

FAUST²: Formal Abstractions of Uncountable-State Stochastic Processes

Sadegh Esmail Zadeh Soudjani¹, Caspar Gevaerts¹, and Alessandro Abate^{2,1,*}

¹ Delft Center for Systems and Control (DCSC),
TU Delft – Delft University of Technology, Delft, The Netherlands
S.EsmailZadehSoudjani@tudelft.nl

² Department of Computer Science, University of Oxford, Oxford, United Kingdom
alessandro.abate@cs.ox.ac.uk

Abstract. FAUST² is a software tool that generates formal abstractions of (possibly non-deterministic) discrete-time Markov processes (dtMP) defined over uncountable (continuous) state spaces. A dtMP model is specified in MATLAB and abstracted as a finite-state Markov chain or a Markov decision process. The abstraction procedure runs in MATLAB and employs parallel computations and fast manipulations based on vector calculus, which allows scaling beyond state-of-the-art alternatives. The abstract model is formally put in relationship with the concrete dtMP via a user-defined maximum threshold on the approximation error introduced by the abstraction procedure. FAUST² allows exporting the abstract model to well-known probabilistic model checkers, such as PRISM or MRMC. Alternatively, it can handle internally the computation of PCTL properties (e.g. safety or reach-avoid) over the abstract model. FAUST² allows refining the outcomes of the verification procedures over the concrete dtMP in view of the quantified and tunable error, which depends on the dtMP dynamics and on the given formula. The toolbox is available at

<http://sourceforge.net/projects/faust2/>

1 Models: Discrete-Time Markov Processes

We consider a discrete-time Markov process (dtMP) $s(k), k \in \mathbb{N} \cup \{0\}$ defined over a general state space, such as a finite-dimensional Euclidean domain [1] or a hybrid state space [2]. The model is denoted by the pair $\mathfrak{S} = (\mathcal{S}, T_{\mathfrak{S}})$. \mathcal{S} is a continuous (uncountable) but bounded state space, e.g. $\mathcal{S} \subset \mathbb{R}^n, n < \infty$. We denote by $\mathcal{B}(\mathcal{S})$ the associated sigma algebra and refer the reader to [2, 3] for details on measurability and topological considerations. The conditional stochastic kernel $T_{\mathfrak{S}} : \mathcal{B}(\mathcal{S}) \times \mathcal{S} \rightarrow [0, 1]$ assigns to each point $s \in \mathcal{S}$ a probability measure $T_{\mathfrak{S}}(\cdot|s)$, so that for any set $A \in \mathcal{B}(\mathcal{S}), k \in \mathbb{N} \cup \{0\}$,

$$\mathbb{P}(s(k+1) \in A | s(k) = s) = \int_A T_{\mathfrak{S}}(d\bar{s}|s).$$

* This work has been supported by the European Commission STREP project MoVeS 257005 and IAPP project AMBI 324432, and by the John Fell OUP Research Fund.

We refer to the code or to the case study for a modeling example. The software allows handling the relevant instance of Stochastic Hybrid Systems (SHS). SHS are discrete-time Markov processes evolving over hybrid state spaces. The hybrid state $s = (q, x)$ of SHS has two components: $q \in \mathcal{Q}$ is the discrete part, and $x \in \mathbb{R}^{n_q}$ is the continuous part. The state space of the SHS is the (disjoint) union of continuous spaces associated to the discrete locations $\mathcal{S} \subset \cup_{q \in \mathcal{Q}} \{q\} \times \mathbb{R}^{n_q}$. The formal definition and characterization of the conditional stochastic kernel of a SHS, along its theoretical analysis and formal verification, are discussed in detail in [2, 4–9].

Implementation: FAUST² is implemented in MATLAB and its user interaction is enhanced by a Graphical User Interface as in Figure 1. A dtMP model is fed into FAUST² as follows. Select the Formula free option in the box Problem selection ① in Figure 1, and enter the bounds on the state space \mathcal{S} as a $n \times 2$ matrix in the prompt Domain in box ⑧. Alternatively if the user presses the button Select ⑧, a pop-up window prompts the user to enter the lower and upper values of the box-shaped bounds of the state space. The transition kernel T_s can be specified by the user (select User-defined ②) in an m-file, entered in the text-box Name of kernel function, or loaded by pressing the button Search for file ⑦. Please open the files `./Templates/SymbolicKernel.m` for a template and `ExampleKernel.m` for an instance of kernel T_s . As a special case, the class of affine dynamical systems with additive Gaussian noise is described by the difference equation $s(k+1) = As(k) + B + \eta(k)$, where $\eta(\cdot) \sim \mathcal{N}(0, \text{Sigma})$. (Refer to the Case Study on how to express the difference equation as a stochastic kernel.) For this common instance, the user can select the option Linear Gaussian model in the box Kernel distribution ②, and input properly-sized matrices `A, B, Sigma` in the MATLAB workspace. FAUST² also handles Gaussian dynamical models $s(k+1) = f(s(k)) + g(s(k))\eta(k)$ with nonlinear drift and variance: select the bottom option in box ② and enter the symbolic function `[f g]` via box ⑦. \square

The software also handles models with non-determinism [9], here regarded as *external* and as such accessible: a *controlled* dtMP is a tuple $\mathfrak{S} = (\mathcal{S}, \mathcal{U}, T_s)$, where \mathcal{S} is as before, \mathcal{U} is a continuous control space (e.g. a bounded set in \mathbb{R}^m), and T_s is a Borel-measurable stochastic kernel $T_s : \mathcal{B}(\mathcal{S}) \times \mathcal{S} \times \mathcal{U} \rightarrow [0, 1]$, which assigns to any state $s \in \mathcal{S}$ and input $u \in \mathcal{U}$ a probability measure $T_s(\cdot | s, u)$.

Implementation: In order to specify a non-deterministic model in FAUST², tick the relevant check Controlled/non-deterministic model ③, and enter the bounds on the space \mathcal{U} as a $m \times 2$ matrix in the window Input set ⑧. \square

2 Formal Finite-State Abstractions of dtMP Models

This section discusses the basic procedure to approximate a dtMP $\mathfrak{S} = (\mathcal{S}, T_s)$ as a finite-state Markov chain (MC) $\mathfrak{P} = (\mathcal{P}, T_p)$, as implemented in FAUST². $\mathcal{P} = \{z_1, z_2, \dots, z_p\}$ is a finite set of abstract states of cardinality p , and $T_p : \mathcal{P} \times \mathcal{P} \rightarrow [0, 1]$ is a transition probability matrix over the finite space \mathcal{P} : $T_p(z, z')$ characterizes the probability of transitioning from state z to state z' . The finite




Fig. 1. Graphical User Interface of FAUST². Overlaid numbered boxes refer to specific description in the text.

state space of \mathfrak{P} is constructed by (arbitrarily) partitioning the state space of \mathfrak{S} and selecting a representative point in each partition set to make up the states in \mathcal{P} . The probability of transitioning from (abstract) state z to state z' , $T_p(z, z')$, is computed by marginalizing the stochastic kernel T_s of \mathfrak{S} , namely computing the probability of jumping from state z to any point in the (concrete) partition set corresponding to the (abstract) state z' . Algorithm 1 describes the abstraction of model \mathfrak{S} as a finite-state MC \mathfrak{P} [6]. In Algorithm 1, $\Xi : \mathcal{P} \rightarrow 2^{\mathcal{S}}$ represents a set-valued map that associates to any point $z_i \in \mathcal{P}$ the corresponding partition set $A_i \subseteq \mathcal{S}$, whereas the map $\xi : 2^{\mathcal{S}} \rightarrow \mathcal{P}$ relates any point s or set in \mathcal{S} to the corresponding discrete state in \mathcal{P} .

Algorithm 1. Abstraction of dtMP \mathfrak{S} by MC \mathfrak{P}

Require: input dtMP $\mathfrak{S} = (\mathcal{S}, T_s)$

1: Select a finite partition of the state space \mathcal{S} as $\mathcal{S} = \cup_{i=1}^p A_i$ (A_i are non-overlapping)

2: For each A_i , select an arbitrary representative point $z_i \in A_i$, $\{z_i\} = \xi(A_i)$

3: Define $\mathcal{P} = \{z_i, i = 1, \dots, p\}$ as the finite state space of the MC \mathfrak{P}

4: Compute the transition probability matrix $T_p(z, z') = T_s(\Xi(z')|z)$ for all $z, z' \in \mathcal{P}$

Ensure: output MC $\mathfrak{P} = (\mathcal{P}, T_p)$

Consider the representation of the kernel T_s by its density function $t_s : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}^{\geq 0}$, namely $T_s(ds'|s) = t_s(s'|s)ds'$ for any $s, s' \in \mathcal{S}$. The abstraction error over the next-step probability distribution introduced by Algorithm 1 depends on the regularity of function t_s : assuming that t_s is Lipschitz continuous, namely that there is a finite positive constant h_s such that

$$|t_s(\bar{s}|s) - t_s(\bar{s}|s')| \leq h_s \|s - s'\|, \quad \forall s, s', \bar{s} \in \mathcal{S}, \quad (1)$$

then the next-step error is $E = h_s \delta_s \mathcal{L}(\mathcal{S})$, where δ_s is the max diameter of the state-space partition sets and $\mathcal{L}(\mathcal{S})$ is the volume of the state space [6]. When interested in working over a finite, N -step time horizon, the error results in the quantity EN . Notice that the error can be reduced via δ_s by considering a finer partition, which on the other hand results in a MC \mathfrak{P} with a larger state space. It is evidently key to obtain error bounds that are as tight as possible: the error bounds on the abstraction can be improved in three different ways [8, 10]. First, by computing a local version of the error; second, by leveraging continuity requirements that go beyond the Lipschitz condition raised in (1); and, finally, by normalizing possibly ill-conditioned dynamics operating on heterogeneous spatial scales.

Implementation: FAUST² enables the user to enter the time horizon N of interest (box Number of time steps ⑤), and a threshold on the maximum allowed error (box Desired abstraction error ⑤). The software generates a Markov chain with the desired accuracy by pressing the button Generate the abstraction ⑥. Among other messages, the user is prompted with an estimated running

time, which is based on an over-approximation of the Lipschitz constant of the kernel, on a uniform partitioning of the space \mathcal{S}^1 , and on the availability of parallelization procedures in MATLAB, and is asked whether to proceed. \square

In the case of a non-deterministic dtMP, the input space is also partitioned as $\mathcal{U} = \cup_{i=1}^q U_i$, and arbitrary points $u_i \in U_i$ are selected. The dtMP \mathfrak{G} is abstracted as a Markov decision process (MDP) $\mathfrak{P} = (\mathcal{P}, \mathcal{U}_p, T_p)$, where now the finite input space is $\mathcal{U}_p = \{u_1, u_2, \dots, u_q\}$, and $T_p(u, z, z') = T_s(\Xi(z')|z, u)$ for all $z, z' \in \mathcal{P}, u \in \mathcal{U}_p$. The abstraction error can be formally quantified as $E = 2(h_s \delta_s + h_u \delta_u) \mathcal{L}(\mathcal{S})$, where δ_u is the max diameter of the input-space partitions and h_u is the Lipschitz constant of the density function with respect to the inputs, that is $|t_s(\bar{s}|s, u) - t_s(\bar{s}|s, u')| \leq h_u \|u - u'\|, \forall u, u' \in \mathcal{U}, s, \bar{s} \in \mathcal{S}$.

Implementation: The user may tick the check in ③ to indicate that the dtMP is controlled (non-deterministic), specify a box-shaped domain for the input in box **Input set** ⑧, enter a time horizon in box **Number of time steps** ⑤, and require an error threshold in box **Desired abstraction error** ⑤. FAUST² automatically generates an MDP according to the relevant formula on the error.

Notice that the quantification of the abstraction error requires state and input spaces to be bounded. In the case of an unbounded state space, the user should truncate it to a bounded, box-shaped domain: selecting the **Formula free** option in the box **Problem selection** ①, the domain is prompted in box **Domain** ⑧. Algorithm 1 is automatically adjusted by assigning an absorbing abstract state to the truncated part of the state space. For details please see [10, 11].

The user may select one of two options in the box **Gridding procedure** ④, dealing with adaptive gridding. FAUST² generates partition sets based on local computation of the error, as follows: a rather course partition of the state space is initially selected and the corresponding local errors are computed; the partition is sequentially refined by splitting the sets if the maximum local error is greater than the threshold entered by the user in box ⑤. In the step, the partition can be refined by splitting the partition set with the largest local error, which results in an abstraction with the least number of states but requires a larger computational time (cf. [10, Algorithm 3]); alternatively, FAUST² obtains faster generation time by splitting all the sets with local errors greater than threshold (cf. [10, Algorithm 4]). Both procedures are guaranteed to result, after a finite number of steps, in the desired selected error. We plan to implement an *anytime algorithm* option, as possible for the discussed adaptive gridding procedures. This option enables the user to externally interrupt the refinement procedure at any time, and returns the last computed abstract model together with its associated error bound. \square

The states of the abstract model \mathfrak{P} may be labeled. The state labeling map $L : \mathcal{P} \rightarrow \Sigma$, where Σ is a finite alphabet, is defined by a set of linear inequalities: for any $\alpha \in \Sigma$ the user characterizes the set of states $L^{-1}(\alpha)$ as the intersection of half-planes (say, as a box or a simplex): the software automatically determines

¹ At the moment we assume to have selected options **Uniform gridding** and **Lipschitz via integral** among the lists in box ④. Comments on further options are in Section 3.

all points $z \in \mathcal{P}$ belonging to set $L^{-1}(\alpha)$. The obtained labeled finite-state model can be automatically exported to well-known model checkers, such as PRISM and MRMC [12, 13], for further analysis. In view of the discussed error bounds, the outcomes of the model checking procedures over the abstract model \mathfrak{P} may be refined over the concrete dtMP \mathfrak{S} – more details can be found in [6, 11].

Implementation: Labels are introduced in FAUST² as follows: suppose that the intersection of half-planes $A_\alpha z \leq B_\alpha$ (where A_α, B_α are properly-sized matrices) tags states z by label $\alpha \in \Sigma$. The user may add such a label by pressing button **Add** (10) and subsequently entering symbol α and matrices A_α, B_α in the pop-up window. The user can also edit or remove any previously defined label using buttons **Edit**, **Remove** in (10), respectively. The button **States with selected label** (10) shows the sets associated to the active label over the plot in (13).

The user may click the buttons in (11) to export the abstracted model to PRISM or to MRMC. Alternatively, FAUST² is designed to automatically check or optimize over quantitative, non-nested PCTL properties, without relying on external model checkers: Section 3 elaborates on this capability. \square

3 Formula-Dependent Abstractions for Verification

Algorithm 1, presented in Section 2, can be employed to abstract a dtMP as a finite-state MC/MDP, and to directly check it against properties such as probabilistic invariance or reach-avoid, that is over (quantitative, non-nested) bounded-until specifications in PCTL [14]. Next, we detail this procedure for the finite-horizon probabilistic invariance (a.k.a. safety) problem, which can be formalized as follows. Consider a bounded continuous set $A \in \mathcal{B}(\mathcal{S})$ representing the set of safe states. Compute the probability that an execution of \mathfrak{S} , associated with an initial condition $s_0 \in \mathcal{S}$ remains within set A during the finite time horizon $[0, N]$, that is $p_{s_0}(A) := \mathbb{P}\{s(k) \in A \text{ for all } k \in [0, N] | s(0) = s_0\}$.

The quantity $p_{s_0}(A)$ can be employed to characterize the satisfiability set of a corresponding bounded-until PCTL formula, namely

$$s_0 \models \mathbb{P}_{\sim \epsilon}\{\text{true } U^{\leq N}(\mathcal{S} \setminus A)\} \Leftrightarrow p_{s_0}(A) \sim 1 - \epsilon,$$

where $\mathcal{S} \setminus A$ is the complement of A over \mathcal{S} , **true** is a state formula valid everywhere on \mathcal{S} , the inequality operator $\sim \in \{>, \geq, <, \leq\}$, and \sim represents its complement.

FAUST² formally approximates the computation of $p_{s_0}(A), \forall s_0 \in \mathcal{S}$, as follows. \mathfrak{S} is abstracted as an MC \mathfrak{P} via Algorithm 1: the bounded safe set A is partitioned as $A = \cup_{i=1}^{p-1} A_i$; representative points $z_i \in A_i$ are selected and, along with an extra absorbing variable ϕ for $\mathcal{S} \setminus A$, make up the state space \mathcal{P} ; the transition probability matrix T_p is obtained by marginalizing the concrete kernel T_s . Given the obtained discrete-time MC $\mathfrak{P} = (\mathcal{P}, T_p)$ and considering the finite safe set $A_p = \{z_1, \dots, z_{p-1}\} \subset \mathcal{P}$, FAUST² internally computes the safety probability over \mathfrak{P} via dynamic programming [6], along with the associated abstraction error which is now tailored over the PCTL formula of interest.

Implementation: The user may select option **PCTL Safety** in the list within box (1), enter the boundaries of the **Safe** set within box (8), and press button (6) to

proceed obtaining the abstraction and computing the probability of the selected formula. The computed value of $p_{s_0}(A)$ is displayed in box **Probability given s0** (14), for any user-selected initial state s_0 that is input in box **Initial condition s0** (14). The user can optionally press button **Properties of s0** (14) to get more information about the concrete state s_0 , including the related discrete state $z = \xi(\Xi(s))$ of the MC, as well as the associated labels. Furthermore, the quantity $p_{s_0}(A)$ can be plotted, as a function of the initial state s_0 , by pressing buttons **Plot grid** and **Color grid** in (13). Clearly these outputs are exclusively available for models of dimensions $n = 1, 2, 3$. \square

It is of interest to obtain tight bounds on the error associated to the abstraction procedure since, given a user-defined error threshold, tighter bounds would generate abstract models \mathfrak{P} with fewer states. The abstraction error bound in FAUST², tailored around the discussed safety problem, can be efficiently decreased under different types of regularity assumptions on the conditional density function of the dtMP \mathfrak{S} [10]. For instance, in contrast to the global continuity assumption in (1), if t_s is locally Lipschitz continuous as

$$|t_s(\bar{s}|s) - t_s(\bar{s}|s')| \leq h(i, j) \|s - s'\|, \quad \forall \bar{s} \in A_j, \forall s, s' \in A_i, \quad (2)$$

(here sets A_i form a partition of A , as from Algorithm 1) then the error is

$$|p_{s_0}(A) - p_{p_0}(A_p)| \leq \max\{\gamma_i \delta_i | i = 1, \dots, p\}, \quad (3)$$

where $p_{p_0}(A_p)$ is initialized at the discrete state $p_0 = \xi(s_0) \in A_p$. Here δ_i is the diameter of the set $A_i \subset A$, and the constants γ_i are given by $\gamma_i = N \sum_{j=1}^m h(i, j) \mathcal{L}(A_j)$. Since $h(i, j) \leq h_s$, the obtained error in (3) is smaller than the older quantity $N h_s \delta_s \mathcal{L}(\mathcal{S})$. Notice that the structure of the error in (3) leads to gridding algorithms for abstraction that are adapted to the formula and can be made sequential [10]: FAUST² initializes the procedure with coarse partition sets (resulting in a small MC abstraction but with a large approximation error), and sequentially refines the partitions adaptively where the local errors are high (leading to an MC abstraction with increasing state space), until the global error becomes less than a user-defined threshold.

Implementation: FAUST² allows the user to select three different gridding procedures in box **Gridding procedure** (4): the reader is referred to [10] for the details of these three options. The **Uniform gridding** option leads to a one-shot (non sequential) procedure, as already discussed in Section 2, whereas the two **Adaptive gridding** options result in sequential and adaptive procedures leading to better errors and to smaller abstractions, but in general requiring more computation time. The error bound quantification hinges on the constant in the right-hand side of (2), which can be computed differently as in box **Assumptions on kernel** (4): tighter errors lead to longer computations [10]. In order to provide full control on the chosen inputs, for any possible selection of gridding procedure, desired abstraction error, and error bound computation, the user is prompted in a pop-up window with an estimated running time, and asked whether to proceed.

This range of algorithms and procedures are also implemented for probabilistic reach-avoid (constrained reachability) problems, which are encompassed by general bounded-until PCTL formulas $\mathbb{P}_{\sim \epsilon}\{\Phi \text{ U}^{\leq N} \Psi\}$. The user can select this option in box **Problem selection** ①, and is asked to input sets Φ, Ψ as safe and target sets in the texts in box ⑧.

Let us remark that the described abstraction algorithms and procedures are also available for the formula-free abstraction discussed in Section 2. \square

The safety problem for a controlled dtMP [9] is defined as follows. Consider the class of deterministic Markov policies $\pi = (\mu_0, \mu_1, \dots)$, where the functions $\mu_k : \mathcal{S} \rightarrow \mathcal{U}$ are properly measurable deterministic functions. The safety probability for a controlled dtMP under a given policy π is given by

$$p_{s_0}^\pi(A) := \mathbb{P}\{s(k) \in A \text{ for all } k \in [0, N] | s(0) = s_0, u(k+1) = \mu_k(s(k))\}.$$

The safety problem deals with the computation of the *maximally safe* deterministic Markov policy π^* , such that $p_{s_0}^{\pi^*}(A) = \sup_{\pi} p_{s_0}^\pi(A)$, $\forall s_0 \in A$. Similarly we can compute the *minimally safe* policy, or an optimal policy related to the reach-avoid problem (defined with the bounded-until operator).

Implementation: FAUST² computes a suboptimal policy for a given problem over an MDP, with a given threshold on the distance to the optimal safety probability, and quantifies the corresponding approximate quantity $p_{s_0}^{\pi^*}(A)$. The approximate optimal policy can be stored by pushing button **Save results** ⑫, which provides the user with two options: either storing it in the disk as a .mat file, or loading it to the workspace. \square

4 Accessing and Testing FAUST²

The toolbox is available at

<http://sourceforge.net/projects/faust2/>

We have successfully tested the toolbox with MATLAB R2012a, R2012b, R2013a, R2013b, on machines running Windows 7, Apple OSX 10.9, and Linux OpenSUSE. FAUST² exploits the command `integral` of MATLAB (introduced in version R2012a) for numerical integrations. (The previous versions of MATLAB contain instruction `quad` and its variations, which will be removed in the future versions of MATLAB – we have thus opted for the most up-to-date version.) Optimization and symbolic computation toolboxes of MATLAB are necessary. FAUST² automatically checks the presence of these packages and displays an error to the user in their absence. The software also takes the advantage of the MATLAB parallel computation toolbox if present. The use of parallel computation toolbox is currently disabled for Apple operating systems due to a conflict.

Please download FAUST² from Sourceforge. The files are organized in the main folder as follows: the sub-folder **Autonomous Models** contains the codes for deterministic systems (without input); the sub-folder **Controlled Models**

includes the codes for non-deterministic systems (input dependent); the sub-folder `Templates` contains templates and examples for the definition of symbolic conditional density functions; the sub-folder `Case Study` contains the files used in the next Section to test the software on a practical study. The file `README` can be opened with your preferred text editor and contains instructions on how to set up and run the software. Alternatively, `FAUST2` can be tested on a case study as elaborated in the next Section. Please set the current directory of `MATLAB` to the folder where the software is stored and run `FAUST2.m` from the `MATLAB` command line.

5 Case Study

In this section we apply `FAUST2` to compute optimal control strategies for the known room temperature regulation benchmark [15]. Probabilistic models for the underlying dynamics are based on [16] and on [2]. We consider the temperature regulation in multiple rooms via cooling water circulation. The amount of extracted heat is changed via a flow-control valve. Then the input signal is the percentage of the valve in the open position. The dynamics of the room temperature evolve in discrete time according to the equations

$$\begin{aligned} s_1(k+1) &= s_1(k) + \frac{\Delta}{C_{ra}}((s_2(k) - s_1(k))k_{cw}u(k) + (T_a - s_1(k))k_{out}) + \eta_{ra}(k), \\ s_2(k+1) &= s_2(k) + \frac{\Delta}{C_{cw}}((s_1(k) - s_2(k))k_{cw}u(k) + Q) + \eta_{cw}(k), \end{aligned} \quad (4)$$

where s_1 is the air temperature inside the room, s_2 is the cooling water temperature, T_a is the ambient temperature, Δ is the discrete sampling time [min], and $\eta_{ra}(\cdot), \eta_{cw}(\cdot)$ are stationary, independent random processes with normal distributions $\mathcal{N}(0, \sigma_{ra}^2 \Delta)$ and $\mathcal{N}(0, \sigma_{cw}^2 \Delta)$, respectively. Equations (4) can be encompassed in the condensed two-dimensional model

$$s(k+1) = f(s(k), u(k)) + \eta(k), \quad \eta(\cdot) \sim \mathcal{N}(0, \Sigma_\eta),$$

which results in a stochastic kernel that is a Gaussian conditional distribution $\mathcal{N}(f(s, u), \Sigma_\eta)$, where $\Sigma_\eta = \text{diag}(\Delta[\sigma_{ra}^2, \sigma_{cw}^2])$. The file `Chiller_Kernel_2d.m` appearing with the release of the software, provides numerical values and physical interpretations of the parameters in equations (4), as well as the symbolic structure of the conditional density function. The dynamical model in (4) can be as well extended to a two-room temperature control (which results in a three-dimensional model), and its conditional density function can be found in file `Chiller_Kernel_3d.m`. We will run `FAUST2` on both 2D and 3D setups.

We are interested in keeping the temperature of the room(s) within a given temperature interval over a fixed time horizon: this can be easily stated as a (probabilistic) safety problem, where we maximize over the feasible inputs to the model. We instantiate and compute this problem over the model above as described in the main text, while providing a step-by-step guide to the user.

In order to select the problem and import the model in FAUST², please follow these steps: select PCTL Safety in box ①, choose User-defined in box ②, tick the check-box ③ to indicate a controlled model, and write the name Chiller_Kernel_2d.m in the text of box ⑦ to load the density function of the two-dimensional model (4).

In the next stage we perform the abstraction and compute the quantity of interest (maximal safety probability). Select the most straightforward (but coarsest) abstraction algorithm, by choosing options Uniform gridding and Lipschitz via integral in ④. Proceed entering the problem parameters as follows: input the number of time steps as 3 and select a desired abstraction error equal to 0.5 in box ⑤; enter the safe temperature interval A as $[19.7, 20.3; 4.7, 5.3]$, as well as the input space \mathcal{U} as $[0, 1]$ in the text within box ⑧.

At this point the software can proceed with the main computations. Please press the button in box ⑥, in order to generate the abstract MDP, to compute the optimal policy and the related maximal safety probability. When the computation is complete, let us proceed with some post-processing: press the buttons Plot grid and Color grid in box ⑬, to generate Figure 2 (left) representing the maximal safety probability. The result of the computation can be stored for further analysis by pressing button ⑫: for instance Figure 2 (right) is generated by retrieving the optimal state-dependent Markov policy at step $N - 1$. The obtained abstract MDP has 144 states and 33 input actions. The experiment has been run in MATLAB 8.4 (R2014b) equipped with parallel computation toolbox on a 12-core Intel Xeon 3.47 GHz PC with 24 GB of memory and Windows 7 operating system. The initial estimate of the required time and the actual simulation time were 5.2 and 3.5 minutes, respectively. The simulation time includes both the generation of the abstract MDP and optimization over the input action.

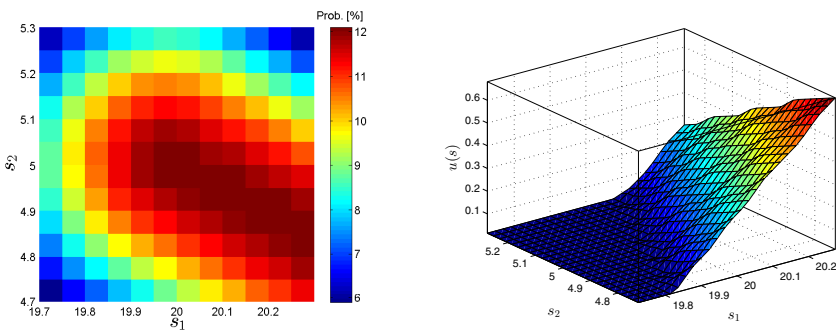


Fig. 2. Room temperature control problem. Left: obtained uniform partition of the safe set, along with optimal safety probability for each partition set (color bar on the right). The safety probability is equal to zero over the complement of the safe set. Right: optimal Markov policy at step $N - 1$, as a function of the model states.

A similar procedure can be followed to study the same probabilistic safety problem over a two-room temperature control, instantiated via the density function `Chiller_Kernel_3d.m`. Figure 3 presents the outcomes obtained using the Adaptive gridding and Lipschitz via integral options, selected in box ④. The abstraction parameters used in this problem is as follows: number of time steps 3, safe temperature interval $[19.5, 20.5; 19.5, 20.5; 4.5, 5.5]$, input space $[0, 1; 0, 1]$. We have selected a large abstraction error equal to 12 in box ⑤ to be able to visualize the adaptive grid generated by the software. The obtained abstract MDP has 45 states and 64 input actions. The experiment has been run on the same computer. The initial estimate of the required time and the actual simulation time were 20.7 and 26.2 minutes, respectively. The user can select a smaller error, at the likely cost of a larger computation time. For this case study, the implemented approach allows for the applicability of the abstraction technique at least to models with dimension 6 (that is, with 6 continuous variables), which is beyond the performance of currently available discretization-based approaches. The reader interested in detailed computational benchmarks for the presented techniques is referred to [8, 10].




Fig. 3. Two-room temperature control problem. Obtained partition of the safe set, together (bar) with optimal safety probability.

6 Summary of the Commands in the Graphical User Interface

We provide a summary of the commands of the GUI in FAUST², as they appear in the boxes highlighted in Figure 1.

- ① The box **Problem selection** provides a list with three options: select **Formula free** to obtain an abstraction of the model which can be exported to PRISM or to MRMC for further analysis; choose **PCTL Safety** in order to abstract the model and compute a safety probability; or opt for **PCTL Reach-Avoid** to get the abstraction tailored around the computation of the reach-avoid probability.
- ② The box **Kernel distribution** gives three options in a list: select **Linear Gaussian model** if the model belongs to the class of Linear Gaussian difference equations (cf. Section 1) and define matrices **A,B,Sigma** in the MATLAB workspace; choose **Non-linear Gaussian model** if the process noise is Gaussian and the drift and variance are non-linear (cf. Section 1), enter the drift and variance as a single symbolic function with two outputs via box ⑦; otherwise choose **User-defined** and enter your kernel as a symbolic function using ⑦.
- ③ Check this box if the model is non-deterministic (controlled).
- ④ Box **Gridding procedure** provides three options: select **Uniform gridding** to generate a grid based on global Lipschitz constant h (cf. Section 2), where the state space is partitioned uniformly along each dimension; choose **Adaptive gridding: local->local** to generate the grid adaptively based on local Lipschitz constants $h(i, j)$ (cf. Section 3), where the size of partition sets is smaller where the local error is higher; select **Adaptive gridding: local->global** to generate the grid adaptively based on local Lipschitz constants $h(i)$ (cf. [10]). The first option is likely to generate the largest number of partition sets and to be the fastest in the generation of the grid. The second option is likely to generate the smallest number of partition sets but to be the slowest in the grid generation. For the detailed comparison of these gridding procedures, please see [10].

The box **Assumptions on kernel** provides three choices: option **Lipschitz via integral** requires the density function $t_s(\bar{s}|s)$ to be Lipschitz continuous with respect to the current state s , and the quantity $T_p(z, z') = T_s(\Xi(z')|z)$ is used in the marginalization (integration) step; option **Lipschitz via sample** requires the density function $t_s(\bar{s}|s)$ to be Lipschitz continuous with respect to both current and the next states s, \bar{s} , and the quantity $T_p(z, z') = T_s(z'|z)\mathcal{L}(\Xi(z'))$ is used in the marginalization step; option **Max-Min** does not require any continuity assumption, but takes longer time in the computation of the error.

- ⑤ The time horizon of the desired PCTL formula or of the problem of interest, and the required upper bound on the abstraction error should be input in these two boxes. For the case of formula-free abstraction you may enter 1 as the number of time steps.
- ⑥ Press this button after entering the necessary data to generate the abstraction: this runs the main code. First, various checks are done to ensure the correctness of the inputted data. Then the partition sets are generated via gridding, the transition matrix is calculated, and the probability and the optimal policy are computed if applicable.
- ⑦ This box is activated for options **User-defined** and **Non-linear Gaussian model** in ②. For the first option, the conditional density function must be an m-file

that generates $t_{\bar{s}}(\bar{s}|s, u)$ symbolically. Please refer to `SymbolicKernel.m` for a template and `ExampleKernel.m` for an example. The name of kernel function should be entered in the text-box or the function should be loaded by pressing the button `Search for file`. For the option `Non-linear Gaussian model`, the non-linear drift and variance must be specified as a single symbolic function with two outputs. Please refer to `NonLinKernel.m` for a template and `NonLinKernelExample.m` for an example.

- ⑧ If the `Formula-free` option is selected in ①, the user can enter the bounds of the state space in the first of the boxes, named `Domain`. In case any of the additional two options in ① are selected, the boundaries of the safe set should be entered in the first text-box named `Safe set`. If the `PCTL Reach-Avoid` option in ① is selected, the second box is activated and the boundaries of the target set should be entered in the text-box named `Target set`. If the model is non-deterministic and the check in box ③ is ticked, the third box is also activated and the boundaries of the Input space may be entered in the box named `Input set`. In all cases the boundaries are to be given as a matrix with two columns, where the first and second columns contain lower and upper bounds, respectively. Alternatively, the user can press the `Select` button and separately enter the lower and upper bounds in the pop-up window.
- ⑨ The resulting error of the abstraction procedure, which is less than or equal to the desired abstraction error introduced in ⑤. This box shows the error associated to the abstracted model.
- ⑩ The user can add, remove, or edit labels associated to the abstract states. The set of states with any label $\alpha \in \Sigma$ can be represented by the intersection of half-planes $A_{\alpha}z \leq B_{\alpha}$. In order to tag these states with the associated label, the user presses button `Add` and subsequently enters symbol α and matrices A_{α}, B_{α} in a pop-up window. The user can also edit or remove any previously defined label by activating its symbol in the static-box and using buttons `Edit`, `Remove`. The button `States with selected label` will show the set of states associated with the active label in ⑬. Adding labels is essential in particular for exporting the result to PRISM or to MRMC.
- ⑪ The abstracted Markov chain or MDP can be exported to PRISM or to MRMC using these buttons. `FAUST2` enables two ways of exporting the result to PRISM: as a `.prism` format that is suitable for its GUI, or as the combination of `.tra` and `.sta` files, which are appropriate for the command line.
- ⑫ Use this button to store the results. A pop-up window appears after pushing the button and the user can opt for storing the date over the workspace, or in memory as an `.mat` file.
- ⑬ The user can plot the generated grid for the state space using the first button. Pressing this button opens a new window showing the partitioned input space for the controlled model. The solution of the safety and of the reach-avoid probability can also be visualized by pressing the second button. This option obviously works exclusively for dimensions $n = 1, 2, 3$.

- ⑭ The user can enter any initial state s_0 in the first box and calculate the safety or the reach-avoid probability of the model starting from that initial state, by pressing the button **Calculate**. The button **Properties of s_0** gives the abstracted state associated to s_0 , namely $z = \xi(\Xi(s_0))$ (cf. Algorithm 1), and all the labels assigned to this state.

7 Extensions and Outlook

There are a number of enticing extensions we are planning to work on, with the objective of rendering FAUST² widely useful and easily deployable.

FAUST² is presently implemented in MATLAB, which is the modeling software of choice in a number of engineering areas. We plan to improve part of its functionalities employing a faster, lower-level programming language, and to enhance the seamless integration with model checking tools.

Furthermore, we plan to extend the functionality of FAUST² by allowing for general label-dependent partitioning, and we are exploring the implementation with differently shaped partitioning sets [10]. We moreover plan to extend the applicability of FAUST² to models with discontinuous and degenerate [17, 18] kernels, to implement higher-order approximations [19], to include abstraction techniques specialized for stochastic max-plus linear systems [20], to embed formal truncations of the model dynamics [21], and to refine techniques and algorithms for non-deterministic (control-dependent) models. Finally, we plan to look into implementing in the tool bounds for infinite horizon properties, as currently investigated theoretically.

References

1. Meyn, S., Tweedie, R.: Markov chains and stochastic stability. Springer (1993)
2. Abate, A., Prandini, M., Lygeros, J., Sastry, S.: Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica* 44(11), 2724–2734 (2008)
3. Bertsekas, D., Shreve, S.: Stochastic Optimal Control: The Discrete-Time Case. Athena Scientific (1996)
4. Abate, A., Amin, S., Prandini, M., Lygeros, J., Sastry, S.: Computational approaches to reachability analysis of stochastic hybrid systems. In: Bemporad, A., Bicchi, A., Buttazzo, G. (eds.) HSCC 2007. LNCS, vol. 4416, pp. 4–17. Springer, Heidelberg (2007)
5. Summers, S., Lygeros, J.: Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem. *Automatica* 46(12), 1951–1961 (2010)
6. Abate, A., Katoen, J.P., Lygeros, J., Prandini, M.: Approximate model checking of stochastic hybrid systems. *European Journal of Control* 6, 624–641 (2010)
7. Abate, A., Katoen, J.P., Mereacre, A.: Quantitative automata model checking of autonomous stochastic hybrid systems. In: ACM Proceedings of the 14th International Conference on Hybrid Systems: Computation and Control, Chicago, IL, pp. 83–92 (2011)

8. Esmail Zadeh Soudjani, S., Abate, A.: Adaptive gridding for abstraction and verification of stochastic hybrid systems. In: Proceedings of the 8th International Conference on Quantitative Evaluation of Systems, pp. 59–69 (September 2011)
9. Tkachev, I., Mereacre, A., Katoen, J., Abate, A.: Quantitative automata-based controller synthesis for non-autonomous stochastic hybrid systems. In: Proceedings of the 16th international conference on Hybrid Systems: Computation and Control, HSCC 2013, pp. 293–302 (2013)
10. Esmail Zadeh Soudjani, S., Abate, A.: Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes. *SIAM Journal on Applied Dynamical Systems* 12(2), 921–956 (2013)
11. Tkachev, I., Abate, A.: Formula-free Finite Abstractions for Linear Temporal Verification of Stochastic Hybrid Systems. In: Proceedings of the 16th International Conference on Hybrid Systems: Computation and Control, Philadelphia, PA, pp. 283–292 (April 2013)
12. Hinton, A., Kwiatkowska, M., Norman, G., Parker, D.: PRISM: A tool for automatic verification of probabilistic systems. In: Hermanns, H., Palsberg, J. (eds.) TACAS 2006. LNCS, vol. 3920, pp. 441–444. Springer, Heidelberg (2006)
13. Katoen, J.P., Khattri, M., Zapreev, I.S.: A Markov reward model checker. In: IEEE Proceedings of the International Conference on Quantitative Evaluation of Systems, Los Alamos, CA, USA, pp. 243–244 (2005)
14. Hansson, H., Jonsson, B.: A logic for reasoning about time and reliability. *Formal Aspects of Computing* 6(5), 512–535 (1994)
15. Fehnker, A., Ivančić, F.: Benchmarks for hybrid systems verification. In: Alur, R., Pappas, G.J. (eds.) HSCC 2004. LNCS, vol. 2993, pp. 326–341. Springer, Heidelberg (2004)
16. Malhamé, R., Chong, C.: Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. *IEEE Transactions on Automatic Control* 30(9), 854–860 (1985)
17. Esmail Zadeh Soudjani, S., Abate, A.: Probabilistic invariance of mixed deterministic-stochastic dynamical systems. In: ACM Proceedings of the 15th International Conference on Hybrid Systems: Computation and Control, Beijing, PRC, pp. 207–216 (April 2012)
18. Esmail Zadeh Soudjani, S., Abate, A.: Probabilistic reach-avoid computation for partially-degenerate stochastic processes. *IEEE Transactions on Automatic Control* 59(2), 528–534 (2014)
19. Esmail Zadeh Soudjani, S., Abate, A.: Higher-Order Approximations for Verification of Stochastic Hybrid Systems. In: Chakraborty, S., Mukund, M. (eds.) ATVA 2012. LNCS, vol. 7561, pp. 416–434. Springer, Heidelberg (2012)
20. Adzkiya, D., Esmail Zadeh Soudjani, S., Abate, A.: Finite Abstractions of Stochastic Max-Plus-Linear Systems. In: Norman, G., Sanders, W. (eds.) QEST 2014. LNCS, vol. 8657, pp. 74–89. Springer, Heidelberg (2014)
21. Esmail Zadeh Soudjani, S., Abate, A.: Precise Approximations of the Probability Distribution of a Markov Process in Time: An Application to Probabilistic Invariance. In: Abraham, E., Havelund, K. (eds.) TACAS 2014 (ETAPS). LNCS, vol. 8413, pp. 547–561. Springer, Heidelberg (2014)