



# A Flexible Toolchain for Symbolic Rabin Games under Fair and Stochastic Uncertainties

Rupak Majumdar<sup>1</sup>, Kaushik Mallik<sup>2(✉)</sup>, Mateusz Rychlicki<sup>3</sup>,  
Anne-Kathrin Schmuck<sup>1</sup>, and Sadegh Soudjani<sup>4</sup>



<sup>1</sup> MPI-SWS, Kaiserslautern, Germany  
{rupak, akschmuck}@mpi-sws.org

<sup>2</sup> ISTA, Klosterneuburg, Austria  
kaushik.mallik@ist.ac.at

<sup>3</sup> School of Computing, University of Leeds, Leeds, UK  
scmkry@leeds.ac.uk

<sup>4</sup> Newcastle University, Newcastle upon Tyne, UK  
Sadegh.Soudjani@newcastle.ac.uk

**Abstract.** We present a flexible and efficient toolchain to *symbolically* solve (standard) Rabin games, fair-adversarial Rabin games, and  $2^{1/2}$ -player Rabin games. To our best knowledge, our tools are the first ones to be able to solve these problems. Furthermore, using these flexible game solvers as a back-end, we implemented a tool for computing correct-by-construction controllers for stochastic dynamical systems under LTL specifications. Our implementations use the recent theoretical result that all of these games can be solved using the same symbolic fixpoint algorithm but utilizing different, domain specific calculations of the involved predecessor operators. The main feature of our toolchain is the utilization of two programming abstractions: one to separate the symbolic fixpoint computations from the predecessor calculations, and another one to allow the integration of different BDD libraries as back-ends. In particular, we employ a multi-threaded execution of the fixpoint algorithm by using the multi-threaded BDD library Sylvan, which leads to enormous computational savings.

## 1 Introduction

Piterman and Pnueli [17] derived the currently best known symbolic algorithm for solving two-player Rabin games over finite graphs with a theoretical complexity of  $O(n^{k+1}k!)$  in time and space, where  $n$  is the number of states and  $k$  is the number of pairs in the winning condition. This work did not provide an

---

Authors ordered alphabetically. R. Majumdar and A.-K. Schmuck are partially supported by DFG project 389792660 TRR 248-CPEC. A.-K. Schmuck is additionally funded through DFG project (SCHM 3541/1-1). K. Mallik is supported by the ERC project ERC-2020-AdG 101020093. M. Rychlicki is supported by the EPSRC project EP/V00252X/1. S. Soudjani is supported by the following projects: EPSRC EP/V043676/1, EIC 101070802, and ERC 101089047.

© The Author(s) 2023

C. Enea and A. Lal (Eds.): CAV 2023, LNCS 13966, pp. 3–15, 2023.

[https://doi.org/10.1007/978-3-031-37709-9\\_1](https://doi.org/10.1007/978-3-031-37709-9_1)

implementation. In a series of papers [3, 4, 15, 16], Mallik et al. showed that this symbolic algorithm can be extended to solve different automated design questions for reactive hardware, software, and cyber-physical systems under fair or stochastic uncertainties. The main contribution of their work is to show that these extensions only require a very mild syntactic change of the Piterman-Pnueli fixed-point algorithm (with very little effect on its overall complexity) and domain-specific realizations of two types of predecessor operators used therein.

Using this insight, we present a *toolchain* for the *efficient symbolic solution of different extensions of Rabin games*. We have created three inter-connected libraries for solving different parts of the problem from different levels of abstraction. The first library, called **Genie**, offers a set of virtual classes to implement the fixpoint algorithm—abstractly, leaving open (i.e. virtual) the predecessor computation. Alongside, we created two other libraries, called **FairSyn** and **Mascot-SDS**, where **FairSyn** solves fair-adversarial [4] and  $2^{1/2}$ -player Rabin games [3], while **Mascot-SDS** solves abstraction-based control problems [15, 16]. **FairSyn** and **Mascot-SDS** use the optimized fixpoint computation provided by **Genie**, with domain specific implementations of the predecessor operations.

The flexibility of our toolchain comes from two different programming abstractions in **Genie**. Firstly, **Genie** offers multiple high-level optimizations for solving the Rabin fixpoint, such as parallel execution (requires a thread-safe BDD library like **Sylvan**) and an acceleration technique [13], while abstracting away from the low-level implementations of the predecessor functions. As a result, any synthesis problem using the core Rabin fixpoint of **Genie** can use the optimizations without spending any extra implementation effort. We used these optimizations from **FairSyn** and **Mascot-SDS**, and achieved remarkable computational savings. Secondly, **Genie** offers easy portability of codes from one BDD library to another, which is important as different BDD libraries have different pros and cons, and the choice of the best library depends on the needs. We empirically showed how switching between the two BDD libraries **Sylvan** and **CUDD** impacts the performance of **FairSyn** and **CUDD**: overall, the **Sylvan**-based experiments were significantly faster, whereas the **CUDD**-based experiments consumed considerably lower amount of memory. Using the combined power of multi-threaded BDD operations using **Sylvan** and the optimizations offered by **Genie**, **Mascot-SDS** was between one and three orders of magnitude faster than the state-of-the-art tool in our experiments.

**Comparison with Existing Tools:** We are not aware of any available tool to directly solve (normal or stochastic) Rabin games *symbolically*. However, it is well-known how to translate *stochastic* Rabin games into (standard) Rabin games [5], and Rabin games into parity games, for which efficient solvers exist, e.g. **oink** [9]. Yet, efficient solutions of stochastic Rabin games via parity games are difficult to obtain, because: (i) the translation from a stochastic Rabin game to a Rabin game involves a quadratic blow-up, and the translation from a Rabin game to a parity game results in an exponential blow-up in the size of the game, (ii) symbolic fixpoint computations become cumbersome very fast for parity games, as the number of vertices and/or colors in the game graph increases, leading to high computation times in practice, and (iii) the only known algorithms capable of handling fair

and stochastic uncertainties efficiently are all *symbolic* in nature, while most of the efficient parity game solvers are non-symbolic. Additionally, unlike the Rabin fixpoint, the nesting of the parity fixpoint does not enable parallel execution.

While it is well known that for normal parity games, computational tractability can be achieved by different non-symbolic algorithms, such as Zielonka’s algorithm [22], tangle learning [8] or strategy-improvement [19], implemented in `oink` [9], it is currently unclear if and how these algorithms allow for the efficient handling of fair or stochastic uncertainties. We are therefore unable to compare our toolchain to the translational workflow via parity games in a fair manner.

In the area of temporal logic control of stochastic systems, `Mascot-SDS` has two powerful features: (a) it can handle synthesis for the rich class of omega-regular (infinite-horizon) specifications, and (b) it provides both over- and under-approximations of the solution, thus enabling a quantitative refinement loop for improving the precision of the approximation. The features of `Mascot-SDS` is compared with other tools in the stochastic category of the recent ARCH competition (see the report [1] for the list of participating tools). As concluded in the report of the competition, other state-of-the-art tools in stochastic category are either limited to a fragment of  $\omega$ -regular specifications or do not provide any indication of the quality of the involved approximations. The only tool [10] that supports  $\omega$ -regular specifications uses a different alternate non-symbolic approach, against which `Mascot-SDS` fares significantly well in our experiments (see Sect. 4.2). Even if we leave stochasticity aside, our tool implements a new and orthogonal heuristic for multi-threaded computation of Rabin fixpoints, which is not considered by other controller synthesis tools [11].

## 2 Theoretical Background

We briefly state the synthesis problems our toolchain is solving. We follow the same (standard) notation for two-player game graphs, winning regions, strategies and  $\mu$ -calculus formulas, as in [4].

### 2.1 Solving Rabin Games Symbolically

Given a game graph  $G = (V, V_0, V_1, E)$ , a Rabin game is specified using a set of Rabin pairs  $\mathcal{R} = \{(Q_1, R_1), \dots, (Q_k, R_k)\}$ , with  $Q_i, R_i \subseteq V$  for every  $i \in [1; k]$ , and  $\varphi := \bigvee_{i \in [1; k]} (\diamond \square \neg R_i \wedge \square \diamond Q_i)$  being the Rabin acceptance condition. Piterman and Pnueli [17] showed that the winning region of a Rabin game can be computed using the  $\mu$ -calculus expression given in (2), where the set transformers  $Cpre : 2^V \rightarrow 2^V$  and  $Apre : 2^V \times 2^V \rightarrow 2^V$  are defined for every  $S, T \subseteq V$  as:

$$Cpre(S) := \{v \in V_0 \mid \exists v' \in S . (v, v') \in E\} \\ \cup \{v \in V_1 \mid \forall v' \in V . (v, v') \in E \implies v' \in S\}, \quad (1a)$$

$$Apre(S, T) := Cpre(T). \quad (1b)$$

**Fair-Adversarial Rabin Games.** A Rabin game is called *fair-adversarial* when there is an additional fairness assumption on a set of edges originating from

The symbolic fixpoint algorithm for solving Rabin games with  $\mathcal{R} = \{(Q_1, R_1), \dots, (Q_k, R_k)\}$  and  $K = [1; k]$ :

$$\nu Y_{p_0} \cdot \mu X_{p_0} \cdot \bigcup_{p_1 \in K} \nu Y_{p_1} \cdot \mu X_{p_1} \cdot \bigcup_{p_2 \in K \setminus \{p_1\}} \nu Y_{p_2} \cdot \mu X_{p_2} \cdot \dots \bigcup_{p_k \in K \setminus \{p_1, \dots, p_{k-1}\}} \nu Y_{p_k} \cdot \mu X_{p_k} \cdot \left[ \bigcup_{j=0}^k C_{p_j} \right], \quad (2)$$

where

$$C_{p_j} := \left( \bigcap_{i=0}^j \bar{R}_{p_i} \right) \cap \left[ (Q_{p_j} \cap Cpre(Y_{p_j})) \cup (Apre(Y_{p_j}, X_{p_j})) \right],$$

and the definitions of  $Cpre$  and  $Apre$  are problem specific.

*Player 1* vertices in  $G$ . Let  $E^\ell \subseteq E \cap (V_1 \times V)$  be a given set of edges, called the *live* edges. Given  $E^\ell$  and a Rabin winning condition  $\varphi$ , we say that *Player 0* wins the *fair-adversarial Rabin game* from a vertex  $v$  if *Player 0* wins the (normal) game for the modified winning condition  $\varphi^\ell := \left( \bigwedge_{e=(v,v') \in E^\ell} (\Box \diamond v \implies \Box \diamond e) \right) \implies \varphi$ . Based on the results of Banerjee et al. [4], fair-adversarial Rabin games can be solved via (2), by defining for every  $S, T \subseteq V$

$$Cpre(S) := \{v \in V_0 \mid \exists v' \in S . (v, v') \in E\} \\ \cup \{v \in V_1 \mid \forall v' \in V . (v, v') \in E \implies v' \in S\}, \quad (3a)$$

$$Apre(S, T) := Cpre(T) \cup \{v \in Cpre(S) \cap V_1 \mid \exists v' \in T . (v, v') \in E^\ell\}. \quad (3b)$$

We see that (3) coincides with (1) if  $E^\ell$  is empty.

**2<sup>1/2</sup>-Player Rabin Games.** A 2<sup>1/2</sup>-player game is played on a game graph  $(V, V_0, V_1, V_r, E)$ , and the only difference from a 2-player game graph is the additional set of vertices  $V_r$  which are called the *random* vertices. The sets  $V_1, V_2$ , and  $V_r$  partition  $V$ . Based on the results of [3] 2<sup>1/2</sup>-Player rabin games can be solved via (2) by defining for all  $S, T \subseteq V$

$$Cpre(S) := \{v \in V_0 \mid \exists v' \in S . (v, v') \in E\} \\ \cup \{v \in V_1 \cup V_r \mid \forall v' \in V . (v, v') \in E \implies v' \in S\}, \quad (4a)$$

$$Apre(S, T) := Cpre(T) \cup \{v \in Cpre(S) \cap V_r \mid \exists v' \in T . (v, v') \in E\}. \quad (4b)$$

## 2.2 Computing Symbolic Controllers for Stochastic Dynamical Systems

A discrete-time stochastic dynamical system  $S$  is represented using a tuple  $(X, U, W, f)$ , where  $X \subseteq \mathbb{R}^n$  is a *continuous* state space,  $U$  is a *finite* set of control inputs,  $W \subset \mathbb{R}^n$  is a *bounded* set of disturbances, and  $f: X \times U \rightarrow X$  is the nominal dynamics. If  $x^k \in X$  and  $u^k \in U$  are the state and control input of  $S$  at some time  $k \in \mathbb{N}$ , then the state at the next time step is given by:

$$x^{k+1} = f(x^k, u^k) + w^k, \quad (5)$$

where  $w^k$  is the disturbance at time  $k$  which is sampled from  $W$  using some (possibly unknown) distribution. Without loss of generality we assume that  $W$

is centered around the origin, which can be easily achieved by shifting  $f$  if needed. A *path* of  $S$  originating at  $x^0 \in X$  is an infinite sequence of states  $x^0 x^1 \dots$  for a given infinite sequence of control inputs  $u^0 u^1 \dots$ , such that (5) is satisfied.

Let  $\varphi$  be a given Rabin specification—called the *control objective*—defined using a finite set of predicates over  $X$ . For every controller  $C: X \rightarrow U$ , the domain of  $C$ , written  $Dom(C)$ , is the set of states from where the property  $\varphi$  can be satisfied with probability 1. For a fixed  $\varphi$ , a controller  $\hat{C}$  is called *optimal* if  $Dom(\hat{C})$  contains the domain of every other controller  $C$ . The problem of computing such an optimal controller for the system in (5) is in general undecidable. Following [15], we compute an approximate solution instead.

This approximate solution is obtained by a discretization of the state space. For this, we assume that the state space  $X$  is a closed and bounded subset of the  $n$ -dimensional Euclidean space  $\mathbb{R}^n$  for some  $n > 0$ , and use the notation  $[[a, b))$  to denote the set  $\prod_{i \in [1:n]} [a_i, b_i)$ . Now, consider a grid-based discretization  $\hat{X}$  of  $X$ , where  $\hat{X} = \{[[a, b)) \mid a, b \in \mathbb{R}^n = X\}$ . One of the key ingredients of our abstraction process is a function  $\hat{f}$  providing hyper-rectangular over-approximation of the one-step reachable set of the nominal dynamics  $f$  of the system  $S$ : for every grid element  $\hat{x} \in \hat{X}$ , we have  $\hat{f}(\hat{x}, u) = [[a', b')) \supseteq \{x' \in X \mid \exists x \in \hat{x}. x' = f(x, u)\}$ . The function  $\hat{f}$  is known to be available for a wide class of commonly used forms of the function  $f$ , and in our implementation we assumed that  $f$  is mixed-monotone and  $\hat{f}$  is the so-called decomposition function (see standard literature for details [7]).

Given the over-approximation of the nominal dynamics obtained through  $\hat{f}$ , we define, respectively, the over- and the under-approximation of the *perturbed* dynamics as  $\bar{g}(\hat{x}, u) := W \oplus \hat{f}(\hat{x}, u)$  and  $\underline{g}(\hat{x}, u) := W \ominus (-\hat{f}(\hat{x}, u))$ , where  $\oplus$  and  $\ominus$  respectively denote the Minkowski sum and the Minkowski difference. Next, we transfer  $\bar{g}$  and  $\underline{g}$  to the abstract state space  $\hat{X}$  to obtain, respectively, the over- and the under-approximation in terms of the *abstract transition function*<sup>1</sup>, i.e.,  $\bar{h}(\hat{x}, u) := \{\hat{x}' \in \hat{X} \mid \bar{g}(\hat{x}, u) \cap \hat{x}' \neq \emptyset\}$  and  $\underline{h}(\hat{x}, u) := \{\hat{x}' \in \hat{X} \mid \underline{g}(\hat{x}, u) \cap \hat{x}' \neq \emptyset\}$ . With  $\bar{h}$  and  $\underline{h}$  available, it was shown by Majumdar et al. [16] that the over-approximation of the optimal controller can be solved by using the fixpoint algorithm in (2), where the predecessor operators are defined for every  $S, T \subseteq \hat{X}$  as

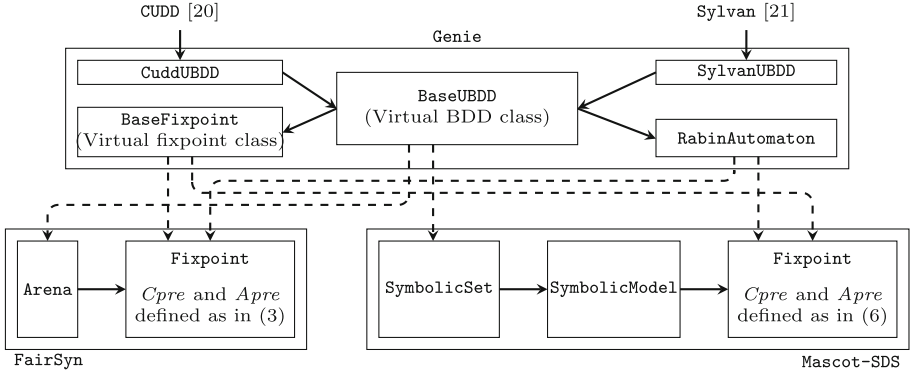
$$Cpre(S) := \{\hat{x} \in \hat{X} \mid \exists u \in U. \bar{h}(\hat{x}, u) \subseteq S\} \quad (6a)$$

$$Apr(S, T) := \{\hat{x} \in \hat{X} \mid \exists u \in U. \bar{h}(\hat{x}, u) \subseteq S \wedge \underline{h}(\hat{x}, u) \cap T \neq \emptyset\}. \quad (6b)$$

### 3 Implementation Details

We develop three interconnected tools, **Genie**, **FairSyn**, and **Mascot-SDS**, which work in close harmony to implement efficient solvers for the solution of (2) with

<sup>1</sup> Here we assume that  $\hat{f}(\hat{x}, u) \subseteq X$ ; otherwise we need to take some extra steps. Details can be found in the work by Majumdar et al. [16].



**Fig. 1.** A schematic diagram of interaction among the three tools. Each block represents one class in the respective tool, and an arrow from class A to class B denotes that B depends on A. The dependency within each tool is shown using solid arrows, while the dependencies of *Mascot-SDS* and *FairSyn* on *Genie* is shown using dashed arrows.

pre-operators defined via (3), (4) and (6), respectively. The tools use binary decision diagrams (BDD) to symbolically manipulate sets of vertices/states of the underlying system, and to manage the BDDs, we offer the flexibility to choose between two of the well-known existing BDD libraries, namely CUDD [20] and Sylvan [21]. The two libraries have their own merits: while CUDD has significantly lower memory footprint, Sylvan offers superior computation speed through multi-threaded BDD operations. Thus, the optimal choice of the library depends on the size of the problem, the computational time limit, and the memory budget, and through our implementation it is possible to choose one or the other by, in some cases, changing only a single line of code and, in the other cases, changing the value of just one flag. Moreover, we expect that integrating other BDD libraries having the same basic BDD operations in our tools will be easy and seamless—thanks to the programming abstraction offered by *Genie*. Such extensions will possibly bring more diverse set of computational strengths for solving the fundamental synthesis problems that we address.

The tools are primarily written using C++, with some small python scripts implementing parts of visualizations of outputs. The main classes of the three tools and their interactions are depicted in Fig. 1. We briefly describe the core functionalities of the tools in the following.

### 3.1 Genie

*Genie* implements the fixpoint algorithm (2) in the class *BaseFixpoint* through two layers of abstraction. One abstraction is through the virtual definitions of the *Cpre* and *Apre* operators, whose concrete implementations are provided in the front-end synthesis tools (in our case *FairSyn* and *Mascot-SDS*). Using this abstraction, we implemented two different optimizations for the efficient iterative computation of the Rabin fixpoint in (2)—independently from the actual implementations of the *Apre* and *Cpre* operators. The first optimization is a

multi-threaded computation of the Rabin fixpoint, exploiting the fixpoint’s inherent parallel structure due to the independence among different sequences of  $(p_1, p_2, \dots)$  used to compute  $\bigcup_{j=0}^k C_{p_j}$ . The second optimization is an accelerated computation of the Rabin fixpoint, achieved through bookkeeping of intermediate values of the BDD variables. The core of the acceleration procedure for general  $\mu$ -calculus fixpoints was proposed by Long et al. [13], and the details specific to the fixpoint in (2) can be found in the paper by Banerjee et al. [4].

The other abstraction in **Genie** is the set of virtually defined low-level BDD operations in the auxiliary class **BaseUBDD**, which enable us to easily switch between different off-the-shelf BDD libraries. The virtual BDD operations in **BaseUBDD** are concretely realized in the classes **CuddUBDD** and **SylvanUBDD**, which work as interfaces between, respectively, the CUDD and the Sylvan BDD libraries. Support for additional BDD libraries can be easily built by creating new interface classes. More details on the functionalities of **Genie** can be found in the longer version of this paper [14].

### 3.2 FairSyn

The core of **FairSyn** is written as a header-only library, which offers the infrastructure to solve (2) with pre-operators defined via (3) and (4). The main component of **FairSyn** is the class **Fixpoint**, which derives from the class **BaseFixpoint** from **Genie**, and implements the concrete definitions of *Cpre* and *Apr* in (3) and (4).

**How to Use:** For computing the winning region and the winning strategy in a fair-adversarial Rabin game (resp. a  $2^{1/2}$ -player Rabin game) using **FairSyn**, one needs to write a program to create the game as a **Fixpoint** object. One possible way of constructing a **Fixpoint** object is through a synchronous product of a game graph (an object of class **Arena**) and a specification Rabin automaton (an object of class **RabinAutomaton**) with an input alphabet of sets of nodes of the **Arena** object. Following is a snippet:

---

```
// typedef Genie::CuddUBDD UBDD; // use this for CUDD
typedef Genie::SylvanUBDD UBDD; // use this for Sylvan
UBDD base;
...
Arena<UBDD> A(base, vars, nodes, sys_nodes, env_nodes, edges,
    live_edges); // the game graph
RabinAutomaton<UBDD> R(base, vars, inp_alphabet, filename); // the
    specification automaton
Fixpoint<UBDD> Fp(base, "under", A, R); // the synchronous product
// UBDD strategy = Fp.Rabin(true, 20, Fp.nodes_, 0); // sequential
    fixpoint solver
UBDD strategy = Fp.Rabin(true, 20, Fp.nodes_, 0,
    Genie::ParallelRabinRecurse); // parallel fixpoint solver
...

```

---

where `vars` is a (possibly initially empty) set of integers which will contain the set of newly created BDD variables, `nodes`, `sys_nodes`, and `env_nodes` are, respectively, vectors of indices of various types of vertices, `edges` and `live_edges` are, respectively, vectors of the respective types of edges, `inp_alphabet` is a `std::map` object that maps input symbols of the Rabin automaton to the respective BDDs representing sets of nodes in the `Arena`, and `filename` is the name of the file in which the Rabin automaton is stored (using the standard HOA format [2]). The game is solved by calling `Fp.Rabin`, a member function of the `Genie::BaseFixpoint` class (see Sect. 3.1).

### 3.3 Mascot-SDS

The core of `Mascot-SDS` is also written as a header-only library. It is built on top of the well-known tool called `SCOTS` [18], with several classes of `Mascot-SDS` still retaining their original identities from `SCOTS`, owing to the close similarity of the basic uniform grid-based abstraction used in both tools. The main difference between the two tools is that `Mascot-SDS` synthesizes controllers for *stochastic* systems, while `SCOTS` synthesizes controllers for only *non-stochastic* systems.

The two main classes of `Mascot-SDS` are called `SymbolicSet` and `SymbolicModel`, which respectively model the abstract spaces obtained through uniform grid-based discretizations (like  $\hat{X}$  in Sect. 2.2) and the abstract transition relations ( $\hat{h}$  and  $\underline{h}$  in Sect. 2.2). The abstract transition relations are computed using an auxiliary class called `SymbolicModelMonotonic` (not shown in Fig. 1). Notice that we offer the flexibility to use both `CUDD` and `Sylvan` while creating objects from `SymbolicSet` and `SymbolicModel`. A `Fixpoint` object is a child of the class `BaseFixpoint` from `Genie`, which is created by taking a synchronous product between a `SymbolicModel` object and a `RabinAutomaton` object specifying the control objective given as user input. The class `Fixpoint` implements the concrete definitions of the *Cpre* and *Apr*e operator according to (6).

**How to Use:** For ease of use, we have written a pair of tools called `Synthesize` and `Simulate` using the library of `Mascot-SDS`. `Synthesize` synthesizes controllers for stochastic dynamical systems whose nominal dynamics is mixed-monotone, and `Simulate` visualizes simulated closed-loop trajectories using the synthesized controller. The inputs to `Synthesize` include the dynamic model of the system and the control objective; the latter can be specified either in LTL or using a Rabin automaton. To use `Synthesize`, simply use the following syntax:

---

```
<path-to-Synthesize binary>/Synthesize <path-to-input-file>/<input.cfg>
    <sylvan/cudd flag>
```

---

where the `<input.cfg>` is an input configuration file containing all the inputs, and the `<sylvan/cudd flag>` is either 1 or 0 depending on whether the parallel version using `Sylvan` is to be run or the sequential version using `CUDD`.

Some of the main ingredients in the `input.cfg` file are: (a) the description of the dynamical system's variable spaces (like state space, input space,



etc.) including their discretization parameters, (b) the file where the decomposition function of the nominal dynamics of the system is stored, (c) the absolute value of maximum disturbance, and (d) the specification either as an LTL formula or as the filename where a Rabin automaton is stored (in HOA format [2]). The decomposition function is required to be given as a C-compatible header file so that `Synthesize` can link to (use) this function at runtime (see the `mascot-sds/examples/` directory for examples). When the specification is given as a Rabin automaton (over a labeling alphabet of the system states), the automaton needs to be stored in a file in the HOA format. Alternatively, an LTL specification can be given, along with a mapping between the atomic predicates and the states of the system. In that case `Synthesize` uses `Owl` [12] to convert the LTL specification to a Rabin automaton.

The output of `Synthesize` is a folder called `data` that contains pieces of the controller encoded in BDDs and stored in binary files as well as various metadata information stored in text files. These files can be processed by `Simulate` to visualize simulated closed-loop trajectories of the system. The usage of `Simulate` is similar to `Synthesize`:

---

```
<path-to-Simulate binary>/Simulate <path-to-input-file>/<input.cfg>
    <sylvan/cudd flag>
```

---

where the `input.cfg` file should, in this case, contain information that are required to simulate the closed-loop, like simulation time steps, the python script that will plot the state space predicates (see the examples), etc.

## 4 Examples

We present experimental results, showcasing practical usability of our tools and comparing performances with the state of the art. All the experiments were run on a computer with Intel Xeon E7-8857 v2 48 core processor and 1.5 TB RAM.

### 4.1 Synthesizing Code-Aware Resource Mangers Using FairSyn

We consider a case study introduced by Chatterjee et al. [6]. In this example, there are two bounded FIFO queues, namely the broadcast and output queues, which interact among each other and transmit and receive data packets through a common network. The two queues are implemented using separate threads running on a single CPU. For this multi-threaded program, we consider the problem of synthesizing a code-aware resource manager, whose task is to grant different threads accesses to different shared synchronization resources (mutexes and counting semaphores). The specification is deadlock freedom across all threads at all time while assuming a fair scheduler (scheduling every thread always eventually) and fair progress in every thread (i.e., taking every existing execution branch always eventually). The resource-manager is code-aware, and has knowledge about the require and release characteristics of all threads for different resources. This enables us to avoid deadlocks more effectively than the

case when the resource-manager does not have access to the code. Chatterjee et al. [6] showed that the synthesis problem (of the resource manager) can be reduced to the problem of computing the winning strategy in a  $2^{1/2}$ -player game, which we solved using FairSyn.

Table 1 compares the computational resources for the CUDD and Sylvan-based implementations of FairSyn; more details can be found in our earlier work [4]. It can be observed that the Sylvan-based implementation is significantly faster, although it consumes much more memory.

**Table 1.** Performance of FairSyn; code-aware resource management benchmark.

Broadcast and Output Queue Capacities	Number of BDD variables	Computation Time (seconds)		Peak Memory Usage	
		CUDD	Sylvan	CUDD	Sylvan
(1, 1)	25	255.33	11.40	292 MiB	671 MiB
(2, 1)	27	957.99	29.20	310 MiB	681 MiB
(3, 1)	27	903.01	31.13	310 MiB	973 MiB
(1, 2)	27	1308.09	39.57	315 MiB	682 MiB
(1, 3)	27	1249.37	41.76	309 MiB	681 MiB
(2, 2)	29	5127.93	111.62	342 MiB	685 MiB
(3, 2)	29	5104.20	114.30	339 MiB	975 MiB
(2, 3)	29	5644.09	118.12	341 MiB	975 MiB
(3, 3)	29	6156.57	137.56	339 MiB	975 MiB

## 4.2 Synthesizing Controllers for Stochastic Dynamical Systems Using Mascot-SDS

We use Mascot-SDS to synthesize controllers for two different applications.

**A Bistable Switch.** First, we compare our tool’s performance against the state-of-the-art tool called StochasticSynthesis (abbr. SS) [10] on a benchmark example that was proposed by the authors of SS. In this example, there is a 2-dimensional nonlinear bistable switch that is perturbed with bounded stochastic noise. There are two synthesis problems with two different control objectives: one, a safety objective, and, two, a Rabin objective with two Rabin pairs. The model of the system and the control objectives can be found in the original paper [10].


The tool SS uses graph theoretic techniques to solve the controller synthesis problem, which is an alternative approach that is substantially different from our symbolic fixpoint based technique. In Table 2, we summarize the performance of Mascot-SDS powered by CUDD and Sylvan, alongside the performance of SS. Both

**Table 2.** Performance comparison between *Mascot-SDS* and *StochasticSynthesis* (abbreviated as *SS*) [10] on the bistable switch. Col. 1 shows the specifications and the respective numbers of Rabin pairs, Col. 2 shows the approximation error ranges (smaller error means more intense computation), Col. 3, 4, and 5 Col. 6, 7, and 8 compare the peak memory footprint (as measured using the “time” command) for *Mascot-SDS* with *CUDD*, *Mascot-SDS* with *Sylvan*, and *SS* respectively. “TO” stands for timeout (5 h of cutoff time).

Spec.	upper bound on approx. error	Total running time			Peak memory footprint		
		<i>Mascot-SDS</i>		<i>SS</i> [10]	<i>Mascot-SDS</i>		<i>SS</i> [10]
		<i>CUDD</i>	<i>Sylvan</i>		<i>CUDD</i>	<i>Sylvan</i>	
$\varphi_1$ (1 Rabin pair)	20%–30%	11 s	<2 s	27 s	351 MiB	79 MiB	223 MiB
	10%–20%	9 s	2 s	43 s	351 MiB	105 MiB	290 MiB
	5%–10%	14 s	4 s	1 h 49 min	405 MiB	251 MiB	25 GiB
	0%–5%	48 s	10 s	TO	553 MiB	759 MiB	TO
$\varphi_2$ (2 Rabin pairs)	20%–30%	21 s	<2 s	21 s	324 MiB	40 MiB	202 MiB
	10%–20%	26 s	2 s	25 s	371 MiB	80 MiB	203 MiB
	5%–10%	37 s	4 s	1 min 17 s	436 MiB	242 MiB	490 MiB
	0%–5%	2 min 24 s	13 s	TO	573 MiB	761 MiB	TO

**Table 3.** Performance of *Mascot-SDS* with *CUDD* and *Sylvan* for the table-serving robot experiment.

	<i>CUDD</i>	<i>Sylvan</i>
Comp. time	1 h 3 min	2 min 55 s
Peak memory	673 MiB	1.1 GiB



**Fig. 2.** Closed-loop trajectories for 100 time steps with *kitchen* (green), *table* (blue), and *obstacle* (black). (Color figure online)

*Mascot-SDS* and *SS* compute controllers whose domains under-approximate the optimal controller domains. The second column of Table 2 shows a measure of the approximation error. For every comparable approximation error bound, both versions of *Mascot-SDS* significantly outperformed *SS*, both time and memory-wise. In fact, *Mascot-SDS* with *Sylvan* was at least an order of magnitude faster in all instances. This is particularly astonishing, since *SS* uses a sophisticated *lazy* abstraction refinement technique, whereas *Mascot-SDS* uses a plain *uniform* abstraction which is typically computationally expensive. This shows the immense potential of our toolchain; we plan to extend *Mascot-SDS* with *lazy* gridding, an orthogonal optimization, in a future release to make further computational savings. For *Mascot-SDS* itself, as expected, *Sylvan* was significantly faster than *CUDD*. On the other hand, though *Sylvan* used less memory than *CUDD* in the simpler setups (the ones with more error), the memory requirement of *Sylvan* quickly grew and surpassed that of *CUDD* for the more complicated setup.

**Table-Serving Robot.** We consider the controller synthesis problem for a table-serving robot that needs to satisfy the following specification:  $\Box\Diamond kitchen \wedge \Box\neg obstacle \wedge (\Box\Diamond request \leftrightarrow \Box\Diamond table)$ , where *table*, *kitchen*, *obstacle*, and *request* are predicates over the state space. The robot itself is modeled as the discrete-time abstraction of the standard 3-dimensional Dubins vehicle [15] with an additional (i.e., 4th) dimension that records if a *request*, which is controlled by the environment, is pending. In Table 3, we summarize the computational resources, and, in Fig. 2, we show a simulated closed-loop trajectory that was plotted using our tool *Simulate*. We observe that *Sylvan* was much faster, but *CUDD* consumed much less memory.

## References

1. Abate, A., et al.: ARCH-COMP21 category report: stochastic models. In: 8th International Workshop on Applied Verification of Continuous and Hybrid Systems, pp. 55–89 (2021)
2. Babiak, T., et al.: The Hanoi omega-automata format. In: Kroening, D., Păsăreanu, C.S. (eds.) CAV 2015. LNCS, vol. 9206, pp. 479–486. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-21690-4\\_31](https://doi.org/10.1007/978-3-319-21690-4_31)
3. Banerjee, T., Majumdar, R., Mallik, K., Schmuck, A.-K., Soudjani, S.: A direct symbolic algorithm for solving stochastic Rabin games. In: TACAS 2022. LNCS, vol. 13244, pp. 81–98. Springer, Cham (2022). [https://doi.org/10.1007/978-3-030-99527-0\\_5](https://doi.org/10.1007/978-3-030-99527-0_5)
4. Banerjee, T., Majumdar, R., Mallik, K., Schmuck, A.K., Soudjani, S.: Fast symbolic algorithms for omega-regular games under strong transition fairness. *TheoretiCS* (to appear) (2023). arXiv preprint [arXiv:2202.07480](https://arxiv.org/abs/2202.07480)
5. Chatterjee, K., de Alfaro, L., Henzinger, T.A.: The complexity of stochastic Rabin and Streett games. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 878–890. Springer, Heidelberg (2005). [https://doi.org/10.1007/11523468\\_71](https://doi.org/10.1007/11523468_71)
6. Chatterjee, K., De Alfaro, L., Faella, M., Majumdar, R., Raman, V.: Code aware resource management. *Formal Methods Syst. Des.* **42**(2), 146–174 (2013)
7. Coogan, S., Arcak, M.: Efficient finite abstraction of mixed monotone systems. In: Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control, pp. 58–67 (2015)
8. Dijk, T.: Attracting tangles to solve parity games. In: Chockler, H., Weissenbacher, G. (eds.) CAV 2018. LNCS, vol. 10982, pp. 198–215. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-96142-2\\_14](https://doi.org/10.1007/978-3-319-96142-2_14)
9. Dijk, T.: Oink: an implementation and evaluation of modern parity game solvers. In: Beyer, D., Huisman, M. (eds.) TACAS 2018. LNCS, vol. 10805, pp. 291–308. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-89960-2\\_16](https://doi.org/10.1007/978-3-319-89960-2_16)
10. Dutreix, M., Huh, J., Coogan, S.: Abstraction-based synthesis for stochastic systems with omega-regular objectives. *Nonlinear Anal. Hybrid Syst.* **45**, 101204 (2022)
11. Geretti, L., et al.: ARCH-COMP20 category report: continuous and hybrid systems with nonlinear dynamics. In: Proceedings of the 7th International Workshop on Applied Verification of Continuous and Hybrid Systems, pp. 49–75 (2020)

12. Křetínský, J., Meggendorfer, T., Sickert, S.: Owl: a library for  $\omega$ -words, automata, and LTL. In: Lahiri, S.K., Wang, C. (eds.) ATVA 2018. LNCS, vol. 11138, pp. 543–550. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-01090-4\\_34](https://doi.org/10.1007/978-3-030-01090-4_34)
13. Long, D.E., Browne, A., Clarke, E.M., Jha, S., Marrero, W.R.: An improved algorithm for the evaluation of fixpoint expressions. In: Dill, D.L. (ed.) CAV 1994. LNCS, vol. 818, pp. 338–350. Springer, Heidelberg (1994). [https://doi.org/10.1007/3-540-58179-0\\_66](https://doi.org/10.1007/3-540-58179-0_66)
14. Majumdar, R., Mallik, K., Rychlicki, M., Schmuck, A.K., Soudjani, S.: A flexible toolchain for symbolic Rabin games under fair and stochastic uncertainties (2023). <https://kmallik.github.io/assets/pdf/cav23-toolpaper.pdf>
15. Majumdar, R., Mallik, K., Schmuck, A.K., Soudjani, S.: Symbolic qualitative control for stochastic systems via finite parity games. IFAC-PapersOnLine **54**(5), 127–132 (2021)
16. Majumdar, R., Mallik, K., Soudjani, S.: Symbolic controller synthesis for büchi specifications on stochastic systems. In: Proceedings of the 23rd International Conference on Hybrid Systems: Computation and Control, pp. 1–11 (2020)
17. Piterman, N., Pnueli, A.: Faster solutions of Rabin and Streett games. In: 21st Annual IEEE Symposium on Logic in Computer Science (LICS’06), pp. 275–284. IEEE (2006)
18. Rungger, M., Zamani, M.: Scots: a tool for the synthesis of symbolic controllers. In: Proceedings of the 19th International Conference on Hybrid Systems: Computation and Control, pp. 99–104 (2016)
19. Schewe, S.: An optimal strategy improvement algorithm for solving parity and payoff games. In: Kaminski, M., Martini, S. (eds.) CSL 2008. LNCS, vol. 5213, pp. 369–384. Springer, Heidelberg (2008). [https://doi.org/10.1007/978-3-540-87531-4\\_27](https://doi.org/10.1007/978-3-540-87531-4_27)
20. Somenzi, F.: Cudd: CU decision diagram package release 3.0.0 (2015). <https://github.com/ivmai/cudd>
21. van Dijk, T., van de Pol, J.: Sylvan: multi-core decision diagrams. In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 677–691. Springer, Heidelberg (2015). [https://doi.org/10.1007/978-3-662-46681-0\\_60](https://doi.org/10.1007/978-3-662-46681-0_60)
22. Zielonka, W.: Infinite games on finitely coloured graphs with applications to automata on infinite trees. Theor. Comput. Sci. **200**(1), 135–183 (1998). [https://doi.org/10.1016/S0304-3975\(98\)00009-7](https://doi.org/10.1016/S0304-3975(98)00009-7), <https://www.sciencedirect.com/science/article/pii/S0304397598000097>

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

