arXiv:2001.04289v1 [cs.LO] 8 Jan 2020

Symblicit Exploration and Elimination for Probabilistic Model Checking*

Ernst Moritz Hahn[1] and Arnd Hartmanns[2]

[1] Queen's University Belfast, Belfast, UK
[2] University of Twente, Enschede, The Netherlands

Abstract. Binary decision diagrams can compactly represent vast sets of states, mitigating the state space explosion problem in model checking. Probabilistic systems, however, require multi-terminal diagrams storing rational numbers. They are inefficient for models with many distinct probabilities and for iterative numeric algorithms like value iteration. In this paper, we present a new "symblicit" approach to checking Markov chains and related probabilistic models: We first generate a decision diagram that *symb*olically collects all reachable states and their predecessors. We then concretise states one-by-one into an exp*licit* partial state space representation. Whenever all predecessors of a state have been concretised, we *eliminate* it from the explicit state space in a way that preserves all relevant probabilities and rewards. We thus keep few explicit states in memory at any time. Experiments show that very large models can be model-checked in this way with very low memory consumption.

1 Introduction

Many of the complex systems that we are surrounded by, rely on, and use every day are inherently probabilistic: The Internet is built on randomised algorithms such as the collision avoidance schemes in Ethernet and wireless protocols, with the latter additionally being subject to random message loss. Hard- and software in cars, trains, and airplanes is designed to be fault-tolerant based on mean-time-to-failure statistics and stochastic wear models. Machine learning algorithms give recommendations based on estimates of the likelihoods of possible outcomes, which in turn may be learned from randomly sampled data.

Given a formal mathematical model of such a system, e.g. in the form of (a high-level description of) a discrete- or continuous-time Markov chain (DTMC or CTMC), probabilistic model checking can automatically compute (an approximation of) the value of a quantity of interest. Such quantities include the probability to finally reach an unsafe state (a measure of reliability), the steady-state probability to be in a failure state (determining availability), the long-run average reward (measuring e.g. throughput or energy consumption), or the accumulated cost up to a certain set of states (where e.g. a job is complete). The

* The authors are listed alphabetically. This work was supported by NWO VENI grant no. 639.021.754.

standard approach is to proceed in two phases: First, *explore* the state space, building a representation of the set of reachable states and the transitions connecting them. Transitions are annotated with rational values for probabilities and rewards, which are usually represented as floating-point numbers. Second, use an iterative numeric algorithm such as value iteration [35] or one of its sound variants [5,21,29,36] to *compute* the quantity of interest. These algorithms in fact compute a value for every state that approximates the quantity starting from that state up to a prescribed error ϵ. In contrast to classic functional model checking, which admits on-the-fly algorithms for e.g. reachability or LTL properties, probabilistic model checking is thus doubly affected by the state space explosion problem: First, the entire state space must be stored in memory, including many numeric values. Second, the numeric computation requires multiple vectors of values to be stored, and updates to be performed on them, for all states.

Current approaches to mitigate the state space explosion problem in probabilistic model checking include the use of partial exploration and learning algorithms, bisimulation minimisation, and compact representations of the state space or value vectors by binary decision diagrams (BDDs). They exploit different structural properties that only sometimes overlap. The learning-based approaches [1, 8] for reachability probabilities work well for models where a small initial subset of the state space determines most of the probability mass. In such cases, which are not abundant among existing case studies [23], they complete in a few seconds while exhaustive approaches run out of time or memory [11, Table 1]. Bisimulation minimisation reduces the state space to a quotient according to a probabilistic bisimulation relation; see [3, Sect. 5.1] for an overview. It has been implemented in Storm [15] and allows certain very large models to be checked efficiently; in general, its impact depends on the amount of bisimilar states in the given system. Finally, BDDs [9, 33] have a long history of use in (discrete-state) model checking [12] to compactly represent state spaces, in good cases reducing memory usage by orders of magnitude. They work well when the state space is structured and exhibits symmetry, which is often the case for real-life case studies modelled by humans (as opposed to randomly generated examples). In probabilistic model checking, however, numeric values from continuous domains are part of state spaces and must be encoded in the decision diagrams. A binary encoding of their floating-point representation does not usually result in compact diagrams; instead, multi-terminal BDDs (MTBDDs), where each of the (unbounded number of) leaves represents one number, have been applied with some success, notably in the probabilistic model checker Prism [32]. They however do not provide much compaction for models with many distinct probabilities or reward values due to the large number of leaves. They also do not work well to represent the large vectors of values used in iterative numeric algorithms such as value iteration, which progress through many very different intermediate values for each state before converging to, but often not reaching, a fixpoint. For this reason, Prism defaults to its *hybrid* engine, which uses MTBDDs for the state space but arrays of double-precision values for iteration. Its fully symbolic *mtbdd* engine only solves specific large structured models in reasonable runtime.

Our contribution is a new approach that combines (MT)BDDs, explicit state representations, and state elimination to tackle the problem of model-checking large probabilistic specifications. Its novelty lies in (1) using BDDs precisely for those tasks that they work well for, and (2) using state elimination instead of the standard iterative algorithms in the computation phase. We work with discrete-state probabilistic systems; in this paper, we use DTMC to explain our approach, but the same techniques apply directly to CTMC, Markov decision processes (MDP) [35], and Markov automata (MA) [16], too.

Our state space exploration performs a standard explicit-state breadth-first search, but we use a decision diagram instead of the standard hashset to store the set of visited states. We do not store transitions, thus no continuous numeric values blow the diagram up. However, we do count the number of predecessors of each state—thus we use an MTBDD. Since this number is a discrete quantity with low variation in most models, the diagrams usually remain compact. For the computation phase, we explore the state space again, this time creating a representation that includes transitions, but that is explicit. During this exploration, we keep track of the number of explored predecessors of each state. Whenever, for some fully-explored state s, this number reaches the predecessor count given by the MTBDD, we apply *state elimination*: we remove s from the (yet incomplete) explicit state space representation, and replace all of its incoming and outgoing transitions by direct transitions between the predecessors and successors of s. By redistributing the original transitions' probabilities and rewards in the right way, the quantity of interest remains unaffected. This method of computation simultaneously avoids the iterative algorithms' convergence and precision issues [20] and keeps memory usage due to the explicit representation low: on most models, most states are eliminated soon after they have been fully explored, thus only few need to be kept in memory at any time. Upon termination, only the initial state and goal state(s) remain, and the value for the quantity of interest can be read off the transitions connecting them.

Two technical insights make our computation phase work well: First, we use an explicit representation not only to avoid storing (continuous) probabilities and rewards in an MTBDD, but also because state elimination tends to create unstructured intermediate state spaces that would blow up any BDD representation. Second, the precomputed predecessor count allows us to eliminate a state at precisely the moment after which we will not encounter it again in our search, avoiding costly re-explorations and re-eliminations.

Related work. State elimination stems from the classic reduction algorithm to convert a finite automaton into a regular expression [10]. It was introduced to probabilistic model checking to solve parametric Markov chains [13] and forms the core of the PARAM [25] and PROPHESY [14] tools. For non-parametric models, it enables efficient computation of reward-bounded reachability probabilities [22]. In this paper, we use it for non-parametric Markov chains and unbounded (infinite-horizon) properties. In all of these settings, its effectiveness crucially depends on the order in which states are eliminated, which is determined by (configurable) heuristics. Symblicit techniques have previously been

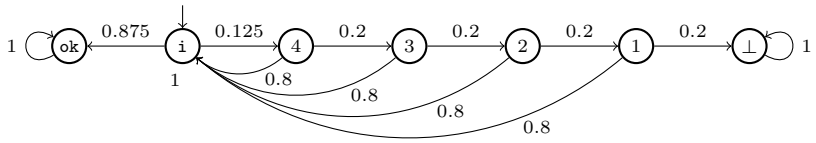

Fig. 1. DTMC M_z for the Zeroconf protocol ($h = 32, a = 2^8, p = 0.2, n = 4$)

used for long-run average properties [38], based on bisimulation minimisation, and later expanded in related settings [7]. A different form of elimination on strongly-connected components was used by Gui et al. [19] to accelerate the (explicit-state) computation of reachability probabilities via value iteration.

2 Background

Mathematical notions. $\mathbb{R}_0^+$ is the set of all non-negative, and $\mathbb{R}^+$ the set of all positive, real numbers. A (discrete) *probability distribution* over S is a function $\mu \in S \to [0,1]$ with countable *support* $spt(\mu) \stackrel{\text{def}}{=} \{\, s \in S \mid \mu(s) > 0 \,\}$ and $\sum_{s \in spt(\mu)} \mu(s) = 1$. $Dist(S)$ is the set of all probability distributions over S.

2.1 Discrete-Time Markov Chains

Definition 1. *A* discrete-time Markov chain *(DTMC) is a tuple* $M = \langle S, s_I, P, R \rangle$ *consisting of a finite set of* states S*, an* initial state $s_I \in S$*, a* transition *function* $P \colon S \to Dist(S)$*, and a* reward *function* $R \colon S \to \mathbb{R}_0^+$.

We also write a transition as $s \xrightarrow{p} s'$ if $p = P(s)(s') > 0$. A transition is uniquely identified by the two states it connects. When in state s of a DTMC, we delay for one time unit before jumping to the next state. *Continuous-time Markov chains* (CTMC) extend DTMC by additionally assigning a *rate* $Q(s) \in \mathbb{R}^+$ to every state. Then the probability to delay for at most t time units is $1 - \mathrm{e}^{-Q(s) \cdot t}$, i.e. the residence time follows the exponential distribution with rate $Q(s)$. In both models, the probability to then move to state s' is given by $P(s)$. When staying for t time units in state s, we incur a reward of $R(s) \cdot t$. To simplify the presentation, we use DTMC throughout this paper, but mention the changes needed in definitions or algorithms to use CTMC, where appropriate.

Example 1. As a running example, we use a very abstract model of the Zeroconf protocol [6], shown as DTMC M_z in Fig. 1 (adapted from [18]). We draw transitions as arrows labelled with their probability. Non-zero rewards are given next to the states. M^z has 7 states and 12 transitions. The protocol is used by hosts joining a network to auto-configure a unique IP address. A new host joining the network of $h = 32$ host starts in state `i`. It selects an address uniformly at random from the space of $a = 2^8$ addresses. The probability that the address is already in use is $\frac{h}{a} = \frac{1}{8}$. The host checks $n = 4$ times whether its address is

already in use. If it is not, all checks will succeed, modelled by state `ok`, to which we moved with probability $1 - \frac{h}{a}$. If it is, then states n down to 1 model the checks. Each check can fail to give the correct negative result due to message loss with probability $p = 0.2$. If all tests do so, then the host incorrectly believes that it has a unique address, in state $\bot$. Otherwise, it retries with a newly chosen address from state `i`. We incur a reward of 1 in state `i`, i.e. for every IP address we try. The size of the DTMC can be blown up arbitrarily via parameter n.

In practice, higher-level modelling languages like MODEST [24] or the PRISM language [32] are used to specify larger DTMC. The semantics of a DTMC is formally captured by its *paths*:

Definition 2. *Given a DTMC M as above, a* finite path *is a sequence $\pi_{\text{fin}} = s_0\, t_0\, s_1\, t_1 \ldots s_n$ of states $s_i \in S$ and delays $t_i \in \mathbb{R}^+$ where $P(s_i)(s_{i+1}) > 0$ and $t_i = 1$ for all $i \in \{0, \ldots, n-1\}$. Let $|\pi_{\text{fin}}| \stackrel{\text{def}}{=} n$, $\text{last}(\pi_{\text{fin}}) \stackrel{\text{def}}{=} s_n$, $\text{dur}(\pi_{\text{fin}}) \stackrel{\text{def}}{=} \sum_{i=0}^{n-1} t_i$, and $\text{rew}(\pi_{\text{fin}}) \stackrel{\text{def}}{=} \sum_{i=0}^{n-1} t_i \cdot R(s_i)$. Π_{fin} is the set of all finite paths starting in s_I. A* path *is an analogous infinite sequence π, and Π are all paths starting in s_I. We define $s \in \pi \Leftrightarrow \exists\, i\colon s = s_i$. Let $\pi_{\to m}$ for $m \in \mathbb{N}$ be the prefix of π of length m, i.e. $|\pi_{\to m}| = m$, and let $\pi_{\to G}$ be the shortest prefix of π that contains a state in $G \subseteq S$, or $\bot$ if π contains no such state. Let $\text{rew}(\bot) \stackrel{\text{def}}{=} \infty$.*

In CTMC, the t_i can be arbitrary numbers in $\mathbb{R}_0^+$. For M as above, following the rules described below Definition 1 and the standard cylinder set construction [4], we obtain a probability measure $\mathbb{P}_M$ on measurable sets of paths starting in s_I.

Definition 3. *Given a set of goal states $G \subseteq S$, the* reachability probability *w.r.t. g is $\mathbb{P}(\diamond\, G) \stackrel{\text{def}}{=} \mathbb{P}_M(\pi \in \Pi \mid \exists\, g \in G\colon g \in \pi)$. Let $r_G\colon \Pi \to \mathbb{R}_0^+$ be the random variable defined by $r_G(\pi) = \text{rew}(\pi_{\to G})$. Then the* expected reward *to reach G is the expected value of r_G under $\mathbb{P}_M$, written as $\mathbb{E}(\blacklozenge\, G)$. Let $r_{lra}\colon \Pi \to \mathbb{R}_0^+$ be defined by $r_{lra}(\pi) = \lim_{i\to\infty} \text{rew}(\pi_{\to i})/\text{dur}(\pi_{\to i})$. Then the* long-run average reward *is the expected value of r_{lra} under $\mathbb{P}_M$, written as $\mathbb{L}$.*

The steady-state probability $\mathbb{S}(S')$ of residing in a state in $S' \subseteq S$ is a special case of the long-run average reward where $R(s) = 1$ if $s \in S'$ and 0 otherwise. Whenever we consider a DTMC with a set of goal states G, we assume that they have been made absorbing, i.e. that for all $g \in G$ we have $P(g)(g) = 1$. Given a CTMC, reachability probabilities and expected rewards can be computed on its *embedded DTMC*, obtained by dividing all rewards by $Q(s)$; only for long-run averages do we need a dedicated treatment of the rates resp. residence times.

Example 2. For our Zeroconf example DTMC M_z from Fig. 1, we may want to compute the probability to eventually pick a unique address $\mathbb{P}(\diamond\, \{\, \texttt{ok}\,\})$, which will be just below 1, and the expected number of addresses that we ever try $\mathbb{E}(\blacklozenge\, \{\, \texttt{ok}, \bot\,\})$. Note that $\mathbb{E}(\blacklozenge\, \{\, \texttt{ok}\,\})$ is ∞ by definition since the set of paths that never reach state `ok` has positive probability.

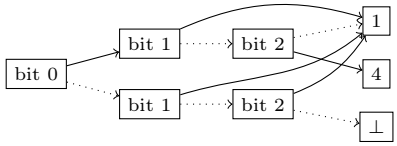

Fig. 2. MTBDD counting the numbers of predecessors for all states of M_z

2.2 Binary Decision Diagrams

Binary decision diagrams (BDDs) [9, 33] represent Boolean functions as rooted directed acyclic graphs. They have two leaf nodes, *true* and *false*. Every inner node is associated to one input bit, and has two children: the high (solid line) and low (dotted line) child. On a path from the root to a leaf, every bit must occur at most one. Such a path corresponds to the inputs in which bit b_i is assigned to *true* (*false*) if we go from a node for b_i to its high (low) child. We typically order the bits on all paths, merge isomorphic subgraphs, and remove redundant nodes. Then a BDD can represent many functions with few nodes.

In model checking, BDDs are used to represent sets of states (by assigning *true* to the binary encoding of a state iff it is in the set) as well as the transition relation (by assigning *true* to the binary encoding of a pair of states if they are connected by a transition). In probabilistic model checking, however, we need to encode functions that map to rational numbers to encode transition probabilities, rewards, and the value vectors in value iteration. Most tools represent them as 64-bit floating-point values, but the corresponding binary representation does not typically allow good compression with BDDs. Symbolic probabilistic model checkers such as PRISM [32] this use *multi-terminal* BDDs (MTBDDs) with one leaf node per number. Since a finite model only contains finitely many probabilities, or values for states, this approach is effective, but often not efficient: for example, when performing value iteration on our example DTMC M_z for $\mathbb{P}(\diamond \{ \texttt{ok} \})$, we have to encode the following function after 5 iterations:

$$\{ \texttt{i} \mapsto 0.99225, 4 \mapsto 0.9716, 3 \mapsto 0.966, 2 \mapsto 0.938, 1 \mapsto 0.784, \texttt{ok} \mapsto 1, \bot \mapsto 0 \}$$

Observe that every state has a distinct value, thus the MTBDD offers no compression. In practice, they only work well for very specific models with few distinct transition probabilities and rewards, and where the iterative numeric algorithms assign the same (intermediate) values to many states.

Example 3. Fig. 2 shows an MTBDD mapping every state of M_z to its number of predecessor states. We have 7 states, thus use 3 bits for their encoding. States 1 through 4 are encoded as that number, `i` is 5 (101_2), `ok` is 6 ((110_2), and $\bot$ is 7 (111_2). There is no (reachable) state encoded as 0, thus we map 0 to the extra $\bot$ leaf node—in this way, such an MTBDD can indicate that certain states are unreachable, or have not been explored yet. Observe that the MTBDD representation achieves some compression by excluding two redundant nodes for bit 2. If we scale the model up by increasing n, the compression would increase.

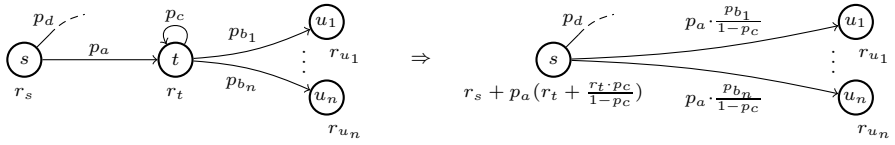

Fig. 3. DTMC state elimination

2.3 State Elimination

State elimination is a process by which a state of a DTMC is removed, i.e. transitions are modified such that it is no longer reachable from the initial state and it is removed from the state set S, in a way that preserves the values of all properties of interest. We show the schematic of state elimination in Fig. 3: we eliminate a state t by redistributing the probability to enter a self-loop onto its other outgoing transitions, then combine its incoming and outgoing transitions. It is easy to see that this preserves the probabilities of all measurable sets of paths that pass through t when projecting t out from every path. In particular, the paths that forever take the self-loop have probability mass zero, which is why we could eliminate the loop. For rewards, the transformation only preserves the *expected* reward values of sets of paths:

1. In t, the expected number of times we take the self-loop is $\frac{p_c}{1-p_c}$, thus the expected reward from passing through t is $\frac{r_t \cdot p_c}{1-p_c}$ (for the loop) plus r_t (for taking one of the other outgoing transitions.
2. Out of s, the probability to enter t next is p_a, thus we multiply the expected reward of passing through t by p_a and add this to the reward of s.

Alg. 1 shows the pseudocode to perform state elimination on a DTMC stored in explicit data structures (i.e. hash sets for states, lists of transitions, etc.).

3 Symblicit Exploration and Elimination

As we explained in Sect. 1 and illustrated in Sect. 2.2, many probabilistic models do not give rise to a compact BDD-based representation if the numeric values—probabilities, rewards, rates for CTMC—are included. Furthermore, the standard iterative numeric algorithms like value iteration usually produce data that is hardly BDD-compressible. In this section, we present a combined symbolic-explicit approach that uses MTBDDs in a way that usually avoids such problems, and that uses state elimination to calculate $\mathbb{P}$, $\mathbb{E}$, and $\mathbb{L}$ values without having to keep (values for all states of) the entire state space in memory.

The pseudocode of our approach is shown as function `ExploreEliminate` in Alg. 2. It uses functions `Explore` of Alg. 3 and `Eliminate` of Alg. 1. We typeset values that represent executable code in `monospace` font: compact specifications in high-level modelling languages are typically compiled to or interpreted as functions that, given an explicit (bit string) representation of a state, enumerate its transitions (`P`), compute its reward (`R`), and return *true* iff it is a goal state

```
1  function Eliminate(⟨S, s_I, P, R⟩, s ∈ S, S_keep ⊆ S)                    // all explicit
2      if s ∈ spt(P(s)) ∧ P(s) < 1 then // if s has a self-loop and other transitions:
3          foreach s' ∈ spt(P(s)) \ { s } do    // redistribute onto other transitions
4              P(s)(s') := P(s)(s')/(1 − P(s)(s))          // the self-loop's probability
5          R(s) := R(s) + R(s) · P(s)(s)/(1 − P(s)(s))   // add the expected reward
6          P(s)(s) := 0                                                  // remove the self-loop
7      foreach s_pre ∈ { s'' | s ∈ spt(P(s'')) } \ { s } do // for every predecessor s_pre:
8          p := P(s_pre)(s), P(s_pre)(s) := 0   // remove the transition from s_pre to s
9          foreach s' ∈ spt(P(s)) do                  // then merge the transitions of s
10             P(s_pre)(s') := P(s_pre)(s') + p · P(s)(s') // into the transitions of s_pre
11         R(s_pre) := R(s_pre) + p · R(s)    // merge the reward of s into that of s_pre
12     if s ∉ S_keep ∧ P(s)(s) = 0 then                  // if s is not needed: remove
13         P := P \ { s ↦ P(s) }, R := R \ { s ↦ R(s) }   // its transitions, reward,
14         S := S \ { s }                                                // and the state itself
```

Alg. 1: State elimination for probabilities and expected rewards

(G). We mark variables storing symbolic data (i.e. BDDs or MTBDDs) with a $\hat{hat}$. All other values typeset in *italics* use explicit data structures such as bit strings for states, hash sets or queues of such bit strings, lists of transitions, etc.

Our first step, in line 2, is to symbolically explore the set of reachable states by calling function `Explore`. This function performs a standard breadth-first search, using a BDD for the set of visited states, and additionally constructs an MTBDD that counts the number of predecessors of each state like the one shown in Fig. 2 for M_z. In our implementation, *seen* and *pre* are actually managed in a single MTBDD as explained in Example 3.

We then, starting from line 3, perform another exploration of the state space. This time, however, we use explicit data structures, and we track the number of *fully explored* predecessors for every state in hash table pre'. A state is fully explored if its reward, all of its transitions, and all successor states, have been added to the explicit representations for S, R, and P. We track the set of fully explored states in hash set *done*. Whenever we are done visiting a state s in this second exploration (i.e. in line 13 and below), it has just become fully explored, and the fully-explored-predecessor count of its successors has changed. We then check which of these changed states fulfils the criteria for being eliminated: It must be fully explored (which only s is for certain), and all of its predecessors must be fully explored (which we determine by comparing pre' and $\hat{pre}$). We call `Eliminate` on these states in line 17. In this way, if we indeed manage to eliminate most states soon after they have been explored, the explicit data structures—S, P, R, pre', *done*, etc.—only track few states at any time and thus consume little memory. The predecessor count in $\hat{pre}$ is crucial for being able to perform efficient elimination; without it, we would have to apply heuristics that could lead to states being eliminated that would later be explored as successors of other states again, leading to costly re-exploration and re-eliminations.

```
1  function ExploreEliminate(s_I, P, R, G)        // explicit s_I, executable P, R, G
2      p̂re := Explore(s_I, P)                      // get predecessor count MTBDD for all states
3      done := ∅, agenda := { s_I }                // done stored as hash set, agenda as queue
4      S := { s_I }, P := ∅, R := ∅, pre' := { s_I ↦ 0 }   // explicit sets and functions
5      while agenda ≠ ∅ do
6          s := next element of agenda, agenda := agenda \ { s }
7          foreach s' ∈ spt(P(s)) do                                   // explore s
8              P(s)(s') := P(s)(s'), R(s) := R(s)
9              if s' ∉ S then
10                 S := S ∪ { s }, agenda := agenda ∪ { s }
11                 pre' := pre' ∪ { s' ↦ 0 }
12             if s' ≠ s then pre'(s') := pre'(s') + 1
13         done := done ∪ { s }                                        // s is now fully explored
14         E := { s_e | s_e ∈ {s} ∪ spt(P(s)) ∩ done }                 // collect just modified states
15         E := { s_e | s_e ∈ E ∧ p̂re(s_e) = pre'(s) }                 // with all predecessors explored
16         foreach s_elim ∈ E do                                       // and eliminate them
17             Eliminate(⟨S, s_I, P, R⟩, s_elim, { s_I })
18             if s_elim ∉ S then                                      // cleanup
19                 pre' := pre' \ { s_elim ↦ pre'(s_elim) }, done := done \ { s_elim }
20     if we compute a reachability probability then return Σ_{g∈G} P(s_I)(g)
21     else if we compute an expected reward then
22         if spt(P(s_I)) \ G ≠ ∅ then return ∞                        // we have ℙ(◇ G) < 1
23         else return R(s_i) + Σ_{g∈G} P(s_I)(g) · R(g)
```

Alg. 2: Symblicit exploration-elimination for probabilities and expected rewards

In `Eliminate`, if a state is part of the set S_{keep}, we still modify and "redirect" the transitions of its predecessors to go around this state, but we do not remove it from the state space. We use this to avoid eliminating the initial state s_I. We also do not eliminate states whose only transition is a self-loop: they do not have successors to which transitions could be redirected. Elimination will thus eventually reduce each bottom strongly connected component (BSCC) of the DTMC to one such self-loop state. Since we assume all goal states to only have a self-loop, each of them is a BSCC. Once the outer loop of line 5 in `ExploreEliminate` finishes, `Eliminate` has been called for all states. Every surviving state at this point is thus the result of eliminating a number of transient states plus a non-goal BSCC or a goal state, and has become a direct successor of the initial state. We can then directly read the value of $\mathbb{P}(\diamond\, \mathtt{G})$ from the transitions to the goal states (line 20). Similarly, the value of $\mathbb{E}(\blacklozenge\, \mathtt{G})$ can be derived directly from the remaining rewards, if it is not ∞ by definition (lines 22-23).

Example 4. For our example DTMC M_z of Fig. 1, we have already shown the predecessor count MTBDD computed by `Explore` in Fig. 2. Let us now step through the rest of `ExploreEliminate` on this model. The partial state spaces that we consider in each step are shown in Fig. 4. Fully explored states are drawn

```
1  function Explore(s_I, P)                          // explicit s_I, executable P
2      seen := { s_I }, agenda := { s_I }      // seen stored as BDD, agenda as queue
3      pre := { s_I ↦ 0 }                      // predecessor count, stored as MTBDD
4      while agenda ≠ ∅ do
5          s := next element of agenda, agenda := agenda \ { s }
6          foreach s' ∈ spt(P(s)) \ { s } do
7              if s' ∉ seen then
8                  seen := seen ∪ { s }, agenda := agenda ∪ { s }
9                  pre := pre ∪ { s' ↦ 0 }
10             pre(s') := pre(s') + 1     // s is a previously-unseen predecessor of s'
11     return pre
```

Alg. 3: Symbolic exploration via breadth-first search with predecessor counting

with solid outlines, all other states (i.e. those in S but not in *done*) with dashed outlines. In step (1), we have just fully explored state `i`, i.e. we just executed line 13 in the first iteration of the outer loop. Since *pre* tells us that `i` still has unexplored predecessors, we cannot eliminate, and next explore `ok` in step (2). We then eliminate `ok`—its only predecessor `i` is fully explored—but since `ok` has just a single self-loop, the elimination has no effect. In step (3), we have just explored state 4, which can now be eliminated. The result is shown as step (4). We proceed in the same pattern in steps (5) through (8). Then, in step (9), we fully explore state 1. Now all predecessors of `i` are fully explored, and we can eliminate both 1 and `i`. For the sake of illustration, let us pick the more complicated ordering and eliminate `i` first. The result is shown as step (10). Since `i` is the initial state, we keep it, but redirect all incoming transitions. We also merge its rewards, which is why 1 now has a non-zero reward. Note that we show rationals in Fig. 4, but our implementation uses floating-point numbers. Remember that, without *pre*, we might have eliminated `i` too early; after any subsequent exploration of a state in $\{1, \ldots, n\}$, we would then have to re-eliminate `i`. We finally eliminate 1 in step (11) and explore state $\bot$ in step (12). At this point, the outer loop terminates; we read $\mathbb{P}(\diamond \{\texttt{ok}\}) = \frac{4375}{4376} \approx 0.999771$ and $\mathbb{E}(\blacklozenge \{\texttt{ok}, \bot\}) = 1 + \frac{119918}{119793} \approx 2.001043$. Observe that, at any time, we kept at most 4 explicit states in memory. We can arbitrarily increase the size of this model by increasing n, but will only ever need at most 4 states in memory.

Long-run average rewards. The algorithm we presented so far computed reachability probabilities and expected rewards. For long-run average reward properties, there are no goal states. In such a case, our state elimination procedure computes the *recurrence reward* for each BSCC [17].

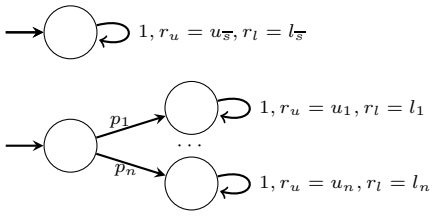

Fig. 5. Computation of long-run averages.

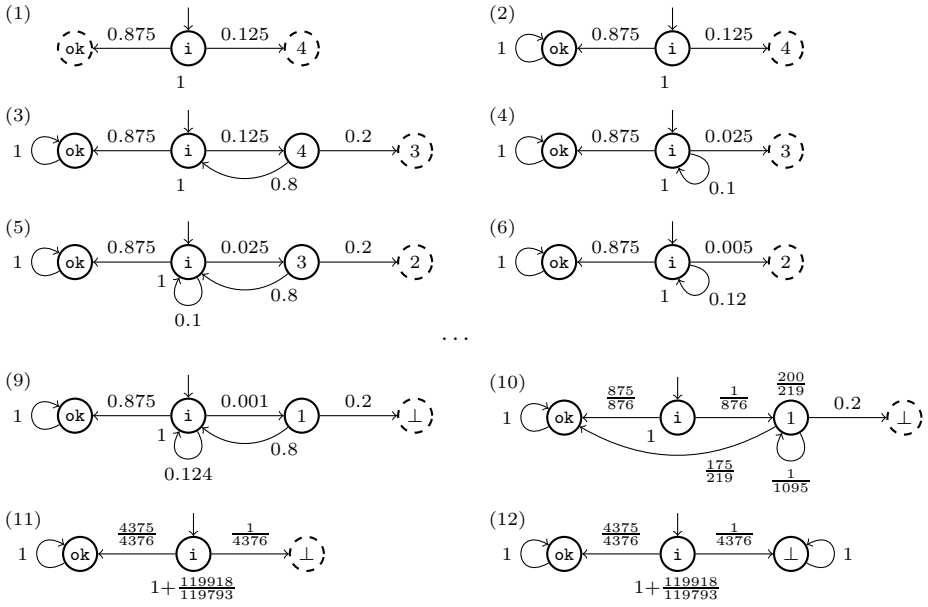

Fig. 4. Example for exploration with interleaved elimination on M_z

To obtain the long-run average reward, we need to divide the recurrence reward for the rewards as given in the DTMC by the recurrence reward that we would obtain if all states had reward 1. We can do so by straightforwardly extending Algs. 1 and 2 to work on two reward structures r_u and r_l in parallel. Upon termination of the outer loop in `ExploreEliminate`, we then have one of the two situations described at the end of Sect. 4 in [17], and can again directly read off the value for our ($\mathbb{L}$-)property. Consider Fig. 5: In the simpler case, the remaining model consists of the initial state $\overline{s}$ with a self-loop with probability one and $r_u = u_{\overline{s}}$, $r_l = l_{\overline{s}}$. In this case, the average value is $\frac{u_{\overline{s}}}{l_{\overline{s}}}$. In the other case, the remaining model consists of the initial state $\overline{s}$ which has a probability of p_i to move to one of the other n remaining states s_i , $i = 1, \ldots, n$, which all have a self-loop with probability one and $r_u(s_i) = u_i$, $r_l(s_i) = l_i$. In this case, the average value is $\sum_{i=1,\ldots,n} p_i \frac{u_i}{l_i}$.

When computing long-run average rewards, the final value for $\mathbb{L}$ may be small, but the two recurrence rewards that we need to divide are often extremely large numbers beyond what can usefully be represented as 64-bit (i.e. double-precision) floating point numbers. We thus implemented a variant of our algorithm that uses the GNU MPFR library (see mpfr.org) for arbitrary-precision floating-point arithmetic, allowing us to use more than 64 bits. We did not find this to significantly affect the performance of the overall approach.

Alternatives and optimisations. So far, we have assumed that we compute successors for each state explicitly and individually. For the state elimination phase,

doing so is indeed necessary. However, for just exploring the states we could also compute the *transition relation* as a BDD, and then use the transition relation to symbolically explore the set of reachable states. This is the standard approach in model checkers such as PRISM [32] and potentially faster than the semi-symbolic approach we have discussed. Also, using the transition relation and according MTBDD operations (in particular sum-abstraction), the number of predecessors of each state can also be computed symbolically.

We have so far assumed that the reachable states and the number of predecessors are stored as (MT)BDDs. An alternative to this approach is to store these numbers on secondary storage (e.g. hard disk) in a similar way as e.g. in [28]. This approach would be useful for models the state space of which is not suitable to be stored as a BDD. This might be the case because of lack of implicit symmetries or because the size of the representation of each state is not constant.

4 Experimental Evaluation

We have implemented a preliminary version of our method which we integrated as a plugin for the probabilistic model checker ePMC [26]. For the analysis, we transform the model and property into C++ code so as to have a means to quickly compute successors of states, similar to the approach used in SPIN [31]. This C++ file is then appended with code so as to achieve the following: In the first phase, we then explore the state space in a breadth-first manner where we explore each state explicitly but store sets of states as BDDs, using the BDD package CUDD [2]. In the second phase, we generate an MTBDD mapping all states to value 0. Then, we iterate over all reachable states, recompute their successors, and increment the value of these successors in the MTBDD by 1 each time. In the third phase, we execute the state elimination algorithm as discussed. This C++ code is then compiled and run in a process separate from ePMC, such that we can measure the memory usage exactly (the memory usage of ePMC itself is not of much interest, because it is rather small and about the same for any analysis).

In the following, we apply our tool on several case studies from the website of the probabilistic model checker PRISM. All experiments were performed on a MacBook Pro with a 2.7 GHz Quad-Core Intel Core i7 processor and 16 GB 2133 MHz LPDDR3 RAM. In the following tables, "model states" is the total number of states the model has for the given parameters, "result", is the value of the property computed, "time" is the total time of the analysis in seconds, "exp. states" ("exp. transitions") are maximal number of states (transitions) being stored explicitly at the same time. By "peak mem" we denote the maximal memory usage in MB of the analysis process.

4.1 Simple Molecular Reactions: $Na + Cl \leftrightarrow Na^+ + Cl^-$

This case study [3] is a CTMC modelling the chemical reaction $Na + Cl \leftrightarrow Na^+ + Cl^-$. The parameters of this case study are *N1*, the initial number of *Na* molecules and *N2*, the initial number of *Cl* molecules. In Table 1, we consider the performance figures for the analysis of `R=?[S]` which describes the expected long-run average number of *Na* molecules. Here, we consider a starting configuration in which initially the number of *Na* and *Cl* is the same, that is, $N1 = N2$.

N1=*N2*	model states	result	time	exp. states	exp. trans	peak mem
10	11	2.2623e+01	8	5	5	22
100	101	2.3894e+01	8	5	5	22
1,000	1,000	2.4012e+01	7	5	5	22
10,000	10,001	2.4024e+01	7	5	5	25
100,000	100,001	2.4025e+01	10	5	5	28
1,000,000	1,000,001	2.4025e+01	32	5	5	27
10,000,000	10,000,001	2.4025e+01	258	5	5	31
100,000,000	100,000,001	2.4025e+01	2,609	5	5	29
1,000,000,000	1,000,000,001	2.4025e+01	18,807	5	5	25

Table 1. $Na + Cl \leftrightarrow Na^+ + Cl^-$ performance figures.

As we see, the model scales well for large numbers of molecules and accordingly large state spaces. The memory usage grows only slowly with increasing model parameters, and the number of states and transitions required to be stored explicitly is constant.

4.2 Bounded Retransmission Protocol

The Bounded Retransmission Protocol [30][4] is a file transmission protocol. Files are divided into N packages, each of which is transferred individually. Data and confirmation packages are sent over unreliable channels, such that they might get lost. Packages can only be resent a number of MAX times. In Table 2, we provide performance figures for the analysis of the property `P=?[ F s=5 ]`, that is the probability that the sender does not eventually report a successful transmission. "N" and "MAX" are as discussed above, the other numbers are as in the previous case study.

Compared to the instances of the PRISM website, we have analysed instances with higher parameter numbers N and MAX because our focus was in the scalability of our method. For comparison, for the first table entry we used the same parameters as the last table entry on the PRISM website.

As we see, we are able to handle instances with several million states with a low memory usage. Even for higher parameter values for which the number of total states the model consists of is in the millions, we never use more than a few thousand explicit states and transitions and less than 100MB.

[3] https://www.prismmodelchecker.org/casestudies/molecules.php
[4] https://www.prismmodelchecker.org/casestudies/brp.php

N	MAX	model states	result	time	exp. states	exp. trans	peak mem
64	5	4,936	4.48e-08	9	12	29	25
64	10	9,101	1.05e-15	9	20	53	25
64	100	84,071	5.03e-153	9	140	517	25
64	1000	833,771	3.14e-1526	25	258	986	33
128	10	18,189	2.11e-15	9	20	53	25
128	100	168,039	1.01e-152	10	140	517	32
128	1000	1,666,539	6.28e-1526	38	514	1,978	36
256	10	36,365	4.21e-15	9	20	53	27
256	100	335,975	2.01e-152	15	150	517	28
256	1000	3,332,075	1.26e-1525	72	1,026	3,962	39
512	10	72,717	8.42e-15	9	20	53	29
512	100	671,847	4.03e-152	18	140	517	32
512	1000	6,663,147	2.51e-1525	140	1,340	5,167	42
1024	10	145,421	1.69e-14	10	20	53	29
1024	100	1,262,631	8.06e-152	29	140	517	30
1024	1000	13,325,291	5.02e-1525	280	1,340	5,167	43
2048	10	290,829	3.37e-14	12	20	53	27
2048	100	2,687,079	1.61e-151	50	140	517	30
2048	1000	26,649,579	1.00e-1524	552	1,340	5,167	41
4096	10	581,645	6.74e-14	17	20	53	27
4096	100	5,374,055	3.22e-151	95	140	517	28
4096	1000	53,298,155	2.01e-1524	1,151	1,340	5,004	39
8192	10	1,163,277	1.35e-13	26	20	53	28
8192	100	10,748,007	6.45e-151	187	140	517	27
8192	1000	106,595,307	4.02e-1524	2,385	1,340	5,004	40
16384	10	2,326,541	2.67e-13	48	20	53	27
16384	100	21,495,911	1.29e-150	392	140	517	28
16384	1000	213,189,611	8.03e-1524	4,534	1,340	5,004	40
32768	10	4,653,069	5.39e-13	91	20	53	26
32768	100	42,991,719	2.58e-150	781	140	517	27
32768	1000	426,378,219	1.61e-1523	9,039	1,340	5,004	41
65536	10	9,306,125	1.08e-12	20	54	156	25
65536	100	85,983,335	5.16e-150	140	503	1,362	29
131072	10	18,612,237	2.1572e-12	312	20	54	26
131072	100	171,966,567	1.03e-149	2,847	140	503	29

Table 2. Bounded Retransmission Protocol performance figures.

4.3 Wireless Communication Cell

This case study is a performance model of wireless communication cells [27][5]. The parameter N describes the number of channels in a cell. We analyse the property `R{"calls"}=? [ S ]` which describes the average number of calls in the cell on the long run. We provide performance figures in Table 3.

N	model states	result	time	exp. states	exp. trans	peak mem
10,000	10,001	7.00e+01	7	5	5	24
100,000	100,001	7.00e+01	9	5	5	24
1,000,000	1,000,001	7.00e+01	24	5	5	25
10,000,000	10,000,001	7.00e+01	183	5	5	26
100,000,000	100,000,001	7.00e+01	1,731	5	5	25

Table 3. Wireless Communication Cell performance figures.

Also for this case study, the approach works fine in that the number of states and transitions to be stored is small as is the peak memory usage.

4.4 Crowds Protocol

The Crowds protocol [37][6] is a means to allow anonymous web browsing. To do so, messages are not directly sent, but forwarded to other users, who might either forward them again or send them to the destination. By doing so, it is hard for attackers to decide whether the sender of a message is the original sender or is just forwarding the message. The model version we consider has two parameters: *TotalRuns* is the number of routing paths of the model instance, and *CrowdSize* is the number of honest participants of the protocol. We consider the property `Pmax=?[true U (new & runCount=0 & observe0 > observe1 & ... & observe0 > observe19)]` which means that an attacker is eventually able to observe the true sender of a message more often than participants just forwarding the message and is thus able to guess the original sender. In Table 4, we provide performance figures.

CrowdSize	*TotalRuns*	model states	result	time	exp. states	exp. trans	peak mem
5	5	8,653	2.7884e-01	9	289	1,278	77
5	6	18,817	2.9791e-01	9	510	2,236	77
5	7	37,291	3.1812e-01	9	823	3,898	164
10	5	111,294	2.1662e-01	19	6,604	33,513	2487
10	6	352,535	2.3162e-01	106	17,824	89,120	10698
15	5	592,060	1.9674e-01	676	44,104	356,143	9240

Table 4. Crowds Protocol performance figures.

As we see, for this model the current implementation does not perform very well. The reason is that too many states and transitions have to be stored explicitly at the same time, leading to a large memory overhead.

[5] https://www.prismmodelchecker.org/casestudies/cell.php
[6] https://www.prismmodelchecker.org/casestudies/crowds.php

4.5 Embedded Control System

This case study [34][7] is an embedded control system which features a cyclic polling process. If a certain component detects that more than a given number MAX_COUNT of cycles have been skipped due to issues,then the system is shut down for safety reasons. We consider the property `R{"danger"}=? [ F "down" ]`, which expresses the expected time the system is in an endangered state before it eventually has to be shut down. We provide performance figures in Table 5. Again, the analysis method works fine, as the number of states and transitions stored explicitly is limited.

MAX_COUNT	model states	result	time	exp. states	exp. trans	peak mem
512	320,316	3.3454e-01	56	267	11,938	80
1024	639,804	3.3731e-01	107	267	11,938	80
2048	1,278,780	3.4283e-01	201	267	11,938	80
4096	2,556,732	3.5371e-01	400	267	11,938	80
8192	5,112,636	3.7485e-01	794	267	11,938	80
16384	10,224,444	4.1469e-01	1,579	267	11,938	80
32768	20,448,060	7.6563e-01	2,914	284	10,513	72

Table 5. Embedded Control System performance figures.

5 Conclusion and Future Work

In this paper, we have discussed a new memory-efficient analysis method for properties of stochastic models. Our method is widely applicable to a large variety of stochastic models and properties, though, for conciseness of presentation we have concentrated on DTMCs (with some of the case studies being CTMCs). Experimental evidence has shown that our approach has the potential to analyse models with millions of states with just a few megabytes. One advantage of our method is that we directly obtain a precise result: Our current implementation is based on (variable-precision) floating-point numbers, and computation precision is limited by the properties of this representation. We could however as well use exact (rational) arithmetic and would obtain exact values without any change in the core algorithms, since all intermediate and final values are rational, though at the cost of increased computation time and memory usage. Alternatively, we could use interval arithmetic so as to obtain precise upper and lower bounds of the values computed, again, without any change in the algorithm, and with only moderate overhead. This is in contrast to methods based on value iteration or state-space abstraction, for which special precaution is required to ensure this. Our current implementation explores and eliminates states in a strict breadth-first search order. Motivated by the problematic performance of the Crowds protocol case study, we also want to consider different search orders so as to improve the behaviour for models for which the strict breadth-first search order does not perform well.

[7] https://www.prismmodelchecker.org/casestudies/embedded.php

References

1. Ashok, P., Butkova, Y., Hermanns, H., Kretínský, J.: Continuous-time Markov decisions based on partial exploration. In: ATVA. LNCS, vol. 11138, pp. 317–334. Springer (2018)
2. Bahar, R.I., Frohm, E.A., Gaona, C.M., Hachtel, G.D., Macii, E., Pardo, A., Somenzi, F.: Algebraic decision diagrams and their applications. Formal Methods in System Design 10(2/3), 171–206 (1997)
3. Baier, C., Hermanns, H., Katoen, J.P.: The 10,000 facets of MDP model checking. In: Computing and Software Science - State of the Art and Perspectives, LNCS, vol. 10000, pp. 420–451. Springer (2019)
4. Baier, C., Katoen, J.P.: Principles of model checking. MIT Press (2008)
5. Baier, C., Klein, J., Leuschner, L., Parker, D., Wunderlich, S.: Ensuring the reliability of your model checker: Interval iteration for Markov decision processes. In: CAV. LNCS, vol. 10426, pp. 160–180. Springer (2017)
6. Bohnenkamp, H.C., van der Stok, P., Hermanns, H., Vaandrager, F.W.: Cost-optimization of the IPv4 Zeroconf protocol. In: DSN. pp. 531–540. IEEE Computer Society (2003)
7. Bohy, A., Bruyère, V., Raskin, J., Bertrand, N.: Symblicit algorithms for mean-payoff and shortest path in monotonic Markov decision processes. Acta Inf. 54(6), 545–587 (2017)
8. Brázdil, T., Chatterjee, K., Chmelik, M., Forejt, V., Kretínský, J., Kwiatkowska, M.Z., Parker, D., Ujma, M.: Verification of Markov decision processes using learning algorithms. In: ATVA. LNCS, vol. 8837, pp. 98–114. Springer (2014)
9. Bryant, R.E.: Binary decision diagrams. In: Handbook of Model Checking, pp. 191–217. Springer (2018)
10. Brzozowski, J.A., McCluskey, E.J.: Signal flow graph techniques for sequential circuit state diagrams. IEEE Trans. Electronic Computers 12(2), 67–76 (1963)
11. Butkova, Y., Hartmanns, A., Hermanns, H.: A Modest approach to modelling and checking Markov automata. In: QEST. LNCS, vol. 11785, pp. 52–69. Springer (2019)
12. Chaki, S., Gurfinkel, A.: BDD-based symbolic model checking. In: Handbook of Model Checking, pp. 219–245. Springer (2018)
13. Daws, C.: Symbolic and parametric model checking of discrete-time Markov chains. In: ICTAC. LNCS, vol. 3407, pp. 280–294. Springer (2004)
14. Dehnert, C., Junges, S., Jansen, N., Corzilius, F., Volk, M., Bruintjes, H., Katoen, J.P., Ábrahám, E.: PROPhESY: A probabilistic parameter synthesis tool. In: CAV. LNCS, vol. 9206, pp. 214–231. Springer (2015)
15. Dehnert, C., Junges, S., Katoen, J.P., Volk, M.: A Storm is coming: A modern probabilistic model checker. In: CAV. LNCS, vol. 10427, pp. 592–600. Springer (2017)
16. Eisentraut, C., Hermanns, H., Zhang, L.: On probabilistic automata in continuous time. In: LICS. pp. 342–351. IEEE Computer Society (2010)
17. Gainer, P., Hahn, E.M., Schewe, S.: Accelerated model checking of parametric Markov chains. In: ATVA. LNCS, vol. 11138, pp. 300–316. Springer (2018)
18. Gainer, P., Hahn, E.M., Schewe, S.: Incremental verification of parametric and reconfigurable Markov chains. In: QEST. LNCS, vol. 11024, pp. 140–156. Springer (2018)
19. Gui, L., Sun, J., Song, S., Liu, Y., Dong, J.S.: SCC-based improved reachability analysis for Markov decision processes. In: ICFEM. LNCS, vol. 8829, pp. 171–186. Springer (2014)

20. Haddad, S., Monmege, B.: Reachability in MDPs: Refining convergence of value iteration. In: RP. LNCS, vol. 8762, pp. 125–137. Springer (2014)
21. Haddad, S., Monmege, B.: Interval iteration algorithm for MDPs and IMDPs. Theor. Comput. Sci. 735, 111–131 (2018)
22. Hahn, E.M., Hartmanns, A.: A comparison of time- and reward-bounded probabilistic model checking techniques. In: SETTA. LNCS, vol. 9984, pp. 85–100 (2016)
23. Hahn, E.M., Hartmanns, A., Hensel, C., Klauck, M., Klein, J., Kretínský, J., Parker, D., Quatmann, T., Ruijters, E., Steinmetz, M.: The 2019 comparison of tools for the analysis of quantitative formal models (qcomp 2019 competition report). In: 25 Years of TACAS: TOOLympics. LNCS, vol. 11429, pp. 69–92. Springer (2019)
24. Hahn, E.M., Hartmanns, A., Hermanns, H., Katoen, J.P.: A compositional modelling and analysis framework for stochastic hybrid systems. Formal Methods in System Design 43(2), 191–232 (2013)
25. Hahn, E.M., Hermanns, H., Wachter, B., Zhang, L.: PARAM: A model checker for parametric Markov models. In: CAV. LNCS, vol. 6174, pp. 660–664. Springer (2010)
26. Hahn, E.M., Li, Y., Schewe, S., Turrini, A., Zhang, L.: iscasmc: A web-based probabilistic model checker. In: FM 2014: Formal Methods - 19th International Symposium, Singapore, May 12-16, 2014. Proceedings. pp. 312–317 (2014)
27. Haring, G., Marie, R., Puigjaner, R., Trivedi, K.: Loss formulae and their application to optimization for cellular networks. IEEE Transaction on Vehicular Technology 50(3), 664 –673 (2000)
28. Hartmanns, A., Hermanns, H.: Explicit model checking of very large MDP using partitioning and secondary storage. In: Automated Technology for Verification and Analysis - 13th International Symposium, ATVA 2015, Shanghai, China, October 12-15, 2015, Proceedings. pp. 131–147 (2015)
29. Hartmanns, A., Kaminski, B.L.: Optimistic value iteration. CoRR abs/1910.01100 (2019), http://arxiv.org/abs/1910.01100
30. Helmink, L., Sellink, M., Vaandrager, F.: Proof-checking a data link protocol. In: Barendregt, H., Nipkow, T. (eds.) Proc. International Workshop on Types for Proofs and Programs (TYPES'93). LNCS, vol. 806, pp. 127–165. Springer (1994)
31. Holzmann, G.J.: The model checker SPIN. IEEE Trans. Software Eng. 23(5), 279–295 (1997)
32. Kwiatkowska, M.Z., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: CAV. LNCS, vol. 6806, pp. 585–591. Springer (2011)
33. Lee, C.Y.: Representation of switching circuits by binary-decision programs. The Bell System Technical Journal 38(4), 985–999 (1959)
34. Muppala, J., Ciardo, G., Trivedi, K.: Stochastic reward nets for reliability prediction. Communications in Reliability, Maintainability and Serviceability 1(2), 9–20 (July 1994)
35. Puterman, M.L.: Markov Decision Processes: Discrete Stochastic Dynamic Programming. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics, John Wiley & Sons Inc., New York (1994)
36. Quatmann, T., Katoen, J.P.: Sound value iteration. In: CAV. LNCS, vol. 10981, pp. 643–661. Springer (2018)
37. Reiter, M., Rubin, A.: Crowds: Anonymity for web transactions. ACM Transactions on Information and System Security (TISSEC 1(1), 66–92 (1998)
38. Wimmer, R., Braitling, B., Becker, B., Hahn, E.M., Crouzen, P., Hermanns, H., Dhama, A., Theel, O.E.: Symblicit calculation of long-run averages for concurrent probabilistic systems. In: QEST. pp. 27–36. IEEE Computer Society (2010)