ARTICLE IN PRESS

Performance Evaluation ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Performance Evaluation

journal homepage: www.elsevier.com/locate/peva

Certified policy synthesis for general Markov decision processes: An application in building automation systems

Sofie Haesaert [a], Nathalie Cauchi [b,*], Alessandro Abate [b]

[a] Department of Electrical Engineering, Technische Universiteit Eindhoven, Eindhoven, Netherlands

[b] Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford, United Kingdom

HIGHLIGHTS

- Extension of theory of policy synthesis and verification for two rooms.
- Introduction of library of thermal model with different levels of abstractions.
- Control policy for abstract model used to design strategy for concrete model.
- New software toolbox supporting the computation of approximate relations.

ARTICLE INFO

Article history:
Available online xxxx

Keywords:
Verification
Synthesis
General Markov decision processes
Safety
Building automation systems
Temperature control

ABSTRACT

In this paper, we present an industrial application of new approximate similarity relations for Markov models, and show that they are key for the synthesis of control strategies. Typically, modern engineering systems are modelled using complex and high-order models which make the correct-by-design controller construction computationally hard. Using the new approximate similarity relations, this complexity is reduced and we provide certificates on the performance of the synthesised policies. The application deals with stochastic models for the thermal dynamics in a "smart building" setup: such building automation system set-up can be described by discrete-time Markov decision processes evolving over an uncountable state space and endowed with an output quantifying the room temperature. The new similarity relations draw a quantitative connection between different levels of model abstraction, and allow to quantitatively refine over complex models control strategies synthesised on simpler ones. The new relations, underpinned by the use of metrics, allow in particular for a useful trade-off between deviations over probability distributions on states and distances between model outputs. We develop a software toolbox supporting the application and the computational implementation of these new relations.

1. Introduction

Buildings consume more than 40% of the energy in Europe [1] and in the United States [2]. The design of their heating, ventilation and air-conditioning (HVAC) systems has become a focal point not only in the buildings and HVAC community [3,4], but also in the control community [5–7].

* Corresponding author.
E-mail address: nathalie.cauchi@cs.ox.ac.uk (N. Cauchi).

https://doi.org/10.1016/j.peva.2017.09.005
0166-5316/

The goal of modern building automation systems (BAS) is to control the climate in an energy-efficient manner, while allowing for diverse and complex functionality imposed by the users. For this, a myriad of possible control strategies have been proposed, the most modern of which hinge on accurate mathematical models of the thermal dynamics [8–10]. The quantification of building dynamics leads to models evolving stochastically over continuous spaces. Stochastic dynamics are needed to capture the effect of external disturbances, such as weather fluctuations and occupancy changes, both of which can hardly be described with deterministic models. While accurate models are needed, their inherent level of complexity, related to stochasticity and dimensionality, limits the type of controllers that can be designed. Recent literature [11–14] has seen control designs targeted to optimally maintain a comfortable inside climate. To satisfy the goal for future modern BAS, such control strategies should not only be optimal but also guarantee user-imposed diverse and complex requirements and specifications. The use of temporal logics allows for a formalised expression of such diverse functional requirements, and enables for the automated verification and synthesis of control strategies [15–17]. An interesting formal approach is the approximation of the original (concrete) models by simpler (abstract) models that are prone to be analysed or algorithmically verified. Further, if a strategy can be synthesised on these abstractions, it should then be certifiably refined over the concrete model. Building on the similarity between the abstract and concrete models, key formal guarantees on the latter can be raised. In this work, we employ notions of approximate probabilistic similarity relations, originally introduced in [18]. Underpinned by the use of metrics, these relations allow in particular for a useful trade-off between deviations over probability distributions on states, and distances between model outputs. Other relations, also targeting general, uncountable-state spaces, such as those established via martingale theory [19,20] require stringent stability criteria, or alternatively enforce structural abstractions [21] based on state-space gridding [22].

Within the control community, the computational complexity of a model is tackled by abstracting it to a lower dimensional one. Numerical model-order reduction (MOR) techniques reduce the dimensionality of models with minimal loss in input–output accuracy [9,23–26]. A number of different MOR techniques are in use [27]: Gramian-based MOR involve balancing the observability and controllability Gramians via truncations [27]; Schur-based MOR techniques compute a reduced-order system through a set of projections [28]; alternative MOR techniques reduce model orders with respect to the Hankel norm [29]. These numerical model-order reduction techniques are oblivious of the underlying physics associated to the dynamics and, as a result, the obtained low-order models often no longer have a physical interpretation. This can be of a disadvantage when faults arise and an analysis on the underlying model is needed to detect and interpret the physical source causing faults. In this paper, we perform model reduction based on assumptions made on the underlying physical quantities. We present a methodology to provide guarantees for control strategy synthesis over a library of thermal models. Models could in general be obtained via MOR techniques, however key to this library of models is that they maintain correspondences via the underlying physical variables. The provided guarantees are computed by a new toolbox that implements the synthesised control strategy via (ε, δ)-approximate simulation relations. The new relations serve as a performance metric between models of different orders of complexity and allow the user to choose the most suitable model abstraction for achieving the required level of performance by the synthesised and refined control policies. The toolbox is configured via a graphical user interface (GUI) that simplifies the certified policy synthesis process. Consequently, the toolbox facilitates the process of designing certified control strategies for BAS.

The used notions of approximate similarity relations for control refinement were first introduced in [18]. The theoretical background is given more extensively in [30], and a case study using FAUST^2 [17]. In this paper, we move beyond the purely theoretical treatment of these relations. For building management, we show how these relations can be used. More precisely, we explain the computational procedures that can be used to compute the accuracy of these relations for given models.

The article has the following structure: Section 2 introduces the problem statement and gives a demonstration of the developed policy synthesis toolbox. Section 3 introduces the library of models that capture the thermal dynamics of building zones. This is followed by the theory behind policy synthesis via approximate simulation relations in Section 4. Section 5 presents how certified control refinement based on the approximate simulation relations can be performed. The case study, which highlights the advantages and applicability of the devised framework, is presented in Section 6. This is then followed by the computational implementation of the toolbox in Section 7.

2. Certified control synthesis

With as goal the future industrial application of certified control synthesis for complex engineering systems, we want to develop a first application-oriented framework to enable the design and evaluation of this methodology. We target the application in BAS and single out an exemplar case study in room heating. We want to provide quantified performance of the synthesised policies with certificates.

Often a single system can be described by several models of varying complexity and accuracy. For the BAS case study, we develop such a library of models, of use for advanced control solutions. Given a library of different but related models, we want to reason over different levels of abstraction, while maintaining guarantees. Such guarantees are established based on the computation of (ε, δ)-approximate (bi)-simulation relations, which are capable of refining assertions on synthesised control policies amongst models. The pair (ε, δ) represent the deviation in the output trajectories between complex and abstract models and the differences in the stochastic transition kernels, respectively. In particular, we aim at automating the