing the approximation error Δ in terms of an appropriate gain stability property (ρ_Δ/μ_Δ gain stability with gain bound γ), and then designing a controller K for the nominal DFM model M that is robust to all admissible uncertainties; in other words, K is designed such that the interconnection of M , Δ and K satisfies the auxiliary

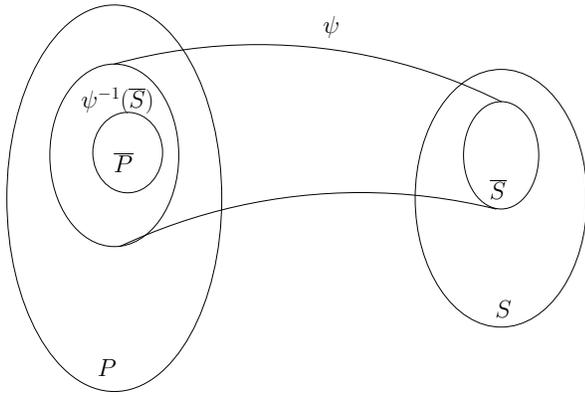


Fig. 2. Approximation for control synthesis: finite alphabet considerations

performance objective (3), for any Δ in the class $\underline{\Delta}$:

$$\begin{aligned} \underline{\Delta} &= \{ \Delta \subset \mathcal{Z}^{\mathbb{Z}^+} \times \mathcal{W}^{\mathbb{Z}^+} \mid \\ &\inf_{T \geq 0} \sum_{t=0}^T \gamma \rho_{\Delta}(z(t)) - \mu_{\Delta}(w(t)) > -\infty \\ &\text{holds } \forall (z, w) \in \Delta \} \end{aligned}$$

Note that this synthesis problem can be elegantly formulated with the help of the ‘Small Gain Theorem’ proposed in [21], and is readily solvable leading to a certified controller K as demonstrated in [18].

In particular, given a choice of functions ρ_{Δ} and μ_{Δ} , let $\underline{\Delta}_i$ be the uncertainty class associated with gain bound γ_i on Δ , and let $S_{\underline{\Delta}_i}$ be the family of systems consisting of the interconnection of M and every $\Delta \in \underline{\Delta}_i$. That is, $S_{\underline{\Delta}_i}$ is the set of all $((\mathbf{u}, \hat{\mathbf{r}}), (\mathbf{y}, \hat{\mathbf{v}})) \in (\mathcal{U} \times \hat{\mathcal{R}})^{\mathbb{Z}^+} \times (\mathcal{Y} \times \hat{\mathcal{V}})^{\mathbb{Z}^+}$ for which there exists \mathbf{w} and \mathbf{z} satisfying:

$$\inf_{T \geq 0} \sum_{t=0}^T \gamma_i \rho_{\Delta}(z(t)) - \mu_{\Delta}(w(t)) > -\infty$$

and such that $((\mathbf{u}, \hat{\mathbf{r}}, \mathbf{w}), (\mathbf{y}, \hat{\mathbf{v}}, \mathbf{z})) \in M$. Since by definition $S \subseteq S_{\underline{\Delta}_i}$, if a controller K is found such that every element of $\bar{S}_{\underline{\Delta}_i}$, defined as:

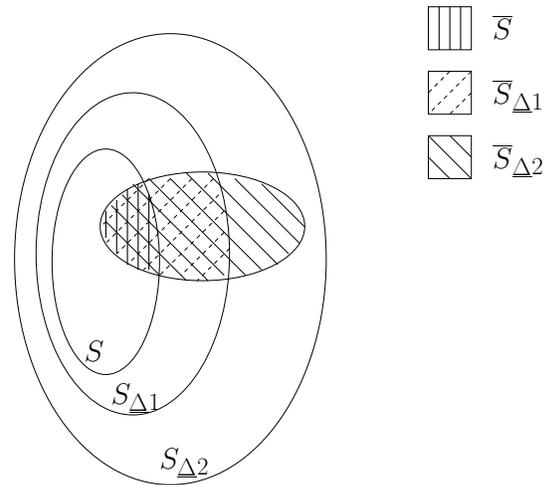
$$\bar{S}_{\underline{\Delta}_i} = \{ ((\mathbf{u}, \hat{\mathbf{r}}), (\mathbf{y}, \hat{\mathbf{v}})) \in S_{\underline{\Delta}_i} \mid (\mathbf{y}, \mathbf{u}) \in K \}$$

satisfies the auxiliary performance objective (3), then so does every element of \bar{S} .

The numerical value of the gain γ , while not directly indicative of the quality of approximation¹, should reflect the trend that quality of approximation improves as γ decreases. The requirement in condition 2 of Definition 5 that ρ_{Δ} and μ_{Δ} are non-negative and are not identically zero allows for a well-defined notion of gain. For two gain bounds $\gamma_1 < \gamma_2$, it follows that (Figure 3):

$$S \subseteq S_{\underline{\Delta}_1} \subseteq S_{\underline{\Delta}_2}$$

¹As is the case in the classical robust control framework, whether a given value of γ is small enough for controller synthesis is not known until an attempt is made to synthesize a controller.


Fig. 3. Approximation for control synthesis: description of Δ

and hence for a given controller $K \subset \mathcal{Y}^{\mathbb{Z}^+} \times \mathcal{U}^{\mathbb{Z}^+}$, we have:

$$\bar{S} \subseteq \bar{S}_{\underline{\Delta}_1} \subseteq \bar{S}_{\underline{\Delta}_2}$$

While the robust control synthesis problem is not discussed here (interested readers are referred to [18] for an exposition of the approach), it should be intuitively clear that the difficulty of finding a controller K to meet the auxiliary performance objective (3) increases as the set of feasible signals of the system considered gets larger, and hence as gain bound γ gets larger.

C. Another Look at the LTI Approximation Problem

In order to get better intuition for the proposed notion of approximation, it may be helpful reinterpret the LTI model order reduction problem in the language of the proposed notion. In this case, P is a stable LTI system of order n , M is a stable LTI system of order $m < n$, $\hat{\mathcal{R}} = \mathcal{R} = \mathcal{U} = \mathcal{Z}$, $\hat{\mathcal{V}} = \mathcal{V} = \mathcal{Y} = \mathcal{W}$, Δ is a stable LTI system given by $\Delta = P - M$, and $S = P$. Thus ψ is simply the identity map, and the relevant picture is the familiar illustration in Figure 3.

It should be intuitively clear that ψ thus captures the necessity, in general, to approximate the *performance objective* in addition to the *plant* for the class of problems considered in this paper, since the original plant and the approximate model do not share the same input and output alphabet sets. The exception is the case where the original plant P is itself a finite state machine.

This simple exercise, in addition to further elucidating the proposed notion of approximation, also clearly suggests the potential for improving the proposed notion by finding a meaningful way of quantifying the difference between \bar{P} and $\psi^{-1}(\bar{S})$, thus quantifying the quality of approximation of the performance objective, in a manner consistent with the goal of robust control synthesis. This will be the subject of future research.

IV. PROPOSED NOTION IN THE CONTEXT OF EXISTING NOTIONS

In this section, we briefly outline the features of the proposed approach in relation to two of the more widely studied approaches to finite state approximation, namely (simulation and/or bisimulation) ‘abstractions’ and ‘qualitative models’.

All three approaches share a similarity in that they enable certified-by-design controller synthesis. In other words, if a sufficiently close finite state approximate model is constructed for a given plant and synthesis is successful, the resulting controller will be guaranteed to satisfy the desired performance criteria, thus bypassing the need for extensive testing and simulation to verify the behavior of the actual closed loop system.

However, owing to the fundamental differences in the formulation of the finite state approximation problem, and consequently in the construction and properties of the finite state approximate models, the process by which the certified-by-design controller is synthesized, and consequently the complexity of the synthesis procedure, is significantly different in the three approaches.

A. Finite Abstractions

For the purpose of this discussion, we restrict our attention to simulation abstractions, though the main observation we will make in this context readily extends to approximate simulation abstractions [5], as well as exact and approximate bisimulation abstractions. The concept of a simulation abstraction originated in computer science where the emphasis is on *verification* of systems (concurrent processes). The verification process can be thought of as checking that all possible output signals of the system satisfy some specifications or constraints. Intuitively, the goal of simulation abstraction is to construct a finite state system that can *exactly match every output signal* of the original system for *some* choice of input (generally different from the corresponding input of the original system). It then follows that if an infinite state system can be abstracted to a finite state one, the verification process is simplified, at the expense of introduced conservatism.

Repeating the above description of simulation abstractions in a slightly more rigorous manner (adapted from [17]), consider the widely employed notation in which a discrete-time dynamical system S is defined to be a sextuple

$$(X, I, U, \rightarrow, Y, H)$$

where X is a set of states; $I \subset X$ is a set of initial states; U and Y are input and output alphabets, respectively; $H : X \rightarrow Y$ is an output map; and $\rightarrow \subset X \times U \times X$ is a transition relation defining the system dynamics. In this setting, a system $S_2 = (X_2, I_2, U_2, \rightarrow_2, Y_2, H_2)$ is said to ‘simulate’ another system $S_1 = (X_1, I_1, U_1, \rightarrow_1, Y_1, H_1)$ if $Y_1 = Y_2$ and there exists a relation $\mathcal{R} \subset X_1 \times X_2$ (referred to as a ‘simulation’ relation) satisfying three properties:

- 1) For every $x_1 \in I_1$, there exists an $x_2 \in I_2$ such that $x_1 \mathcal{R} x_2$.

- 2) $x_1 \mathcal{R} x_2 \Rightarrow H_1(x_1) = H_2(x_2)$.

- 3) If $x_1 \mathcal{R} x_2$ and there exists an input $u_1 \in U_1$ such that $(x_1, u_1, x_1^+) \in \rightarrow_1$, then there exists an $x_2^+ \in X_2$ and a $u_2 \in U_2$ such that $(x_2, u_2, x_2^+) \in \rightarrow_2$ and $x_1^+ \mathcal{R} x_2^+$.

In particular, the setups of interest here are ones in which state set X_2 is finite while state set X_1 is infinite (continuous or hybrid).

There are two critical differences between the notion of approximation proposed in this paper and the notion of simulation abstraction reviewed briefly above. Indeed, the proposed approach aims to find a finite state model that *approximates* the *input/output* feasible signals of the plant. In contrast, simulation abstractions seek to find a finite state model whose feasible output signals *contain*² the feasible *output* signals of the plant. While this difference may be of little consequence in analysis and verification problems, it greatly impacts the process by which a controller is synthesized. Specifically, in the former case, the problem of controller synthesis for the finite state model reduces to solving a discrete min-max problem. Moreover, by construction, the solution to this optimization problem immediately yields a certified-by-design finite state controller for the original plant. In contrast, in the latter case, the process of synthesizing a controller is a two step procedure. First, a controller is synthesized for the finite abstraction using supervisory control techniques. This finite state controller cannot be directly plugged back into the original system; rather, it needs to be subsequently refined to a hybrid controller that can then be used to close the loop around the original plant. The differences in the process of synthesizing controllers for the finite state models are due to the ‘approximation’ versus ‘containment’ (or more generally ‘approximate containment’) aspect. The need to further refine the controller in the latter case arises due to the fact that only the outputs are considered, rather than the input/output map: As a result, the inputs that result in identical outputs in the original plant and its abstraction are not identical in general.

These observations suggest that the proposed notion of approximation may be better suited for the problem of control synthesis as it leads to a more streamlined synthesis procedure. However, this is not an absolute statement, as a fundamental assumption in the proposed notion is that the desired performance objective can be adequately captured by a gain stability condition. It is unclear at this point whether performance objectives involving temporal logic specifications, which have been demonstrated in the context of simulation/bisimulation abstractions, can be handled by the proposed approach. This will be one direction of future work.

B. Qualitative Models

The literature on qualitative models [8], [9], [14], [10] describes an older direction of research into the problem of constructing finite approximations of hybrid systems with

²Exactly, in this case, or approximately, in the case of approximate simulation abstractions. In the case of (approximate) bisimulation abstractions, the goal is to (approximately) match the output sets

input and output quantizers. In this setting, ‘qualitative models’ are understood to be finite state automata with non-deterministic state transitions, whose feasible output set or feasible input/output set contains that of the original system. As such, this notion of finite-state approximation clearly differs from both the notion proposed in this paper and the notion of abstraction, in that the *class of nominal models is different*. A certified-by-design supervisory controller for the original plant can be designed using standard supervisory control techniques applied to the qualitative model. However, the exact implications of the class of nominal models on the complexity of the controller synthesis problem are unclear, and will be the subject of further investigation.

V. CONCLUSIONS AND FUTURE WORKS

In this paper, we considered the problem of approximating plants with discrete sensors and actuators (termed ‘systems over finite alphabets’) by deterministic finite memory systems for the purpose of certified-by-design controller synthesis. We proposed a new, control-oriented notion of input/output approximation for these systems, and we briefly discussed some of its key features relative to those of two existing notions of finite state approximation and abstraction, from the standpoint of control synthesis.

Future work will explore the possibility of further refining this notion of approximation, specifically by quantifying the error introduced due to the approximation of the performance objective in a manner consistent with robust control synthesis. Another direction of future work will focus on further exploring this notion of approximation for problems involving broader classes of performance objectives, including linear temporal logic constraints.

VI. ACKNOWLEDGMENTS

The author is indebted to A. Megretski for many stimulating (and often heated!) discussions that inspired this work. The author thanks M. A. Dahleh for feedback on early versions of some of the ideas presented here.

REFERENCES

[1] V. M. Adamjan, D. Z. Arov, and M. G. Krein. Analytic properties of Schmidt pairs for a Hankel operator and the generalized Schur-Takagi problem. *Math. USSR Sbornik*, 15(1):31–73, 1971.

[2] R. Alur, T. Henzinger, G. Lafferriere, and G. Pappas. Discrete abstractions of hybrid systems. *Proceedings of the IEEE*, 88(2):971–984, 2000.

[3] A. C. Antoulas. *Approximation of Large Scale Dynamical Systems*. Cambridge University Press, 2005.

[4] C. Belta, A. Bicchi, M. Egerstedt, E. Frazzoli, E. Klavins, and G. J. Pappas. Symbolic planning and control of robot motion: State of the art and grand challenges. *IEEE Robotics and Automation Magazine*, 14(1):61–70, March 2007.

[5] A. Girard, A. A. Julius, and G. J. Pappas. Approximate simulation relations for hybrid systems. *Discrete Event Dynamic Systems*, 18:163–179, 2008.

[6] A. Girard and G. Pappas. Approximation metrics for discrete and continuous systems. *IEEE Transactions on Automatic Control*, 52(5):782–798, 2007.

[7] K. Glover. All optimal Hankel norm approximations of linear multi-variable systems and their \mathcal{L}^∞ error bounds. *International Journal of Control*, 39(6):1145–1193, 1984.

[8] J. Lunze. On the stability of qualitative models. *Systems & Control Letters*, 21(2):137–142, 1993.

[9] J. Lunze. Qualitative modeling of linear dynamical systems with quantized state measurements. *Automatica*, 30:417–431, 1994.

[10] J. Lunze and J. Raisch. Discrete models for hybrid systems. In *Modeling, Analysis and Design of Hybrid Systems*, volume 279 of *Lecture Notes in Control and Information Sciences*, pages 67–80. Springer, Berlin/Heidelberg, 2002.

[11] B. Moore. Principal component analysis in linear systems: Controllability, observability and model reduction. *IEEE Transactions on Automatic Control*, AC-26:17–32, 1981.

[12] G. Obinata and B. D. O. Anderson. *Model Reduction for Control System Design*. Springer, 2001.

[13] L. Pernebo and L. Silverman. Model reduction via balanced state space representation. *IEEE Transactions on Automatic Control*, AC-27:382–387, 1982.

[14] J. Raisch and S. D. O’Young. Discrete approximation and supervisory control of continuous systems. *IEEE Transactions on Automatic Control*, 43(4):569–573, April 1998.

[15] P. J. Ramadge and W. M. Wonham. Supervisory control of a class of discrete event processes. *SIAM Journal on Control and Optimization*, 25(1):206–230, 1987.

[16] P. Tabuada. An approximate simulation approach to symbolic control. *IEEE Transactions on Automatic Control*, 53(6):1406–1418, July 2008.

[17] P. Tabuada. *Verification and Control of Hybrid Systems: A Symbolic Approach*. Springer, 2009.

[18] D. C. Tarraf, A. Megretski, and M. A. Dahleh. Finite approximations of switched homogeneous systems for controller synthesis. Submitted for journal publication, preprint available at <http://arxiv.org/abs/0903.5535>.

[19] D. C. Tarraf, A. Megretski, and M. A. Dahleh. Finite automata approximations with error bounds for systems with quantized actuation and measurement: A case study. In *Proc. CDC Decision and Control 43rd IEEE Conference on*, volume 2, pages 1436–1441, December 2004.

[20] D. C. Tarraf, A. Megretski, and M. A. Dahleh. Finite state controllers for stabilizing switched systems with binary sensors. In A. Bemporad, A. Bicchi, and G. Buttazzo, editors, *Hybrid Systems: Computation and Control*, volume 4416 of *Lecture Notes in Computer Science*, pages 543–556. Springer, April 2007.

[21] D. C. Tarraf, A. Megretski, and M. A. Dahleh. A framework for robust stability for systems over finite alphabets. *IEEE Transactions on Automatic Control*, 53(5):1133–1146, June 2008.