

A modular library of stochastic models from building automation systems (Benchmark Proposal)

Nathalie Cauchi¹ and Alessandro Abate¹

Department of Computer Science, University of Oxford, Oxford, U.K
`name.surname@cs.ox.ac.uk`

Abstract

We provide a modular library of stochastic models drawn from Building Automation Systems (BAS), specifically constructed from expertise developed on a real BAS setup. The library allows for the creation of a wide spectrum of stochastic models, representing either physical quantities or digital control modules, with differing levels of complexity. The structure, operation, and dynamics of the models can incorporate (i) stochasticity stemming from different sources (e.g. process or observation noise on the continuous variables; probabilistic terms in the digital modules), (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. We further describe two realistic case studies, where corresponding models are built and engaged under different analysis goals.

1 Introduction

In order to advance research on stochastic systems, it is essential to have a practical way to model and to reason about systems endowed with uncertainty. Uncertainty can be pervasive, and possibly cover both the discrete components and the continuous variables of the system. Currently, there are few and limited tools for simulation, verification, and policy synthesis over models of stochastic (hybrid) systems [4, 7]. These tools are typically specific to a particular type of model structure, are often limited to a small number of continuous variables and require expert knowledge on the specific formalism the tool relies on.

In this paper, we describe a library of stochastic models from the Building Automation Systems (BAS) area, resulting in general benchmarks that can ease the test and development of tools for simulation and verification of such models. The models are inspired by and built around an experimental setup within the Department of Computer Science at the University of Oxford, UK. The setup is part of on-going research in collaboration with estate engineers and industrial partners in the sector. The library allows one to construct models with different configurations and features: for instance, we can build low- to high-dimensional ones, with discrete or continuous inputs and states. The models are endowed with stochasticity which can represent un-modelled components, unknown parameters, random continuous effects, or likelihoods of transitioning to a certain discrete state from the current location. In order to




Figure 1: Building automation system setup

display the versatility of the new library of BAS models, we present two case studies that are built from its components. We focus on modelling temperature dynamics, a key element for ensuring thermal comfort. We employ the two generated models for different analysis goals, comprising simulation, reachability, and control synthesis.

This contribution has the following structure: Section 2 introduces the BAS modelling framework together with the library of models. We develop and analyse two case studies in Section 3. Finally, we present current challenges in modelling and analysis of stochastic systems in Section 4.

2 Building Automation Systems

2.1 BAS: structure and components

BAS models clearly depend on the size and topology of the building, and on its climate control setup. In this work, we consider the BAS setup shown in Figure 2a, which consists of two teaching rooms that are connected to a boiler-heated system. The boiler supplies heat to the heating coil within the air handling unit (AHU) and to two radiators. Valves control the rate of water flow within the heating coils and the radiators. The AHU supplies air to the two zones, which are connected back to back, and adjacent both to the outside and to an interior hall (see Figure 2a). 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 2b presents the Resistor Capacitance (RC) network circuit of the two zones [3], 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

We define models for the individual components in the BAS system. 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 [3]. We obtain models with a number of unknown parameters: these are estimated and validated using data collected from the BAS setup [6]. 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: algebraic relations define static couplings, whereas differential relations define the dynamics for the corresponding variables. The structure in Figure 2a, the quantities in Table 2, and the variables (with associated dynamics) in Table 3, together allow to construct global models for the complete 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 and models from the library, and build different BAS configurations.

Index	Reference	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	$jw \in \{5, 6\}$	zone walls with windows
$j \in \{jn \cup jw\}$	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

Table 1: Indices

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 2: List of variables, inputs, and parameters

A global model of the BAS set-up can be complex, comprising both algebraic and differential relations that are further affected by process noise. A model also contains a number of inputs which can either be construed as control signals or as exogenous signals. Some of the dynamics are non-linear in view of continuous variables that are bi-linearly coupled (cf. AHU air duct model in Table 3). Furthermore, the model features multiple components that present switching discrete behaviours, affecting the dynamics of the continuous variables. In order to tackle the complexity of global BAS models and to add a level of flexibility to the modelling framework, we consider each BAS component as a separate module, characterised by inputs and output elements, and by internal variables. We make use of individual modules describing component type, and then connect different modules based on possible physical couplings. Coupling 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

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	$dT_{z_i}(t) = (C_{z_i})^{-1} \left[\frac{T_{w_{j_n}}(t) - T_{z_i}(t)}{R_{ij}} + Q_{rw,r,i}(t) + Q_{occ,i}(t) + Q_{sa,i}(t) \right] dt + \sigma_{z_i}dW$ $dT_{w_{j_n}}(t) = (C_{w_{j_n}})^{-1} \left[\frac{T_{adj,out}(t) - T_{z_i}(t)}{R_{out}} + \sum_l \frac{T_{adj,l}(t) - T_{w_{j_n}}(t)}{R_{lj}} + Q_{rw,a_{j_n}}(t) \right] dt + \sigma_{w_{j_n}}dW$ $dT_{w_{j_w}}(t) = (C_{w_{j_w}})^{-1} \left[\frac{T_{adj,out}(t) - T_{z_i}(t)}{R_{out}} + \sum_l \frac{T_{adj,l}(t) - T_{w_{j_w}}(t)}{R_{ljw}} + Q_{solar,j_w}(t) + Q_{rw,a_{j_w}}(t) \right] dt + \sigma_{w_{j_w}}dW$ $Q_{rw,r,i}(t) = P_{rad,i}(\alpha_2(T_{rw,r,i}(t) - T_{z_i}(t)) + \alpha_1), Q_{occ,i}(t) = \mu_i(CO_{2,i}(t)) + \beta_{1,i}, Q_{sa,i}(t) = m_a(t)C_{pa}(T_{sa,i}(t) - T_{z_i}(t))$ $Q_{rw,a_{j_i}}(t) = \alpha_3(T_{rw,a}(t) - T_{w_{j_i}}(t)), Q_{solar,j_w}(t) = (\alpha_0 A_i T_{out}(t) + \beta_2)$	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

Table 3: Dynamics and functional relations among component variables

for the wall separating the two zones (cf. W_7 in Figure 2b) corresponds to the individual zone temperatures of the two zone modules (cf. Table 3 zone equations). Having such a modular structure for the individual components provides an added level of versatility, since we can connect different components to create various 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.

2.3 BAS: 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 bilinear depending on the component they represent. They are defined using the symbolic toolbox, are parametrised and can be described both in discrete and continuous time. We provide the parameters which we estimated from data gathered from the BAS lab to construct the individual models. However, users can easily input their own parameters and construct their own model. Different components can be connected together based on their input-output relations by cascading the different symbolic models for each component. The provided scripts also allow you to simulate the models and plots for the defined output variables are presented.

3 Case studies

We set up two case studies and for each (i) we establish the dynamics of the models, (ii) how they are constructed from the library of components, and (iii) we describe the results obtained.

3.1 Two-zone heating setup with stochastic dynamics




Figure 3: BAS setup for the first case study

We consider two zones, each heated by one radiator and with a common supply air, as portrayed in Figure 3. 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 $x^T = [T_{z_1} \ T_{z_2} \ T_{rw,r_1} \ T_{rw,r_2}]^T$, a common supply temperature $u = T_{sa}$ as an input and process noise associated with each state. For this setup, we discretise using a Euler-Maruyama scheme having a uniform sampling time $\Delta = 15$ minutes, and obtain a stochastic linear discrete-time model expressed using (2). Here, the matrices A , B are properly sized and constructed based on the models in Table 3, $Q = [\frac{T_{w,ss}\Delta}{C_{z_1}R_1} \ \frac{T_{w,ss}\Delta}{C_{z_2}R_2} \ \frac{C_{pw}w_{r_1}\Delta}{C_{pw}\rho_hV_{r_1,b}}T_{sw,b,ss} \ \frac{C_{pw}w_{r_2}\Delta}{C_{pw}\rho_hV_{r_2,b}}T_{sw,b,ss}]^T$ with R_i representing the mean resistance offered by the walls; and $\Sigma = \text{diag}([\sqrt{\Delta}\sigma_{z_1}]^2 \ (\sqrt{\Delta}\sigma_{z_2})^2 \ (\sqrt{\Delta}\sigma_{rw,r_1})^2 \ (\sqrt{\Delta}\sigma_{rw,r_2})^2]$ encompasses the variances of the process noise for each state. $W = [w_1 \ w_2 \ w_3 \ w_4]^T$ are independent Gaussian random variables, which are also independent of the initial condition of the process. A simulation run is depicted in Figure 5a.

3.1.1 Probabilistic reachability analysis

For this case study we would like to 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 probabilistic reachability analysis. The safe set is described as an interval around the temperature set-point $T_{SP} = 20^\circ\text{C} \pm 0.5^\circ\text{C}$. 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 use FAUST² [7] to perform probabilistic reachability analysis of $\mathbf{M}_s$. We define the safe set $[T_{z_1} \ T_{z_2} \ T_{rw,r_1} \ T_{rw,r_2}]^T = [18 \ 18 \ 35 \ 35]^T$ and assume an input set of $[15 \ 22]$. The resulting adaptive partition of the safe set along with the optimal safety probability for each partition set is depicted in Figure 5b. We can deduce that the model $\mathbf{M}_s$ has a high probability of being within the required safe set, specifically to have $T_{z_1} \in [19.5 \ 20]$ and $T_{z_2} \in [19.5 \ 20.5]$.

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

In this second case study we focus on the stochastic dynamics of the zone component from Table 3 and consider the two zones shown in Figure 2b. We assume that (i) a central fan pumps in air in both rooms with a common supply temperature $15^\circ\text{C} \leq T_{sa} \leq 30^\circ\text{C}$, (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




Figure 4: First case study

heating coils is fixed ($T_{rw,ass} = 35^\circ C$). The selected model is discretised using Forward-Euler, with a sampling time $\Delta = 15$ minutes, to obtain the discrete-time model

$$\mathbf{M}_c : \begin{cases} x_c[k+1] &= A_c x_c[k] + B_c u_c[k] + F_c d_c[k] + Q_c \\ y_c[k] &= \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} x_c[k]. \end{cases} \quad (2)$$

Here the variables are $x_c = [T_{z_1} \ T_{z_2} \ T_{w_5} \ T_{w_6} \ T_{w_2} \ T_{w_3} \ T_{w_7}]^T$, and a common fan supplies the two zones with a supply rate $u_c = T_{sa}$, whereas $d_c = [T_{out} \ T_{hall} \ CO_{2_1} \ CO_{2_2} \ T_{rw,r_1} \ T_{rw,r_2}]^T$ and $Q_c = \left[\frac{q_{c0}\Delta}{C_{z_1}} \ \frac{q_{c0}\Delta}{C_{z_2}} \ \frac{q_{c2}\Delta}{C_{w_5}} \ \frac{q_{c2}\Delta}{C_{w_6}} \ \frac{q_{c1}\Delta}{C_{w_2}} \ \frac{q_{c1}\Delta}{C_{w_3}} \ \frac{q_{c1}\Delta}{C_{w_7}} \right]$ where $q_{c0} = \beta_{1_i} + \alpha_1 P_{rad_i}$, $q_{c1} = \alpha_3 T_{rw,ass}$ and $q_{c2} = \alpha_0 A_2 \beta_2 + q_{c1}$. Matrices A_c, B_c, F_c are properly sized. The vector d_c corresponds to the disturbance signals, while Q_c represents constant additive terms within the model. We model the disturbances as random external effects following a Gaussian distributions with a mean μ and variance σ , affecting the room temperature dynamics as $T_{out}[k] \sim \mathcal{N}(\mu = 9, \sigma = 1)$, $T_{hall}[k] \sim \mathcal{N}(\mu = 15, \sigma = 1)$, $CO_{2_i}[k] \sim \mathcal{N}(\mu = 500, \sigma = 100), i \in \{1, 2\}$, $T_{rw,r_i}[k] \sim \mathcal{N}(\mu = 35, \sigma = 5), i \in \{1, 2\}$.

3.2.1 Policy synthesis and refinement

For $\mathbf{M}_c$ we would like to synthesise a policy ensuring that the temperature within zone 1 does not deviate from the set point by more than $0.5^\circ C$ over a time horizon equal to four hours (i.e $N = 16$). This requirement can be translated into the following PCTL property $\Phi := \mathbb{P}_{=p}[\Box^{\leq N=16} |T_{z_1} - T_{SP}| \leq 0.5]$. We then aim at synthesising a policy maximising the safety probability p . This synthesis goal can be computationally hard due to the number of continuous variables making up $\mathbf{M}_c$. To mitigate this limitation, we perform policy synthesis via abstractions [3]. We simplify (2) into four abstract models and are labelled as $\mathbf{M}_{c_{a=\{4,\dots,1\}}}$, where a represents the number of continuous variables of the corresponding abstract model. We can quantify the error in the output variable, which has been introduced by the different levels of abstractions, through the use of (ε, δ) -approximate simulation relations [3]. The pair (ε, δ) represents the deviation in the output trajectories between complex and abstract models and the differences in probability distribution of the processes, respectively. Such metrics allows the designer to select which of the considered abstract models provides the best trade off in precision: it is desirable to achieve little deviation in both the output trajectories (small ε) and in the probability distributions (small δ). We compute (ε, δ) -approximate simulation relations between $\mathbf{M}_c$ and the set of abstract models $\mathbf{M}_{c_{a=\{4,\dots,1\}}}$, as presented in Table 4. The (ε, δ)

δ	1	$10^{-\frac{1}{2}}$	10^{-1}	$10^{-\frac{3}{2}}$	10^{-2}	$10^{-\frac{5}{2}}$	10^{-3}
$\mathbf{M}_{c_4}$	0.0008	0.1754	0.2084	0.2339	0.2555	0.2745	0.2910
$\mathbf{M}_{c_3}$	0.0006	0.1933	0.2312	0.2598	0.2831	0.3065	0.3241
$\mathbf{M}_{c_2}$	0.0011	0.1950	0.2373	0.2681	0.2928	0.3155	0.3278
$\mathbf{M}_{c_1}$	0.0010	0.1953	0.2371	0.2595	<u>0.2854</u>	0.3103	0.3254

Table 4: Second case study: error metrics (ε, δ) for concrete and abstract models

pair providing the optimal trade off is obtained with the abstract model $\mathbf{M}_{c_1}$ and corresponds to $(0.2854, 10^{-2})$. Next, we use FAUST² to perform a grid-based computation of the safety probability for $\mathbf{M}_{c_1}$ and obtain a model of size 14893 with an overall accuracy of 0.005. Over this approximation we synthesise the optimal policy for the abstract model which results in a safety probability of $p' = 0.9257$. We refine the obtained policy [3], which results in one that can be used with $\mathbf{M}_c$. The overall process results in Φ being satisfied with a safety probability of $p = p' - \eta - N\delta = 0.7657$, where η is the abstraction error introduced by FAUST². The results obtained further highlight that by trading off the complexity in the number of continuous variables and computing (ε, δ) -simulation relations, we can synthesise policies using simpler models, yet achieve high performance still when the refined policy is applied to the original model.

4 Challenges

The generality and richness of stochastic models introduces a number of difficulties when performing analysis. These difficulties further reflect on the absence, to the best of the authors' present knowledge, of general and scalable software tools for verification of and synthesis over stochastic models. More specifically, we note that

- in contrast to the deterministic case, verification of continuous models with stochastic and hybrid elements has not matured to the point where the community has agreed upon an input standard that can be used to exchange problems. Consequently, benchmarking different tools over stochastic models is at present a manual task. A first step towards achieving automation is the development of a common description language, of which JANI [2] is a notable first attempt.
- general verification goals are bound to be undecidable, and their approximated (and finite) versions are frequently stymied by state-space explosion. This restricts the level of detail and complexity attainable for the verification of stochastic models. Furthermore, at present there exist different stochastic verification tools dealing with specific dynamical features and complexities, which makes it hard to compare results in the area. For instance, Modest [4] deals with stochastic hybrid models in continuous time under the assumption that the continuous variables are described by ordinary differential equations. On the other hand, FAUST² [7] deals with stochastic hybrid models in discrete time with the continuous variables endowed with process noise following a Gaussian distribution. Different still, [8] perform stochastic reachability analysis on uncontrolled, non-hybrid, linear models using Fourier transforms.

- Similar to the previously raised point, the synthesis of control policies on such models is also hindered by the curse of dimensionality. To this end, different techniques aimed at speeding up formal algorithms through abstraction [7], or compositionality [3], or by data-driven approaches, such as reinforcement learning [1, 5] have been developed. However, they all require specific types of models and work under their own set of assumptions.

These considerations highlight the evident need for further research in the area of computational verification and control synthesis over stochastic models. With this grand aim in mind, in this paper we propose a set of general benchmarks for stochastic models: we see this as a catalyst for further work on development of algorithms and software tools. In order to bolster this action, we have further set up a new category (*stochastic models*) within the ARCH friendly competition, and are aggregating interests of key players - often working separately on similar problems in view of their different research areas and communities. We are as well working towards obtaining a compositional software tool that is able to allow for easy construction of stochastic models and for a seamless interfacing with different verification and synthesis tools.

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 Sofie Haesaert, Sadegh Soudjani, and Honeywell Laboratories (Prague) for their fruitful feedback.

References

- [1] Tomáš Brázdil, Krishnendu Chatterjee, Martin Chmelik, Vojtech Forejt, Jan Kretínský, Marta Z. Kwiatkowska, David Parker, and Mateusz Ujma. Verification of markov decision processes using learning algorithms. In *Automated Technology for Verification and Analysis - 12th International Symposium, ATVA 2014, Sydney, NSW, Australia, November 3-7, 2014, Proceedings*, pages 98–114, 2014.
- [2] Carlos E Budde, Christian Dehnert, Ernst Moritz Hahn, Arnd Hartmanns, Sebastian Junges, and Andrea Turrini. Jani: Quantitative model and tool interaction. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 151–168. Springer, 2017.
- [3] Sofie Haesaert, Nathalie Cauchi, and Alessandro Abate. Certified policy synthesis for general markov decision processes: An application in building automation systems. *Performance Evaluation*, 117:75–103, 2017.
- [4] Ernst Moritz Hahn, Arnd Hartmanns, Holger Hermanns, and Joost-Pieter Katoen. A compositional modelling and analysis framework for stochastic hybrid systems. *Formal Methods in System Design*, 43(2):191–232, 2013.
- [5] Mohammadhosein Hasanbeig, Alessandro Abate, and Daniel Kroening. Logically-correct reinforcement learning. *CoRR*, 2018.
- [6] Niels Rode Kristensen, Henrik Madsen, and Sten Bay Jørgensen. Parameter estimation in stochastic grey-box models. *Automatica*, 40(2):225–237, February 2004.
- [7] Sadegh Esmaeil Zadeh Soudjani, Caspar Gevaerts, and Alessandro Abate. Faust²: Formal Abstractions of Uncountable-STate STOchastic processes. In *TACAS*, volume 15, pages 272–286, 2015.
- [8] Abraham P Vinod, Baisravan Homchaudhuri, and Meeko MK Oishi. Forward stochastic reachability analysis for uncontrolled linear systems using fourier transforms. In *Proceedings of the 20th International Conference on Hybrid Systems: Computation and Control*, pages 35–44. ACM, 2017.