

Benchmarks for cyber-physical systems: A modular model library for building automation systems

Nathalie Cauchi* Alessandro Abate*

* *Department of Computer Science, University of Oxford, Oxford, UK*
(e-mail: *name.surname@cs.ox.ac.uk*).

Abstract: Building Automation Systems (BAS) are exemplars of Cyber-Physical Systems (CPS), incorporating digital control architectures over underlying continuous physical processes. We provide a modular model library for BAS drawn from expertise developed on a real BAS setup. The library allows to build models comprising either physical quantities or digital control modules. The structure, operation, and dynamics of the model can be complex, incorporating (i) stochasticity, (ii) non-linearities, (iii) numerous continuous variables or discrete states, (iv) various input and output signals, and (v) a large number of possible discrete configurations. The modular composition of BAS components can generate useful CPS benchmarks. We display this use by means of three realistic case studies, where corresponding models are built and engaged with different analysis goals. The benchmarks, the model library and associated data collected from the BAS setup at the University of Oxford, are kept on-line at <https://github.com/natchi92/BASBenchmarks>

Keywords: cyber-physical systems, building automation systems, thermal modelling, hybrid models, simulation, reachability analysis, probabilistic safety, control synthesis

1. INTRODUCTION

This paper describes a library of models for Building Automation Systems (BAS), which can be employed to create benchmarks for verification, control synthesis, or simulation purposes of Cyber-Physical Systems (CPS). The models are inspired by and built around an experimental setup within the Department of Computer Science at the University of Oxford, which is part of on-going research in collaboration with service engineers and industrial partners in the sector. This library allows to create numerous meaningful models for BAS, which are examples of CPS integrating continuous dynamics and discrete modes.

Interest in BAS, also colloquially known as *smart buildings*, is gaining rapid momentum, in particular as a means for ensuring thermal comfort (Belkhouane et al. [2017]), minimising energy consumption (Soudjani and Abate [2015], Yu et al. [2015]), and ascertaining reliability (Zhang and Hong [2017]). Quantitative models are needed to evaluate system performance, to verify correct behaviour, and to develop specific control algorithms. An overview of the different BAS modelling techniques used in literature is presented in Privara et al. [2013]. Several simulation tools (see Crawley et al. [2008]) have been devised to aide in the development and analysis of models for BAS. Attempting a multi-dimensional characterisation of the broad spectrum of existing BAS models, we can find either deterministic or stochastic ones, low- to high-dimensional ones, with discrete or continuous inputs and states. The choice of a model is an art and a craft (Ma et al. [2012]): one must select simplifying assumptions

that accurately reflect the operational performance of the BAS in specific real-world environments, and introduce uncertainty to represent un-modelled components, unknown parameters or random occupants. We therefore aim to simplify the modelling process such that simulation, verification or strategy synthesis can be carried out seamlessly. Different verification and policy synthesis tools exist in literature (Dragomir et al. [2017], Holub et al. [2016]). They are typically specific to a particular type of model structure and in the case of stochastic or hybrid models are often limited to a small number of continuous variables (Fehnker and Ivančić [2004]). The use of such tools also requires expert knowledge on the specific formalism the tool makes use of.

In order to display the versatility of the library of BAS models, we present three case studies that are built from its components. We focus on modelling temperature dynamics, a key element for ensuring thermal comfort. We employ the three generated models for different analysis goals, comprising simulation, reachability, and control synthesis. The models and the case-studies are kept on-line at the URL above, and at (Cauchi and Abate [2018]). This is to allow their use or modification for different applications and for comparison with other modelling approaches in BAS. The repository also contains real data gathered from the BAS lab at Oxford, which can be employed for further modelling studies. This article has the following structure: Section 2 introduces the BAS modelling framework for CPS. We identify three modelling trade-offs that introduce different complexities on the model dynamics. Based on these trade-offs, we develop and analyse three case studies in Section 3.

Table 1. Indices

Index	Reference	Index	Reference
a	AHU	adj	adjacent zone
adj, out	adjacent exterior zones	b	boiler
d	mixer	$hall$	hallway
$i \in \{1, 2\}$	individual zones	$jn \in \{2, 3, 7\}$	zone walls with no windows
$iw \in \{5, 6\}$	zone walls with windows	$j \in \{jn \cup iw\}$	all zone walls
$l \in \{1, 2\}$	adjacent interior zone	occ	occupants
out	outside	r	radiator
ref	reference	rw	return water
sa	supply air	$solar$	solar energy
sw	supply water	v	collector
w	wall	z	zone
h	water	ar	air

2. BUILDING AUTOMATION SYSTEMS

2.1 BAS: structure and components

BAS models clearly depend on the size and topology of the building (Kim and Katipamula [2017]), and on its climate control setup. In this work, we consider the BAS setup in the Department of Computer Science, at the University of Oxford. A graphical depiction is shown in Figure 1a. The BAS consists of two teaching rooms that are connected to a boiler-heated system. The boiler supplies heat to the heating coil within the AHU and to two radiators. Valves control the rate of water flow in the heating coil and in radiators. The AHU supplies air to the two zones, which are connected back to back, and are adjacent both to the outside and to an interior hall (cf. Figure 1a). The zone air of both rooms can mix with the outside air and exchanges circulating air with the AHU. Return water from the AHU heating coils and radiators is collected and pumped back to the boiler.

Figure 1b presents the Resistor Capacitance (RC) network circuit of the two zones (Haesaert et al. [2017]), which underpins the dynamics for temperature in the zone component - corresponding equations are in Table 3. The heat level in each room is modified by (i) radiative solar energy absorbed through the walls, (ii) occupants, (iii) AHU input supply air, (iv) radiators, and (v) AHU return water. The effect of heat stored in the walls and in rooms is depicted with capacitors, whereas thermal resistance to heat transfer by the walls is depicted by resistor elements.

2.2 BAS: dynamics and configurations

Single components are intended as separate physical structures within the BAS. Their models are built from the underlying physics and are improved via industrial feedback and from existing literature (Ferrari et al. [2017], Haesaert et al. [2017]). We obtain models with a number of unknown parameters: these are estimated and validated using data collected from the BAS setup (Kristensen et al. [2004]). We list indices in Table 1, while all the quantities (variables, parameters, inputs) are listed in Table 2. Table 3 presents all the relations among variables in the model components: algebraic relations define static couplings, whereas differential relations define the dynamics for the corresponding variables. The structure in Figure 1a, the quantities in Table 2, and the variables (with associated dynamics) in Table 3, together allow to construct global models for the BAS setup. We refer to the set of models describing the individual components (cf. Table 3) as a “library of

models”: one can select the individual components from the library, and build different BAS configurations and models.

Models of the BAS setup can be complex, consisting of both algebraic and differential relations, process noise, and possibly numerous continuous variables and discrete modes. Some of the dynamics can be non-linear, especially in view of continuous variables or inputs that are bilinearly coupled (cf. AHU air duct model in Table 3). The number of continuous variables also increase substantially when considering a BAS setup with multiple zones.

We look at the dynamics of the library of BAS components from three different perspectives: (i) presence of stochasticity; (ii) number of continuous variables; and (iii) number of discrete modes. Using the identified sources of complexity, models of different types can be constructed. In order to add a level of flexibility to the modelling framework, we consider each BAS component as a separate module, characterised by input and output elements, and internal variables. We make use of individual modules describing component type, and then connect different modules based on possible physical couplings. Coupling between different modules is also achieved via input-output relationships: e.g., in the zone module we have coupling between two zones through the continuous variable $T_{adj,l}$ corresponding to the adjacent zones, which for the wall separating the two zones (cf. W_7 in Figure 1b) corresponds to the individual zone temperatures of the two zone modules (cf. zone equations in Table 3). Having such a modular structure for the individual components provides an added level of versatility, since we can connect different components to create new models. Modularisation also allows (i) to perform analysis of the whole setup by executing analysis of individual modules and (ii) to extend the library of models by defining new modules that connect to existing modules via their input-output relations (for example, in order to model buildings with larger number of rooms).

2.3 BAS: software description of model library

The library of BAS components comes in the form of MATLAB scripts. Each script represents an individual BAS component. The models are in state-space form and are of two types linear or non-linear depending on the component they represent. They are defined using the symbolic toolbox, and can be described both in discrete and continuous time. We provide the parameters which we estimated from real data gathered from the BAS setup at the University of Oxford, to construct the individual models. Different components can be connected based on their input-output relations by cascading the different symbolic models for each component. Once this is done, the provided scripts allow one to simulate the models and to generate plots for the defined output variables.

3. CASE STUDIES

We set up three case studies and present the trade off between the discussed kinds of complexity. For each of the case studies, we also introduce instances of problems that we solve using these models, and describe the results obtained.

Fig. 1. Building automation system setup

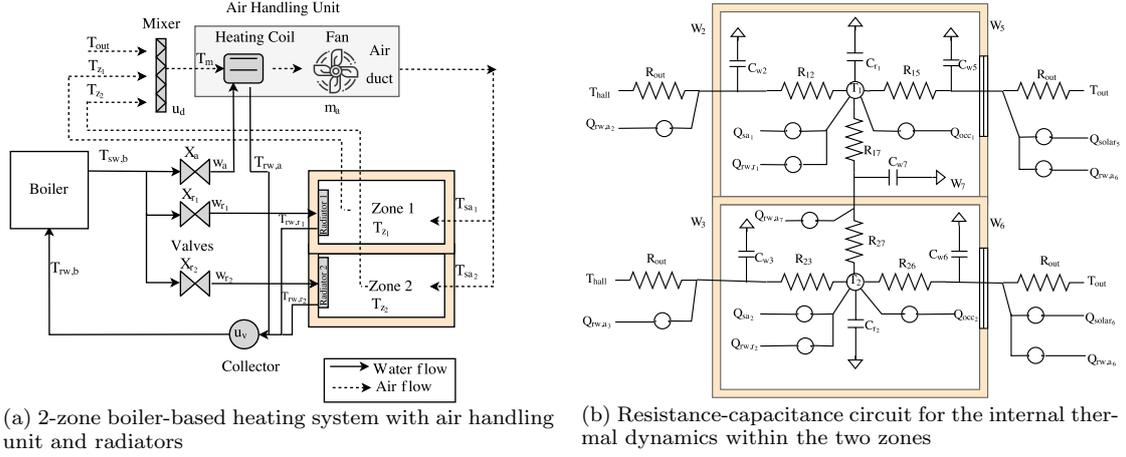


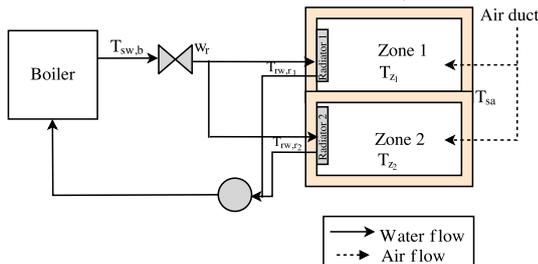
Table 2. List of variables, inputs, and parameters

Symbol	Quantity	Type	Symbol	Quantity	Type
A_i	area of windows of each zone	constant	B_{en}	boiler switch	discrete
C	capacitance	constant	C_{pa}, C_{pw}	specific heat capacity of air and water	constant
$CO_{2,i}$	carbon-dioxide measurements in each zone	input	k_b	steady-state of the boiler	constant
m	mass air flow rate	input	n	number of zones	constant
P_{out}	radiator rated power output	constant input	Q	heat gain	input
R	thermal resistance to heat from walls	constant	T	temperature	state input
u	mixing ratio	input	(UA)	overall transmittance factor of	constant
V	volume of	constant	w	water flow rate	input
w_{max}	maximum water-flow permitted by the valve	constant	X	valve position	input
$\{\alpha, \beta, \mu\}$	de-rating and offset factors	constants	σ	process noise	constant
ρ	density	constant	τ	time constant	constant

Table 3. Dynamics and functional relations among variables of components

Component	Continuous variables	Relation
Boiler	$dT_{sw,b}(t) = \begin{cases} 0 & B_{en}(t) = 0 \\ (\tau_{sw})^{-1} [(-T_{sw,b}(t) + k_b)dt] + \sigma_{sw}dW & B_{en}(t) = 1 \end{cases}$	differential
Valve	$w(t) = (\tau)^{-1} [\exp(\ln(\tau)X(t))w_{max}]$	algebraic
Mixer	$T_d(t) = u_d T_{out}(t) + (1 - u_d) \left(\sum_i T_{z_i}(t) \right) (n)^{-1}$	algebraic
AHU heating coil	$dT_{rw,a}(t) = (C_{pw} \rho_h V_a)^{-1} [(C_{pw} w_a(t)(T_{sw,b}(t) - T_{rw,a}(t)) + (UA)_a(T_d(t) - T_{rw,a}(t)))] dt + \sigma_{rw,a} dW$	differential
AHU air duct	$dT_{sa_i}(t) = (C_a \rho_a V_a)^{-1} [m_a(t) C_{pa}(T_d(t) - T_{sa_i}(t)) + (UA)_a(T_{z_i}(t) - T_{sa_i}(t))] dt + \sigma_{sa_i} dW$	differential
Radiator	$dT_{rw,r_i}(t) = (C_{pw} \rho_h V_{r_i})^{-1} [(C_{pw} w_{r_i}(t)(T_{sw,b}(t) - T_{rw,r_i}(t)) + (UA)_{r_i}(T_{z_i}(t) - T_{rw,r_i}(t))] dt + \sigma_{rw,r_i} dW$	differential
Zone	$dT_{z_i}(t) = (C_{z_i})^{-1} \left[\frac{T_{w_{j_n}}(t) - T_{z_i}(t)}{R_{ij}} + Q_{rw,r_i}(t) + Q_{occi}(t) + Q_{sa_i}(t) \right] dt + \sigma_{z_i} dW$ $dT_{w_{j_n}}(t) = (C_{w_{j_n}})^{-1} \left[\frac{T_{adj,out}(t) - T_{z_i}(t)}{R_{out}} + \sum_l \frac{T_{adj,l}(t) - T_{w_{j_n}}(t)}{R_{lj}} + Q_{rw,a_{j_n}}(t) \right] dt + \sigma_{w_{j_n}} dW$ $dT_{w_{j_w}}(t) = (C_{w_{j_w}})^{-1} \left[\frac{T_{adj,out}(t) - T_{z_i}(t)}{R_{out}} + \sum_l \frac{T_{adj,l}(t) - T_{w_{j_w}}(t)}{R_{lj_w}} + Q_{solar_{j_w}}(t) + Q_{rw,a_{j_w}}(t) \right] dt + \sigma_{w_{j_w}} dW$ $Q_{rw,r_i}(t) = P_{radi} (\alpha_2(T_{rw,r_i}(t) - T_{z_i}(t)) + \alpha_1), Q_{occi}(t) = \mu_i(CO_{2,i}(t)) + \beta_1, Q_{sa_i}(t) = m_a(t) C_{pa}(T_{sa_i}(t) - T_{z_i}(t))$ $Q_{rw,a_j}(t) = \alpha_3(T_{rw,a}(t) - T_{w_j}(t)), Q_{solar_{j_w}}(t) = (\alpha_0 A_i T_{out}(t) + \beta_2)$	differential
Collector	$T_{rw,b}(t) = u_v T_{rw,a}(t) + (1 - u_v) \left(\sum_i T_{rw,r_i}(t) \right) (n)^{-1}$	algebraic

Fig. 2. BAS setup for the first case study



3.1 Two-zone heating setup with deterministic or stochastic dynamics

We consider two zones, each heated by one radiator and with a common supply air, as portrayed in Figure 2. From Table 3, we select two components and corresponding models: the radiator and the zone. We simplify these models with the following assumptions: (i) the wall temperature is constant across the zones and is a fixed value ($T_{w,ss}$); (ii) the boiler is switched ON providing a supply temperature $T_{sw,b,ss}$; (iii) we fix both the mass air flow rate m_a and the radiator water flow rate w_r ; and (iv) we do not include the heat gain from the windows and the AHU heating coils ($T_{w,ss}$) in each zone. We obtain a model

with the four state variables $x^T = [T_{z_1} \ T_{z_2} \ T_{rw,r_1} \ T_{rw,r_2}]^T$ and with a common supply temperature $u = T_{sa}$ as an input. For this setup we further consider three different dynamics: (i) purely deterministic ones; (ii) a deterministic model with additive disturbance; and (iii) a stochastic model. For (i) and (ii), we thus remove the process noise in the template model, while for models (i) and (iii) we do not include the occupancy heat gain. We also discretise the dynamics by a Forward-Euler scheme (for the deterministic models) and a Euler-Maruyama scheme (for the stochastic model Abate et al. [2008]), using a uniform sampling time $\Delta = 15$ minutes, and obtain a set of linear discrete-time models. One should note that the models being considered are not directly fully observable since for all the individual zone temperatures (variables of interest) are the only variables provided as outputs. The dynamics of the deterministic model (i) are described by

$$\mathbf{M}_d : \begin{cases} x[k+1] &= Ax[k] + Bu[k] + Q_d \\ y_d[k] &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} x[k], \end{cases} \quad (1)$$

where again the matrices are taken from the models in Table 3 and Q_d is vector of constant gains. The deterministic model with additive disturbance is

$$\mathbf{M}_{da} : \begin{cases} x[k+1] &= Ax[k] + Bu[k] + F_{da}d_{da}[k] \\ &\quad + Q_{da} \\ y_{da}[k] &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} x[k]. \end{cases} \quad (2)$$

We have extended (1) with additive noise vector d_{da}^T (cf. Q_{occi} in Table 3) representing the different CO_2 levels in each zone. Q_{da} and F_{da} are a properly sized matrices. The stochastic model is expressed by extending (1) to include process noise, as

$$\mathbf{M}_s : \begin{cases} x[k+1] &= Ax[k] + Bu[k] + Q_d + \Sigma W[k] \\ y_s[k] &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} x[k], \end{cases} \quad (3)$$

where Σ encompasses the variances of the process noise for each state. $W = [w_1 \ w_2 \ w_3 \ w_4]^T$ is a vector of independent Gaussian random variables, which are also independent of the initial condition of the process. Details on the individual matrices for all the models are presented in (Cauchi and Abate [2018]).

Reachability analysis For this case study we perform the following verification task: to decide whether traces generated by the models remain within a specified safe set for a given time period. This is achieved by reachability analysis, which takes a probabilistic flavour for the stochastic model. The safe set is described as an interval around the temperature set-point $T_{SP} = 20^\circ C \pm 0.5^\circ C$. We constrain the input u to lie within the set $\{T_{sa} \in \mathbb{R} | 15 \leq T_{sa} \leq 22\}$ for all models and employ a fixed time horizon $N = 6 \times \Delta = 1.5$ hours. We perform reachability analysis of the models M_d and M_s with Axelerator (Cattaruzza et al. [2015]), while we use FAUST² (Soudjani et al. [2015]) to perform probabilistic reachability analysis of \mathbf{M}_s . In order to perform reachability analysis using Axelerator, for each of the models we set the initial condition as $[T_{z_1} \ T_{z_2} \ T_{rw,r_1} \ T_{rw,r_2}]^T = [18 \ 18 \ 35 \ 35]^T$. The reach tube for model \mathbf{M}_d over the whole time horizon is shown in

Figure 3: it encompasses the union of all reachable states over that horizon. The results obtained using Axelerator (cf. Figure 3) are conservative, but can confirm that the model indeed stays in the required safe set for some initial states, but can also exit it. One can note the coupling between the two zones and that zone 1 tends to stay at higher zone temperatures than zone 2. A similar reach tube is obtained for model \mathbf{M}_{da} . We perform probabilistic reachability analysis on model \mathbf{M}_s by defining the same safe set and assuming as input set the interval $[15, 22]$. The resulting adaptive partition of the safe set along with the optimal safety probability for each partition, is depicted in Figure 4. When performing probabilistic reachability analysis using model \mathbf{M}_s , we deduce that the model has a high probability of being within the required safe set, specifically to have $T_{z_1} \in [19.5 \ 20.5]$ and $T_{z_2} \in [19.5 \ 20.5]$.

Fig. 3. First case study: reach tube of M_d over whole time horizon

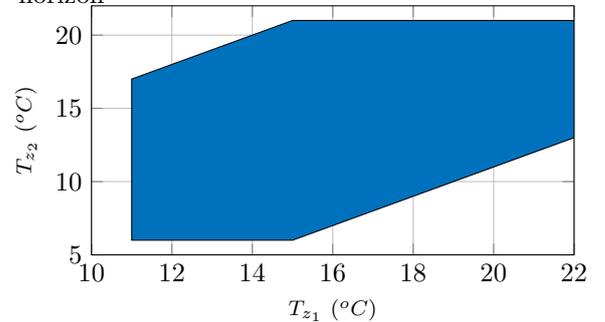
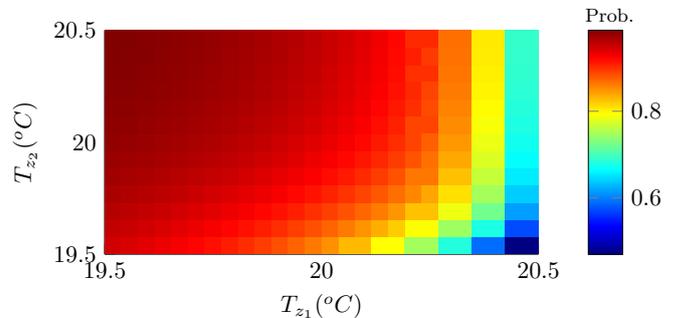


Fig. 4. First case study: partition of the safe set for model M_s , along with optimal safety probability for each partition set



3.2 Two-zone heating setup with a large number of continuous variables

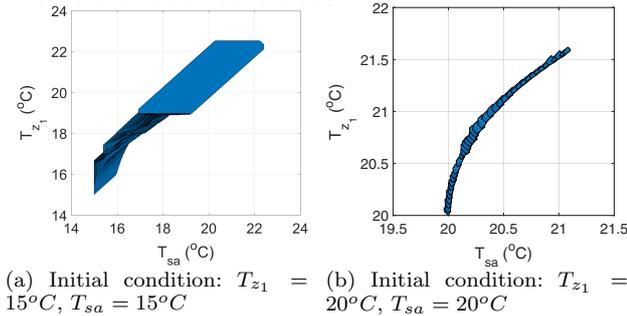
In this second case study we focus on the dynamics of zone components from Table 3 and consider the two zones shown in Figure 1b. We assume that (i) a central fan pumps air into both rooms with a common supply temperature $15^\circ C \leq T_{sa} \leq 30^\circ C$, (ii) the input mass airflow m_a is fixed to $10 \text{ m}^3/\text{hour}$ and (iii) the return water temperature of the AHU heating coils is fixed ($T_{rw,a_{ss}} = 35^\circ C$). As in the previous case study, the selected model is discretised using Forward-Euler, with a sampling time $\Delta = 15$ minutes, to obtain the discrete-time model

$$\mathbf{M}_c : \begin{cases} x_c[k+1] &= A_c x_c[k] + B_c u_c[k] + F_c d_c[k] + Q_c \\ y_c[k] &= [1 \ 0 \ 0 \ 0 \ 0 \ 0] x_c[k]. \end{cases} \quad (4)$$

Here, $m_{a,med}, m_{a,high}$ correspond to the air flow rates when the fan is operating at medium and high speeds.

Reachability analysis We perform reachability analysis of the hybrid model, which we run using SpaceEx (Frehse et al. [2011]). Notice that we do not discretise time and consider a continuous time horizon of 2 hours. We select two different initial conditions: in the first experiment we pick an initial condition equal to $T_{z_1} = 15^\circ C$ and $T_{sa} = 15^\circ C$, while in the second we set $T_{z_1} = 20^\circ C$ and $T_{sa} = 20^\circ C$. The resulting reach tube for the both experiments is shown in Figure 7a.

Fig. 7. Third case study: reach tubes obtained from two different initial conditions



In Figure 7a we can see that the model initially is in state $(O, -)$ and jumps to a new state (H, Op) such that warm outside air is pumped into the zone (due to the low temperature of the initial conditions). From (H, Op) it switches to (M, Op) (notice the reduction in the gradient between the variables $T_{sa} \in [15, 18], T_{z_1} \in [17, 19]$) and eventually switches back to $(O, -)$ in order to maintain the temperature within the comfort region. For Figure 7b, the reach tube shows that the system remains within the initial state $(O, -)$ over the whole time horizon.

ACKNOWLEDGEMENTS

This work is in part supported by the Alan Turing Institute, UK and Malta's ENDEAVOUR Scholarships Scheme. The authors would also like to thank Dario Cattaruzza, Sofie Haesaert and Honeywell Laboratories, Prague for their fruitful feedback.

REFERENCES

Abate, A., Prandini, M., Lygeros, J., and Sastry, S. (2008). Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 44(11), 2724–2734.

Belkhouane, H., Hensen, J., and Attia, S. (2017). Thermal comfort models for net zero energy buildings in hot climates. In *Second International Conference on Energy and Indoor Environment for Hot Climates*. Doha.

Cattaruzza, D., Abate, A., Schrammel, P., and Kroening, D. (2015). Unbounded-time analysis of guarded LTI systems with inputs by abstract acceleration. In *International On Static Analysis*, 312–331. Springer.

Cauchi, N. and Abate, A. (2018). Benchmarks for cyber-physical systems: A modular model library for building automation systems (extended version). *CoRR*, abs/1803.06315.

Crawley, D.B., Hand, J.W., Kummert, M., and Griffith, B.T. (2008). Contrasting the capabilities of building energy performance simulation programs. *Building and environment*, 43(4), 661–673.

Dragomir, I., Preoteasa, V., and Tripakis, S. (2017). The refinement calculus of reactive systems toolset. *arXiv preprint arXiv:1710.08195*.

Fehnker, A. and Ivančić, F. (2004). *Benchmarks for Hybrid Systems Verification*, 326–341. Springer Berlin Heidelberg, Berlin, Heidelberg. doi:10.1007/978-3-540-24743-2_22.

Ferrari, R.M.G., Dibowski, H., and Baldi, S. (2017). A message passing algorithm for automatic synthesis of probabilistic fault detectors from building automation ontologies. *IFAC-PapersOnLine*.

Frehse, G., Le Guernic, C., Donzé, A., Cotton, S., Ray, R., Lebeltel, O., Ripado, R., Girard, A., Dang, T., and Maler, O. (2011). SpaceX: Scalable verification of hybrid systems. In *Computer Aided Verification*, 379–395. Springer.

Haesaert, S., Cauchi, N., and Abate, A. (2017). Certified policy synthesis for general markov decision processes: An application in building automation systems. *Performance Evaluation*, 117, 75–103.

Holub, O., Zamani, M., and Abate, A. (2016). Efficient hvac controls: A symbolic approach. In *Control Conference (ECC), 2016 European*, 1159–1164. IEEE.

Kim, W. and Katipamula, S. (2017). A review of fault detection and diagnostics methods for building systems. *Science and Technology for the Built Environment*, 1–18.

Kristensen, N.R., Madsen, H., and Jørgensen, S.B. (2004). Parameter estimation in stochastic grey-box models. *Automatica*, 40(2), 225–237.

Ma, Y., Kelman, A., Daly, A., and Borrelli, F. (2012). Predictive control for energy efficient buildings with thermal storage: modeling, stimulation, and experiments. *Control Systems, IEEE*, 32(1), 44–64. doi:10.1109/MCS.2011.2172532.

Privara, S., Cigler, J., Vana, Z., Oldewurtel, F., Sager-schnig, C., and Zacekova, E. (2013). Building modeling as a crucial part for building predictive control. *Energy and Buildings*, 56, 8–22.

Soudjani, S.E.Z. and Abate, A. (2015). Aggregation and Control of Populations of Thermostatically Controlled Loads by Formal Abstractions. *Control Systems Technology, IEEE Transactions on*, 23(3), 975–990. doi:10.1109/TCST.2014.2358844.

Soudjani, S.E.Z., Gevaerts, C., and Abate, A. (2015). Faust²: Formal Abstractions of Uncountable-STate STochastic processes. In *TACAS*, volume 15, 272–286.

Yu, W., Li, B., Jia, H., Zhang, M., and Wang, D. (2015). Application of multi-objective genetic algorithm to optimize energy efficiency and thermal comfort in building design. *Energy and Buildings*, 88, 135–143.

Zhang, R. and Hong, T. (2017). Modeling of {HVAC} operational faults in building performance simulation. *Applied Energy*, 202, 178 – 188.