

Are Good-for-Games Automata Good for Probabilistic Model Checking?*

Joachim Klein, David Müller, Christel Baier, and Sascha Klüppelholz

Institute of Theoretical Computer Science
Technische Universität Dresden, 01062 Dresden, Germany
{klein,david.mueller,baier,klueppel}@tcs.inf.tu-dresden.de

Abstract. The potential double exponential blow-up for the generation of deterministic ω -automata for linear temporal logic formulas motivates research on weaker forms of determinism. One of these notions is the *good-for-games* property that has been introduced by Henzinger and Piterman together with an algorithm for generating good-for-games automata from nondeterministic Büchi automata. The contribution of our paper is twofold. First, we report on an implementation of this algorithms and exhaustive experiments. Second, we show how good-for-games automata can be used for the quantitative analysis of systems modeled by Markov decision processes against ω -regular specifications and evaluate this new method by a series of experiments.

1 Introduction

The automata-theoretic approach to formal verification relies on the effective translation of specifications, e.g., formulas of some temporal logic such as linear temporal logic (LTL) into automata over infinite words (ω -automata) [34,6,12]. The verification problem for finite-state system models is then solvable by analyzing the product of the system model and the automaton for the formula. In the classical setting where the system model can be seen as a nondeterministic automaton, *nondeterministic* ω -automata suffice. For some applications, such as game-based synthesis and probabilistic model-checking problems, the nondeterminism of the ω -automaton poses a problem. Used as a monitor to determine the winning strategies of turn-based two-player games, the lack of look-ahead beyond the players' choices in general precludes the use of nondeterministic automata. Similarly, in probabilistic model checking, the lack of look-ahead beyond the probabilistic choices renders nondeterministic automata unsuitable in general. In these settings, the use of *deterministic* ω -automata resolves these problems at the cost of a further worst-case exponential determinization construction [26,31,25]. Thus, there is considerable interest in methods that tackle the worst-case double exponential time-complexity of algorithms caused by the

* Funded by the DFG through the Graduiertenkolleg 1763 (QuantLA), the CRC 912 HAEC, the cluster of excellence cFAED and the project QuaOS and by the ESF young researcher group IMData 100098198 and the EU-FP-7 grant 295261 (MEALS).

construction of deterministic ω -automata for LTL formulas. This includes for example variants of the determinization construction for nondeterministic Büchi automata (NBA) [28,32], heuristics [14,15] and the direct translation from fragments of LTL to deterministic automata [24,17,1]. Instead of reducing the number of states, [27] provides a translation from non-confluent NBA that aims to generate a compact symbolic representation of the generated deterministic automata based on binary decision diagrams (BDDs).

There are also several attempts to avoid determinization in certain scenarios [21,18] and provide better theoretical complexity and performance in practice. Henzinger and Piterman [13] introduce a special property for nondeterministic automata, being *good-for-games* (GFG), that is fulfilled by all deterministic automata but still permits nondeterministic choices. [13] proposes an algorithm, called the HP-algorithm here, for the construction of a nondeterministic GFG automaton with parity acceptance from an NBA that is amenable to a symbolic representation. The number of states in the constructed GFG automaton is still exponential in the number of states of the given NBA, but a smaller worst-case bound on the number of states can be provided than for Safra's determinization algorithm [31]. Among others, [4] introduced the notion *determinizable-by-pruning* for automata that have an embedded deterministic automaton for the same language. [4] states the existence of GFG automata that are not determinizable-by-pruning, but we are not aware of any result stating the existence of languages where GFG automata are more succinct than deterministic ones. To the best of our knowledge, the HP-algorithm is the sole published algorithm for the construction of GFG automata and it has not been implemented or experimentally evaluated yet.

In the context of probabilistic model checking for finite-state Markov chains, [8,3] propose the use of unambiguous automata that can be generated from LTL formulas with a single exponential blow-up in the worst case. Alternative approaches that also lead to single exponential-time model-checking algorithms for Markov chains and LTL specifications have been presented in [5] using weak alternating automata and in [7] using an iterative approach to integrate the effect of the temporal modalities of a given LTL formula φ in the Markov chain. Given that the analogous problem is 2EXPTIME-complete for models where nondeterministic and probabilistic choices alternate [7], there is no hope to generalize these results for Markov decision processes (MDPs). Only for the qualitative analysis of MDPs where the task is to show that an ω -regular path property holds with probability 1, no matter how the nondeterminism is resolved, Büchi automata that are deterministic-in-limit are shown to be sufficient [34,7].

Contribution. The purpose of our paper is to study whether GFG automata are adequate in the context of probabilistic model checking, both at the theoretical and the practical level. At the theoretical level, we answer in the affirmative and provide algorithms for the computation of maximal or minimal probabilities for path properties specified by GFG automata in finite-state Markov decision processes (MDPs). The time complexity of our algorithm is polynomial in the size of the given MDP and GFG automaton. To evaluate the GFG-based approach

empirically, we have implemented the HP-algorithm (and various variants) symbolically using binary decision diagrams (BDDs). In a series of experiments, we study the performance of the HP-algorithm – from LTL formula via NBA to GFG automaton – compared to the determinization implementation of LTL2DSTAR [14,15] based on Safra’s construction. We have furthermore implemented the GFG-based approach for the analysis of MDPs in the popular probabilistic model checker PRISM [22] and evaluated its performance in practice.

Outline. Section 2 briefly introduces our notations for ω -automata and MDPs. The applicability of GFG automata for the quantitative analysis of MDPs is shown in Section 3. In Section 4, we study the HP-algorithm in detail and present a few heuristics that have been integrated in our implementation. Section 5 reports on our experiments, Section 6 contains some concluding remarks. Omitted proofs and other additional material can be found in the technical report [16].

2 Preliminaries

Throughout the paper, the reader is supposed to be familiar with the basic principles of ω -automata, games and temporal logics. For details we refer to [6,12]. We briefly summarize our notations for ω -automata, present the definition of good-for-games automata [13] and provide a condensed survey of the relevant principles of Markov decision processes (MDPs). Further details on MDPs and their use in the context of model checking can be found e.g. in [30,2].

Automata over Infinite Words. An ω -automaton $\mathcal{A} = (Q, \Sigma, \delta, q_0, Acc)$ is a tuple, where Q is a finite set of states, Σ is a finite alphabet, $\delta : Q \times \Sigma \rightarrow 2^Q$ is the (nondeterministic) transition function and $q_0 \in Q$ is the initial state. The last component Acc is the acceptance condition (see below). The size of $|\mathcal{A}|$ denotes the number of states in \mathcal{A} . \mathcal{A} is said to be *complete*, if $\delta(q, \sigma) \neq \emptyset$ for all states $q \in Q$ and all symbols $\sigma \in \Sigma$. \mathcal{A} is called *deterministic*, if $|\delta(q, \sigma)| \leq 1$ for all $q \in Q$ and $\sigma \in \Sigma$. A *run* in \mathcal{A} for an infinite word $w = \sigma_0 \sigma_1 \sigma_2 \dots \in \Sigma^\omega$ is a sequence $\rho = q_0 q_1 q_2 \dots \in Q^\omega$ starting in the initial state q_0 such that $q_{i+1} \in \delta(q_i, a)$ for all $i \in \mathbb{N}$. We write $\text{inf}(\rho)$ to denote the set of all states occurring infinitely often in ρ . A run ρ is called *accepting*, if it meets the acceptance condition Acc , denoted $\rho \models Acc$. We consider here the following three types of acceptance conditions and describe their constraints for infinite runs:

- *Büchi*: $Acc = F$ is a set of states, i.e., $F \subseteq Q$, with the meaning $\Box \Diamond F$
- *parity*: Acc is a function $col : Q \rightarrow \mathbb{N}$ assigning to each state q a parity color and requiring that the least parity color appearing infinitely often is even
- *Rabin*: Acc is a set consisting of pairs (E, F) with $E, F \subseteq Q$, imposing the constraint $\bigvee_{(E,F) \in Acc} (\Diamond \Box \neg E \wedge \Box \Diamond F)$

Büchi acceptance can be seen as a special case of parity acceptance which again can be seen as a special case of Rabin acceptance. We use the standard notations NBA (NRA, NPA) for nondeterministic Büchi (Rabin, parity) automata and

DBA, DRA, DPA for their deterministic versions. The language of \mathcal{A} , denoted $\mathcal{L}(\mathcal{A})$, consists of all infinite words $w \in \Sigma^\omega$ that have at least one accepting run in \mathcal{A} , i.e., $w \in \mathcal{L}(\mathcal{A})$ iff there exists a run ρ for w with $\rho \models Acc$.

It is well-known that the classes of languages recognizable by NBA, NRA, NPA, DRA or DPA are the same (the so-called ω -regular languages), while DBA are less powerful. For each LTL formula φ with atomic propositions in some finite set AP , the semantics of φ can be described as an ω -regular language $\mathcal{L}(\varphi)$ over the alphabet $\Sigma = 2^{AP}$ and there is an NBA \mathcal{A} for φ (i.e., $\mathcal{L}(\varphi) = \mathcal{L}(\mathcal{A})$) whose size is at most exponential in the formula length $|\varphi|$.

Good-for-Games (GFG) Automata. The formal definition of GFG automata [13] relies on a game-based view of ω -automata. Given an ω -automaton \mathcal{A} as before, we consider \mathcal{A} as the game arena of a turn-based two-player game, called *monitor game*: if the current state is q then player 1 chooses a symbol $\sigma \in \Sigma$ whereas the other player (player 0) has to answer by a successor state $q' \in \delta(q, \sigma)$, i.e., resolve the nondeterminism. In the next round q' becomes the current state. A *play* is an alternating sequence $\varsigma = q_0 \sigma_0 q_1 \sigma_1 q_2 \sigma_2 \dots$ of states and (action) symbols in the alphabet Σ starting with the initial state q_0 . Intuitively, the σ_i 's are the symbols chosen by player 1 and the q_i 's are the states chosen by player 0 in round i . Player 0 wins the play ς if ς is infinite and if $\varsigma|_\Sigma = \sigma_0 \sigma_1 \sigma_2 \dots \in \mathcal{L}(\mathcal{A})$ then $\varsigma|_Q = q_0 q_1 q_2 \dots$ is an accepting run. A strategy for player 0 is a function $f : (Q \times \Sigma)^+ \rightarrow Q$ with $f(\dots q \sigma) \in \delta(q, \sigma)$. A play $\varsigma = q_0 \sigma_0 q_1 \sigma_1 q_2 \dots$ is said to be *f-conform* or a *f-play* if $q_i = f(\varsigma \downarrow i)$ for all $i \geq 1$ where $\varsigma \downarrow i = q_0 \sigma_0 \dots \sigma_{i-2} \dots q_{i-1} \sigma_i$ is the prefix of ρ that ends with the chosen symbol in round i . An automaton \mathcal{A} is called *good-for-games* if there is a strategy f such that player 0 wins each f -play. Such strategies will be called *GFG-strategies* for \mathcal{A} . Obviously, each deterministic automaton enjoys the GFG property. GFG automata with Rabin or parity condition cover the full class of ω -regular languages, while GFG automata with Büchi acceptance do not [4]. For illustrating examples of GFG automata see [16].

Markov Decision Processes (MDP). MDPs are an operational model for systems that exhibit nondeterministic and probabilistic choices. For the purposes of this paper, we formalize an MDP by a tuple $\mathcal{M} = (S, Act, P, s_0, AP, \ell)$ where S is a finite set of states, $s_0 \in S$ is the initial state, Act a finite set of actions and $P : S \times Act \times S \rightarrow [0, 1]$ is the transition probability function satisfying:

$$\sum_{s' \in S} P(s, \alpha, s') \in \{0, 1\} \quad \text{for all } s \in S, \alpha \in Act.$$

We write $Act(s)$ for the set of actions α that are enabled in s , i.e., $P(s, \alpha, s') > 0$ for some $s' \in S$, in which case $s' \mapsto P(s, \alpha, s')$ is a distribution formalizing the probabilistic effect of taking action α in state s . We refer to the triples (s, α, s') with $P(s, \alpha, s') > 0$ as a step. The choice between the enabled actions is viewed to be nondeterministic. For technical reasons, we require $Act(s) \neq \emptyset$ for all states s . The last two components AP and ℓ serve to formalize properties of paths in \mathcal{M} . Formally, AP is a finite set of atomic propositions and $\ell : S \rightarrow 2^{AP}$ assigns to each state s the set $\ell(s)$ of atomic propositions that hold in s . Paths

in \mathcal{M} are finite or infinite sequences $\pi = s_0 \alpha_0 s_1 \alpha_2 s_2 \alpha_3 \dots$ starting in the initial state s_0 that are built by consecutive steps, i.e., $P(s_i, \alpha_i, s_{i+1}) > 0$ for all i . The trace of π is the word over the alphabet $\Sigma = 2^{AP}$ that arises by taking the projections to the state labels, i.e., $trace(\pi) = \ell(s_0) \ell(s_1) \ell(s_2) \dots$. For an LTL formula φ over AP we write $\pi \models \varphi$ if $trace(\pi) \in \mathcal{L}(\varphi)$.

As the monitor game in nondeterministic automata, MDPs can be seen as stochastic games, also called a $1\frac{1}{2}$ -player games. The first (full) player resolves the nondeterministic choice by selecting an enabled action α of the current state s . The second (half) player behaves probabilistically and selects a successor state s' with $P(s, \alpha, s') > 0$. Strategies for the full player are called *schedulers*. Since the behavior of \mathcal{M} is purely probabilistic if some scheduler \mathfrak{s} is fixed, one can reason about the probability of path events. If L is an ω -regular language then $\Pr_{\mathcal{M}}^{\mathfrak{s}}(L)$ denotes the probability under \mathfrak{s} for the set of infinite paths π with $trace(\pi) \in L$. In notations like $\Pr_{\mathcal{M}}^{\mathfrak{s}}(\varphi)$ or $\Pr_{\mathcal{M}}^{\mathfrak{s}}(\mathcal{A})$ we identify LTL formulas φ and ω -automata \mathcal{A} with their languages. For the mathematical details of the underlying sigma-algebra and probability measure, we refer to [30,2].

For a worst-case analysis of a system modeled by an MDP \mathcal{M} , one ranges over all initial states and all schedulers (i.e., all possible resolutions of the nondeterminism) and considers the maximal or minimal probabilities for some ω -regular language L . Depending on whether L represents a desired or undesired path property, the quantitative worst-case analysis amounts to computing $\Pr_{\mathcal{M}}^{\min}(\varphi) = \min_{\mathfrak{s}} \Pr_{\mathcal{M}}^{\mathfrak{s}}(L)$ or $\Pr_{\mathcal{M}}^{\max}(L) = \max_{\mathfrak{s}} \Pr_{\mathcal{M}}^{\mathfrak{s}}(L)$.

3 Automata-Based Analysis of Markov Decision Processes

We address the task to compute the maximal or minimal probability in an MDP \mathcal{M} for the path property imposed by a nondeterministic ω -automaton \mathcal{A} . The standard approach, see e.g. [2], assumes \mathcal{A} to be deterministic and relies on a product construction where the transitions of \mathcal{M} are simply annotated with the unique corresponding transition in \mathcal{A} . Thus, $\mathcal{M} \otimes \mathcal{A}$ can be seen as a refinement of \mathcal{M} since \mathcal{A} does not affect \mathcal{M} 's behaviors, but attaches information on \mathcal{A} 's current state for the prefixes of the traces induced by the paths of \mathcal{M} .

We now modify the standard definition of the product for nondeterministic ω -automaton. The crucial difference is that the actions are now pairs $\langle \alpha, p \rangle$ consisting of an action in \mathcal{M} and a state in \mathcal{A} , representing the nondeterministic alternatives in both the MDP \mathcal{M} and the automaton \mathcal{A} . Formally, let $\mathcal{M} = (S, Act, P, s_0, AP, \ell)$ be an MDP and $\mathcal{A} = (Q, \Sigma, \delta, q_0, Acc)$ a complete nondeterministic ω -automaton with $\Sigma = 2^{AP}$. The product MDP is

$$\mathcal{M} \otimes \mathcal{A} = (S \times Q, Act \times Q, P', \langle s_0, q_0 \rangle, AP, \ell')$$

where the transition probability function P' is given by $P'(\langle s, q \rangle, \langle \alpha, p \rangle, \langle s', q' \rangle) = P(s, \alpha, s')$ if $p = q' \in \delta(q, \ell(s))$. In all other cases $P'(\langle s, q \rangle, \langle \alpha, p \rangle, \langle s', q' \rangle) = 0$. The assumption that \mathcal{A} is complete yields that for each $\alpha \in Act(s)$ there is some action $\langle \alpha, q' \rangle \in Act(\langle s, q \rangle)$ for all states s in \mathcal{M} and q in \mathcal{A} . The labeling function is given by $\ell'(\langle s, q \rangle) = \{q\}$. Thus, the traces in $\mathcal{M} \otimes \mathcal{A}$ are words over the

alphabet Q . Likewise, \mathcal{A} 's acceptance condition Acc can be seen as a language over Q , which permits to treat Acc as a property that the paths in $\mathcal{M} \otimes \mathcal{A}$ might or might not have. We prove in [16]:

Theorem 1. *For each MDP \mathcal{M} and nondeterministic ω -automaton \mathcal{A} as above:*

- (a) $\Pr_{\mathcal{M} \otimes \mathcal{A}}^{\max}(Acc) \leq \Pr_{\mathcal{M}}^{\max}(\mathcal{A})$
- (b) *If \mathcal{A} is good-for-games then: $\Pr_{\mathcal{M} \otimes \mathcal{A}}^{\max}(Acc) = \Pr_{\mathcal{M}}^{\max}(\mathcal{A})$*

Theorem 1 (b) shows that with a slightly modified definition of the product, the techniques that are known for the quantitative analysis of MDPs against deterministic ω -automata specifications are also applicable for GFG automata. The computation of maximal probabilities for properties given by an ω -regular acceptance condition Acc (e.g., Büchi, Rabin or parity) can be carried out by a graph analysis that replaces Acc with a reachability condition and linear programming techniques for computing maximal reachability probabilities. See e.g. [2]. The time complexity is polynomial in the size of the \mathcal{M} and \mathcal{A} . Thus, if the specification is given in terms of an LTL formula φ then the costs of our GFG-based approach are dominated by the generation of a GFG automaton for φ . Minimal probabilities can be handled by using $\Pr_{\mathcal{M}}^{\min}(\varphi) = 1 - \Pr_{\mathcal{M}}^{\max}(\neg\varphi)$.

[20,19] proves that a double exponential blow-up for translating LTL to deterministic ω -automata (of any type) is unavoidable. We adapted the proof in [19] for GFG automata (see [16]). Thus, the double exponential time complexity of the GFG-based approach is in accordance with the known 2EXPTIME-completeness for the analysis of MDPs against LTL specifications [7].

Theorem 2. *There exists a family of LTL formulas $(\varphi)_{n \in \mathbb{N}}$ such that $|\varphi_n| = \mathcal{O}(n)$, while every GFG automaton \mathcal{A}_n for φ_n has at least $2^{2^{\Omega(n)}}$ states.*

4 From LTL to GFG Automata

We have previously shown [14,15] that it is possible in practice, using the tool LTL2DSTAR, to obtain deterministic ω -automata for a wide range of LTL formula φ via the translation to an NBA and Safra's determinization construction [31] refined by various heuristics. Here, we are interested in replacing Safra's determinization algorithm with the HP-algorithm [13] to generate a GFG automaton instead of a deterministic automaton. We first provide an outline of the HP-algorithm and then explain a few new heuristics.

The HP-algorithm transforms an NBA $\mathcal{B} = (Q, \Sigma, \delta, q_0, F)$ with $|Q| = n$ states into a GFG automaton \mathcal{A} with parity acceptance and at most $2^n \cdot n^{2^n}$ states and $2n$ parity colors (or an NRA with n Rabin pairs), which improves on the upper bound given for Safra's determinization algorithm. We recall here the main concepts, for a formal description we refer to [13]. Like Safra's construction, the HP-algorithm relies on the simultaneous tracking of multiple subset constructions to determine acceptance or rejection in the NBA. However, while the states of Safra's DRA organize the subsets in trees, the HP-algorithm uses

a simpler, linear arrangement of the subsets. The state space $P = (2^Q \times 2^Q)^n$ of the GFG automaton \mathcal{A} consists of n pairs of subsets of NBA states Q , i.e., states of the form $p = \langle (A_1, B_1), \dots, (A_n, B_n) \rangle$ where $B_i \subseteq A_i \subseteq Q$, plus some additional constraints on the state space. Each set B_i serves to mark those states in A_i that were reached via some accepting state in F of the NBA. The successor state in \mathcal{A} for symbol σ is obtained by applying the transition function δ to each of the subsets and adding states in F to the B_i subsets. In crucial difference to Safra’s construction, the HP-algorithm however then introduces significant nondeterminism by allowing \mathcal{A} to discard an arbitrary number of states in any of the subsets. For $p = \langle \dots (A_i, B_i) \dots \rangle$, the set A'_i in a σ -successor p' of \mathcal{A} thus does not correspond to $A'_i = \delta(A_i, \sigma)$ but there is a nondeterministic choice between any A'_i satisfying $A'_i \subseteq \delta(A_i, \sigma)$, including the empty set. Whenever some A_i is empty, \mathcal{A} can “reset” A_i by setting A_i to some subset of the first set A_1 . Such resets are reflected in the acceptance condition of \mathcal{A} as “bad” events for the pair i , as they signify that the previously tracked runs terminated. The “good” events in the acceptance condition occur whenever all states in an A_i are marked as having recently visited F , i.e., whenever $A_i = B_i \neq \emptyset$. In the next step, B'_i is then cleared and the tracking of visits to F starts anew. Infinitely many “good” events without “bad” events then correspond to the existence of an accepting run in the NBA \mathcal{B} . The HP-algorithm relies on the GFG-strategy to resolve the nondeterminism in the constructed automaton \mathcal{A} , i.e., which states in the subsets are kept, which are dropped and when to reset. There is a large amount of nondeterminism and a lot of combinatorial possibilities in the reachable state space of \mathcal{A} . This is confirmed by our experiments, e.g., applying the construction to the two-state NBA for $\diamond\Box a$ already yields a GFG automaton with 16 states, where LTL2DSTAR generates a two-state DRA. As stated in [13], the HP-algorithm is thus not well-suited for an explicit representation for \mathcal{A} , but is intended for a symbolic implementation. In this context, [13] briefly discusses the possibility of variants of the transition function in the GFG automaton that either apply more or less strict constraints on the relationship enforced between the (A_i, B_i) pairs in each state. In particular, [13] posits that introducing even further nondeterminism (and increasing the number of possible states) by loosening a disjunctness requirement on the A_i may lead to a smaller symbolic representation. In our experiments, we will refer to this as the *loose variant*.

Iterative Approach. In the context of games, [13] proposes an iterative approach to the HP-algorithm by successively constructing the automata \mathcal{A}^m obtained by using only the first m of the n pairs, i.e., by setting $A_i = B_i = \emptyset$ for all $m < i \leq n$. In the acceptance condition this reduces the number of required parity colors to $2m$ and Rabin pairs to m as well. For these automata, $\mathcal{L}(\mathcal{A}^m) = \mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{B})$, but there is no guarantee that \mathcal{A}^m for $m < n$ is good-for-games by construction. We start with $m = 1$ and increase m until early success or reaching $m = n$. Our experimental results indeed show that early termination appears rather often.

We now explain how the iterative approach of [13] can be integrated in the GFG-based quantitative analysis of MDPs against LTL specifications. Suppose,

e.g., that the task is to show that $\Pr_{\mathcal{M}}^{\max}(\varphi) \geq \theta$ for some LTL formula φ and threshold $\theta \in]0, 1]$. Let \mathcal{B} be an n -state NBA with $\mathcal{L}(\mathcal{B}) = \mathcal{L}(\varphi)$ and \mathcal{A}^m the automaton obtained using only the first $m \leq n$ pairs in the HP-algorithm applied to \mathcal{B} . Let Acc^m denote the acceptance condition of \mathcal{A}^m . By Theorem 1 (a):

$$\text{If } \Pr_{\mathcal{M} \otimes \mathcal{A}^m}^{\max}(Acc^m) \geq \theta \text{ for some } m \leq n \text{ then } \Pr_{\mathcal{M}}^{\max}(\varphi) \geq \theta.$$

Moreover, $\Pr_{\mathcal{M} \otimes \mathcal{A}^m}^{\max}(Acc^m) \leq \Pr_{\mathcal{M} \otimes \mathcal{A}^{m+1}}^{\max}(Acc^{m+1})$ for $m < n$. These observations suggest an approach that resembles the classical *abstraction-refinement* schema: starting with $m = 1$, we carry out the quantitative analysis of $\mathcal{M} \otimes \mathcal{A}^m$ against Acc^m and successively increase m until $\Pr_{\mathcal{M} \otimes \mathcal{A}^m}^{\max}(Acc^m) \geq \theta$ or \mathcal{A}^m is GFG (which is the case at the latest when $m = n$). As an additional heuristic to increase the performance of the linear programming techniques that are applied for the quantitative analysis of $\mathcal{M} \otimes \mathcal{A}^m$ against Acc^m , one can reuse the results computed for $\mathcal{M} \otimes \mathcal{A}^{m-1}$ and Acc^{m-1} as initial values.

It remains to explain how to check whether \mathcal{A}^m has the GFG property. For details we refer to [16]. In this aspect, our prototype implementation departs from [13] and checks whether \mathcal{A}^m is GFG by solving a Rabin game (itself an NP-complete problem) constructed from \mathcal{A}^m and a DRA for $\neg\varphi$ constructed with LTL2DSTAR while [13] proposes an algorithm based on checking fair simulation. To study the impact of the iterative approach in terms of the number of required iterations and the size of the resulting GFG automata, the choice of the GFG test is irrelevant.

Union Operator for Disjunctive Formulas. For generating a deterministic automaton from an LTL formula, we have shown in [14] that optionally handling disjunctive LTL formulas of the form $\varphi = \varphi_1 \vee \varphi_2$ by constructing DRA \mathcal{A}_1 and \mathcal{A}_2 for the subformulas φ_1 and φ_2 and then obtaining the DRA $\mathcal{A}_1 \cup \mathcal{A}_2$ for the language $\mathcal{L}(\mathcal{A}_1) \cup \mathcal{L}(\mathcal{A}_2)$ via a product construction can be very beneficial in practice. The definition of $\mathcal{A}_1 \cup \mathcal{A}_2$ used in [14] can easily be extended to NRA. The GFG property is preserved by the union construction. See [16].

5 Implementation and Experiments

We have implemented the HP-algorithm in a tool we refer to as LTL2GFG. Based on LTL2GFG, we have additionally implemented the GFG-based quantitative analysis of MDPs in PRISM. After a brief overview of LTL2GFG, we report on our experiments and comparison with the determinization approach of LTL2DSTAR.

LTL2GFG. Given an LTL formula φ , our implementation LTL2GFG constructs a symbolic, BDD-based representation of a GFG-NPA for φ . It first converts φ into an (explicitly represented) NBA \mathcal{B} . In our experiments, we use LTL2BA v1.1 [11] for this task. To facilitate an efficient symbolic representation of the various subsets used in the HP-algorithm, \mathcal{B} is then converted to a symbolic representation, using a unary encoding of the $|Q| = n$ states of \mathcal{B} , i.e., using one boolean variable q_i per state. The state space of the GFG-automaton \mathcal{A} , i.e., the n pairs (A_i, B_i) is likewise encoded by n^2 boolean variables $a_{i,j}$ and $b_{i,j}$, i.e.,

Table 1. Statistics for the automata \mathcal{A}_φ constructed for the 94 benchmark formulas. Number of \mathcal{A}_φ constructed within a given timeframe and a given range of BDD sizes.

	aborted	\mathcal{A}_φ with constr. time				\mathcal{A}_φ with BDD size					
		< 1s	< 10s	< 1m	< 30m	< 10	< 10 ²	< 10 ³	< 10 ⁴	< 10 ⁵	$\geq 10^5$
LTL2DSTAR std.	0	90	91	92	94	4	65	87	90	91	3
no opt.	0	90	90	92	94	3	48	78	89	90	4
LTL2GFG std.	39	40	47	48	55	3	6	19	26	36	19
std., dynamic	45	34	36	48	49	5	8	19	36	39	10
loose, dynamic	34	43	49	56	60	5	14	31	47	56	4
lo., union, dyn.	29	52	59	61	65	4	13	35	54	60	5
lo., iterative	20	74	74	74	74	3	19	39	60	74	0
lo., it., un., dyn.	18	70	72	74	76	4	32	63	70	76	0

$a_{i,j}$ is true iff NBA state $q_j \in A_i$ and $b_{i,j}$ is true iff $q_j \in B_i$ for $1 \leq i, j \leq n$. To allow the encoding of the transition relations of \mathcal{A} and \mathcal{B} , each state variable has a primed copy, i.e., $q'_i, a'_{i,j}$ and $b'_{i,j}$ and each of the k atomic proposition in φ is represented by a boolean variable l_i . For a BDD-based symbolic representation, the order of the variables is crucial. The state variables and their copies are always kept adjacent. The standard variable ordering used by LTL2GFG is then an interleaving of the $a_{i,j}$ and $b_{i,j}$ variables with the q_j variables, i.e., $l_1 < \dots < l_k < q_1 < \dots < q_j < a_{1,j} < b_{1,j} < a_{2,j} < b_{2,j} < \dots < q_{j+1} < \dots$. LTL2GFG uses the JINC C++ BDD library for the symbolic representation.

Experimental Results for the HP-algorithm. We report here on a number of experiments with LTL2GFG using the benchmark formulas used in the evaluation of LTL2DSTAR in [14,15], i.e., 39 LTL formulas from the literature [10,33] and 55 pattern formulas [9] that represent common specification patterns. All our experiments were carried out on a computer with 2 Intel E5-2680 8-core CPUs at 2.70 GHz with 384GB of RAM running Linux and with a memory limit of 10 GB and a time-out of 30 minutes for each formula.

For the automata \mathcal{A}_φ , we report on the number of BDD nodes in the encoding of the transition function, as this the most crucial aspect. To allow a fair comparison with the explicit determinization in LTL2DSTAR, we consider symbolic encodings of the DRA \mathcal{A}_φ obtained from LTL2DSTAR 0.5.1. This encoding uses $\lceil \log_2 n \rceil$ boolean variables to straightforwardly encode the n state indices in \mathcal{A}_φ , which is the same encoding employed in PRISM for its DRA-based approach to LTL model checking.

Table 1 presents statistics for the construction of DRA with LTL2DSTAR and GFG-NPA with LTL2GFG for the benchmark formulas. The LTL2DSTAR results are given once with standard settings and for a variant where all optimizations are disabled, i.e., with purely Safra’s construction. For LTL2GFG, we start with the pure HP-algorithm and consider variants with the “loose” transition definition, the union construction, and with dynamic reordering of the variable order. We also give statistics for the iterative approach, where LTL2GFG constructs the partial automata \mathcal{A}^m until it can be shown (via solving a Rabin game [29]) that the automaton is GFG.

Table 2. Results of the iterative approach in LTL2GFG, for the loose variant. M is the minimal value $m \leq n$ for which the partial NPA \mathcal{A}^m could be shown to be GFG.

	with n NBA states											
	2	3	4	5	6	7	8	9	10	11	12	> 12
number of φ	13	17	13	9	8	3	3	1	4	2	4	11
number of φ , $M < n$	11	17	13	8	8	2	2	1	0	0	1	3
number of φ , $M = 1$	11	8	5	4	2	1	1	0	0	0	1	3
number of φ , $M = 2$	2	9	8	4	6	1	1	1	0	0	0	0
number of φ , GFG check aborted	0	0	0	1	0	1	1	0	4	2	3	8

LTL2DSTAR constructed most of the automata in a few seconds, the most difficult was constructed in 95s and had 1.2 million BDD nodes. Apart from the most difficult automata, the BDD sizes range in the hundreds and thousands. For all the LTL2GFG variants, a significant fraction of automata could not be constructed in the time and memory limits, around 40% for the standard HP-algorithm, and dropping to around 20% for the best variant. The loose variant by itself had a mixed effect, but in conjunction with dynamic reordering was generally beneficial. The union construction was very beneficial for the disjunctive formulas. For example, the automata for $\Box\Diamond a \rightarrow \Box\Diamond b$ could not be constructed in the time limits with the standard HP-algorithm but could be handled using the union construction. The iterative approach was successful as well in obtaining smaller automata, which is explained by the fact that for a large number of formulas it could be shown that the partial automata \mathcal{A}^1 or \mathcal{A}^2 were already GFG, as detailed in Table 2. For the iterative approach we were mostly focused on experimental data for the minimal value m for which \mathcal{A}^m becomes GFG and the effect on the BDD size. Different algorithms or implementations for the GFG check than the one used in LTL2GFG lead to the same final GFG automata, but could improve the performance. At the end, despite the various approaches implemented in LTL2GFG, there were only 6 formulas with relatively small automata where the BDD size of the smallest GFG automaton was smaller than that of the DRA obtained from LTL2DSTAR (172 nodes instead of 229 nodes, 219 instead of 347, and the other 4 automata differing by 1 or 2 at a size of less than 20 nodes). We do not report here in detail on the number of reachable states in the automata, as none of the GFG automata had a smaller number of states than the DRA generated by LTL2DSTAR. In particular, the automata obtained without the iterative approach often had millions and more states.

Implementation in PRISM. We have extended the MTBDD-based, symbolic engines of PRISM 4.1 with an implementation of our algorithm for computing $\Pr_{\mathcal{M}}^{\max}(\varphi)$ using GFG automata for φ (and $\Pr_{\mathcal{M}}^{\min}(\varphi)$ using a GFG automaton for $\neg\varphi$). We import the BDD of \mathcal{A} generated with LTL2GFG into PRISM and perform the product with \mathcal{M} and analysis in $\mathcal{M} \otimes \mathcal{A}$ symbolically. In its standard approach, PRISM constructs an explicit DRA with an integrated version of LTL2DSTAR, which is then symbolically encoded as described before. The analysis is then carried out symbolically as well.

Experiments in PRISM. As a benchmark, we used a PRISM model [23] for parts of the WLAN carrier-sense protocol of IEEE 802.11. As was to be expected given our results on the automata construction, the GFG-based analysis did not improve on the standard approach. Even using the optimal variant of LTL2GFG for each formula, ignoring the automata construction times, and for cases where the product $\mathcal{M} \otimes \mathcal{A}$ had a comparable BDD size for the GFG- and DRA-based approach, the model checking using the GFG automata took significantly longer. For further details, we refer to [16].

6 Conclusion

We have shown that GFG automata can replace deterministic automata for the quantitative analysis of MDPs against ω -regular specifications without increasing the asymptotic worst-case time complexity. To evaluate the GFG-based approach from the practical side, we implemented the HP-algorithm, integrated several heuristics, and performed exhaustive experiments for the LTL-to-GFG construction and for probabilistic model checking. Our experimental results are a bit disappointing, as the generated GFG automata were often larger than DRA generated by the implementation of Safra's algorithms in LTL2DSTAR, both in the number of states and in the symbolic BDD-based representations. Thus, our empirical results are in contrast to the expectation that the HP-algorithm yields GFG automata that are better suited for symbolic approaches rather than DRA generated by Safra's algorithm. Also in the context of probabilistic model checking, the GFG-based approach turned out to be more time- and memory-consuming than the traditional approach with deterministic automata. However, it is still too early to discard the concept of GFG automata for practical purposes. Our negative empirical results might be an artefact of the HP-algorithm, which is – to the best of our knowledge – the only known algorithm for the generation of GFG automata that are not deterministic. Future directions are the design of other algorithms for the construction of succinct GFG automata. Alternatively, one might seek for automata types that are still adequate for probabilistic model checking and other areas, but rely on weaker conditions than the GFG property.

References

1. Babiak, T., Blahoudek, F., Křetínský, M., Strejček, J.: Effective translation of LTL to deterministic rabin automata: Beyond the (F,G)-fragment. In: Van Hung, D., Ogawa, M. (eds.) ATVA 2013. LNCS, vol. 8172, pp. 24–39. Springer, Heidelberg (2013)
2. Baier, C., Katoen, J.P.: Principles of Model Checking. MIT Press (2008)
3. Benedikt, M., Lenhardt, R., Worrell, J.: Two variable vs. linear temporal logic in model checking and games. Logical Methods in Computer Science 9(2) (2013)
4. Boker, U., Kuperberg, D., Kupferman, O., Skrzypczak, M.: Nondeterminism in the presence of a diverse or unknown future. In: Fomin, F.V., Freivalds, R., Kwiatkowska, M., Peleg, D. (eds.) ICALP 2013, Part II. LNCS, vol. 7966, pp. 89–100. Springer, Heidelberg (2013)

5. Bustan, D., Rubin, S., Vardi, M.Y.: Verifying ω -regular properties of Markov chains. In: Alur, R., Peled, D.A. (eds.) CAV 2004. LNCS, vol. 3114, pp. 189–201. Springer, Heidelberg (2004)
6. Clarke, E., Grumberg, O., Peled, D.: Model Checking. MIT Press (2000)
7. Courcoubetis, C., Yannakakis, M.: The complexity of probabilistic verification. *Journal of ACM* 42(4), 857–907 (1995)
8. Couvreur, J.M., Saheb, N., Sutre, G.: An optimal automata approach to LTL model checking of probabilistic systems. In: Vardi, M.Y., Voronkov, A. (eds.) LPAR 2003. LNCS, vol. 2850, pp. 361–375. Springer, Heidelberg (2003)
9. Dwyer, M., Avrunin, G., Corbett, J.: Patterns in property specifications for finite-state verification. In: ICSE 1999, pp. 411–420. ACM (1999)
10. Etesami, K., Holzmann, G.J.: Optimizing Büchi automata. In: Palamidessi, C. (ed.) CONCUR 2000. LNCS, vol. 1877, pp. 153–167. Springer, Heidelberg (2000)
11. Gastin, P., Oddoux, D.: Fast LTL to Büchi automata translation. In: Berry, G., Comon, H., Finkel, A. (eds.) CAV 2001. LNCS, vol. 2102, pp. 53–65. Springer, Heidelberg (2001)
12. Grädel, E., Thomas, W., Wilke, T. (eds.): Automata, Logics, and Infinite Games. LNCS, vol. 2500. Springer, Heidelberg (2002)
13. Henzinger, T., Piterman, N.: Solving games without determinization. In: Ésik, Z. (ed.) CSL 2006. LNCS, vol. 4207, pp. 395–410. Springer, Heidelberg (2006)
14. Klein, J., Baier, C.: Experiments with deterministic ω -automata for formulas of linear temporal logic. *Theoretical Computer Science* 363(2), 182–195 (2006)
15. Klein, J., Baier, C.: On-the-fly stuttering in the construction of deterministic ω -automata. In: Holub, J., Žďárek, J. (eds.) CIAA 2007. LNCS, vol. 4783, pp. 51–61. Springer, Heidelberg (2007)
16. Klein, J., Müller, D., Baier, C., Klüppelholz, S.: Are good-for-games automata good for probabilistic model checking (extended version). Tech. rep., Technische Universität Dresden (2013), <http://www.tcs.inf.tu-dresden.de/ALGI/PUB/LATA14/>
17. Křetínský, J., Garza, R.L.: Rabinizer 2: Small deterministic automata for LTL\GU. In: Van Hung, D., Ogawa, M. (eds.) ATVA 2013. LNCS, vol. 8172, pp. 446–450. Springer, Heidelberg (2013)
18. Kupferman, O., Piterman, N., Vardi, M.: Safraless compositional synthesis. In: Ball, T., Jones, R.B. (eds.) CAV 2006. LNCS, vol. 4144, pp. 31–44. Springer, Heidelberg (2006)
19. Kupferman, O., Rosenberg, A.: The blow-up in translating LTL to deterministic automata. In: van der Meyden, R., Smaus, J.-G. (eds.) MoChArt 2010. LNCS, vol. 6572, pp. 85–94. Springer, Heidelberg (2011)
20. Kupferman, O., Vardi, M.: From linear time to branching time. *ACM Transactions on Computational Logic* 6(2), 273–294 (2005)
21. Kupferman, O., Vardi, M.: Safraless decision procedures. In: FOCS 2005, pp. 531–542. IEEE Computer Society (2005)
22. Kwiatkowska, M., Norman, G., Parker, D.: Probabilistic symbolic model checking with PRISM: A hybrid approach. *International Journal on Software Tools for Technology Transfer (STTT)* 6(2), 128–142 (2004)
23. Kwiatkowska, M., Norman, G., Sproston, J.: Probabilistic model checking of the IEEE 802.11 wireless local area network protocol. In: Hermanns, H., Segala, R. (eds.) PROBMIV 2002, PAPM-PROBMIV 2002, and PAPM 2002. LNCS, vol. 2399, pp. 169–187. Springer, Heidelberg (2002)
24. Latvala, T.: Efficient model checking of safety properties. In: Ball, T., Rajamani, S.K. (eds.) SPIN 2003. LNCS, vol. 2648, pp. 74–88. Springer, Heidelberg (2003)

25. Löding, C.: Optimal bounds for transformations of ω -automata. In: Pandu Rangan, C., Raman, V., Sarukkai, S. (eds.) FST TCS 1999. LNCS, vol. 1738, pp. 97–109. Springer, Heidelberg (1999)
26. Michel, M.: Complementation is more difficult with automata on infinite words. CNET, Paris (1988)
27. Morgenstern, A., Schneider, K.: From LTL to symbolically represented deterministic automata. In: Logozzo, F., Peled, D.A., Zuck, L.D. (eds.) VMCAI 2008. LNCS, vol. 4905, pp. 279–293. Springer, Heidelberg (2008)
28. Piterman, N.: From nondeterministic Büchi and Streett automata to deterministic parity automata. *Logical Methods in Computer Science* 3(3:5), 1–21 (2007)
29. Piterman, N., Pnueli, A.: Faster solutions of Rabin and Streett games. In: LICS 2006, pp. 275–284. IEEE (2006)
30. Puterman, M.: *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York (1994)
31. Safra, S.: On the complexity of ω -automata. In: FOCS, pp. 319–327. IEEE (1988)
32. Schewe, S.: Tighter bounds for the determinisation of Büchi automata. In: de Alfaro, L. (ed.) FOSSACS 2009. LNCS, vol. 5504, pp. 167–181. Springer, Heidelberg (2009)
33. Somenzi, F., Bloem, R.: Efficient Büchi automata from LTL formulae. In: Emerson, E.A., Sistla, A.P. (eds.) CAV 2000. LNCS, vol. 1855, pp. 248–263. Springer, Heidelberg (2000)
34. Vardi, M., Wolper, P.: An automata-theoretic approach to automatic program verification. In: LICS 1986, pp. 332–344. IEEE Computer Society (1986)