

Good-for-MDP State Reduction for Stochastic LTL Planning

Christoph Weinhuber¹, Giuseppe De Giacomo¹, Yong Li^{2*}, Sven Schewe³, Qiyi Tang³

¹University of Oxford, Oxford, UK

²Key Laboratory of System Software (Chinese Academy of Sciences),
Institute of Software Chinese Academy of Sciences, PRC

³University of Liverpool, UK

{christoph.weinhuber, giuseppe.degiacom}@cs.ox.ac.uk, liyong@ios.ac.cn, {sven.schewe, qiyi.tang}@liverpool.ac.uk

Abstract

We study stochastic planning problems in Markov Decision Processes (MDPs) with goals specified in Linear Temporal Logic (LTL). The state-of-the-art approach transforms LTL formulas into good-for-MDP (GFM) automata, which feature a restricted form of nondeterminism. These automata are then composed with the MDP, allowing the agent to resolve the nondeterminism during policy synthesis. A major factor affecting the scalability of this approach is the size of the generated automata. In this paper, we propose a novel GFM state-space reduction technique that significantly reduces the number of automata states. Our method employs a sophisticated chain of transformations, leveraging recent advances in good-for-games minimisation developed for adversarial settings. In addition to our theoretical contributions, we present empirical results demonstrating the practical effectiveness of our state-reduction technique. Furthermore, we introduce a direct construction method for formulas of the form $\text{GF}\varphi$, where φ is a co-safety formula. This construction is provably single-exponential in the worst case, in contrast to the general doubly-exponential complexity. Our experiments confirm the scalability advantages of this specialised construction.

Code — <https://zenodo.org/records/17582751>

Extended version — <https://arxiv.org/pdf/2511.09073>

1 Introduction

Planning with temporal objectives has a long tradition in artificial intelligence. Early pioneering work, such as (Bacchus, Boutilier, and Grove 1996, 1997; Littman 1997; Littman, Goldsmith, and Mundhenk 1998; Thiébaux et al. 2006), defined key formalisms like *Markov decision processes* (MDPs), analysed the computational complexity of planning problems and introduced methods for handling *non-Markovian rewards*.

These foundations paved the way for more expressive and structured approaches, most notably the use of *Linear Temporal Logic* (LTL) (Pnueli 1977), to formally specify complex goals (Lacerda, Parker, and Hawes 2014; Brafman and Giacomo 2024). This field is now even pushing into areas like reinforcement learning (Yang, Littman, and Carbin

2022), multi-agent systems (Schillinger, Bürger, and Dimarogonas 2019), and handling uncertainty (Yu et al. 2025).

Traditionally, in order to solve an MDP with an LTL goal φ , we first build a *nondeterministic Büchi automaton* (NBA) for φ . Then we convert the NBA to an equivalent *deterministic automaton*, e.g., a deterministic Rabin automaton (DRA) (Safra 1988). Recent progresses in this direction include direct translations of LTL to DRAs (Esparza and Kretínský 2014). Afterwards, we need to perform the Cartesian product of the MDP and the DRA, identify *accepting end-components* (ECs) and compute a (*memoryless*) *strategy* with maximal probability of reaching accepting ECs.

State-of-the-art approaches (Hahn et al. 2015; Sickert et al. 2016) proposed the usage of *limit-deterministic Büchi automata* (LDBAs) to avoid the notorious problem of obtaining DRAs. LDBAs give the burden of resolving the *nondeterminism* in automata to the agent, without changing the optimal satisfaction probability. Yet, only a certain type of LDBAs works for all finite MDPs. The precise criterion for NBAs is expressed by being “*good-for-MDP*” (GFM), where all nondeterministic choices are *angelic* in the sense that they can be resolved by the agent, while preserving the optimal satisfaction probability (Hahn et al. 2020). While the LTL to GFM construction does not change the complexity of the overall planning problem, in practice, the difference among using GFMs and DRAs is significant (Meyer, Sickert, and Luttenberger 2018).

A major factor affecting the scalability of solving MDPs with LTL goals is the size of the GFM automata corresponding to the LTL goals (Hahn et al. 2015; Sickert and Kretínský 2016). In this paper, we propose a novel GFM state-space reduction technique that significantly reduces the number of automata states. Our method employs a sophisticated chain of transformations, which allows us to leverage recent advances in *good-for-games* minimisation developed for adversarial settings (Radi and Kupferman 2022) to our context. We provide experimental evidence that our state-reduction technique effectively reduces the automata state space, by taking benchmark examples of LTL goals from eight sources, including influential papers in planning, verification and reinforcement learning.

Furthermore, we introduce a *direct* method for constructing the GFM automata of formulas of the form $\text{GF}\varphi$, where φ is a *co-safety* property. This kind of formulas is common

*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

in both verification (Holeček et al. 2004) and reinforcement learning (Jackermeier and Abate 2025). We show that the GFM automaton obtained through our specialised construction incurs only a *singly* exponential blow-up, in contrast to the doubly exponential blow-up in the general case. Moreover, we provide experimental evidence that this theoretical advantage indeed translates into significantly fewer states.

Extended Version. We refer to the extended version for all missing proofs and the full benchmark set.

2 Preliminaries

Markov Decision Processes. Following (Baier and Katoen 2008), a Markov decision process (MDP) \mathcal{M} is a tuple $(S, \text{Act}, \Sigma, P, s_0, L)$ with a finite set of states S , a set of actions Act , a set of labels Σ , a transition probability function $P : S \times \text{Act} \times S \rightarrow [0, 1]$, an initial state $s_0 \in S$ and a labelling function $L : S \times \text{Act} \rightarrow \Sigma$ that labels state-action pairs to the set of propositions that hold in that state¹. We abuse the notation by writing $\text{Act}(s)$ to denote the set of available actions at state s . A path ξ of \mathcal{M} is an (in)finite sequence of alternating states and actions $\xi = s_0 a_0 s_1 a_1 \dots$, ending with a state if finite, such that for all $i \geq 0$, $a_i \in \text{Act}(s_i)$ and $P(s_i, a_i, s_{i+1}) > 0$. The sequence $L(\xi) = L(s_0, a_0)L(s_1, a_1), \dots$ over Σ is called the *trace* induced by the path ξ over \mathcal{M} . We denote by FPaths and IPaths the set of all finite and infinite paths of \mathcal{M} .

A (finite-memory) strategy μ of \mathcal{M} is a function $\mu : \text{FPaths} \rightarrow \text{Distr}(\text{Act})$ such that, for each $\xi \in \text{FPaths}$, $\mu(\xi) \in \text{Distr}(\text{Act}(\text{lst}(\xi)))$, where $\text{lst}(\xi)$ is the last state of the finite path ξ and $\text{Distr}(\text{Act})$ denotes the set of all possible distributions over Act . Let $\Omega_\mu^{\mathcal{M}}(s_0)$ denote the subset of infinite paths of \mathcal{M} that correspond to strategy μ and initial state s_0 . A strategy μ is *memoryless* if $\mu : S \rightarrow \text{Distr}(\text{Act})$, meaning it selects actions independently of the past.

A strategy μ of \mathcal{M} is able to resolve the nondeterminism of an MDP and induces a Markov chain (MC) $\mathcal{M}^\mu = (\text{FPaths}, \Sigma, P_\mu, L')$ where, for $\xi = s_0 a_0 \dots s_{n-1} a_{n-1} s_n \in \text{FPaths}$ and $a_n \in \text{Act}(s_n)$, $P_\mu(\xi, \xi \cdot a_n \cdot s_{n+1}) = P(s_n, a_n, s_{n+1}) \cdot \mu(\xi)(a_n)$ and $L'(\xi) = L(s_{n-1}, a_{n-1})$.

A *sub-MDP* of \mathcal{M} is an MDP $\mathcal{M}' = (S', \text{Act}', \Sigma, P', L')$ where $S' \subseteq S$, $\text{Act}' \subseteq \text{Act}$ is such that, for every $s \in S'$, $\text{Act}'(s) \subseteq \text{Act}(s)$, and P' and L' are analogous to P and L when restricted to S' and Act' . In particular, \mathcal{M}' is closed under probabilistic transitions, i.e., for all $s \in S'$ and $a \in \text{Act}'$ we have that $P(s, a, s') > 0$ implies $P'(s, a, s') > 0$. An *end-component* (EC) of an MDP \mathcal{M} is a sub-MDP \mathcal{M}' of \mathcal{M} such that its underlying graph is strongly connected. A maximal end-component (MEC) is an EC $\mathcal{E} = (E', \text{Act}', \Sigma, P', L')$ such that there is no other EC $\mathcal{E}'' = (E'', \text{Act}'', \Sigma, P'', L'')$ such that $E' \subset E''$. An MEC \mathcal{E} that cannot reach states outside E' is a *leaf* component.

Theorem 1 ((de Alfaro 1997; Baier and Katoen 2008)). *Once an EC \mathcal{E} of an MDP is entered, there is a strategy that visits every state-action combination in \mathcal{E} with probability 1 and stays in \mathcal{E} forever. Moreover, for every strategy the union*

¹This generalises the usual labelling function $L : S \rightarrow \Sigma$.

of the end-components is visited with probability 1. An infinite path of an MC \mathcal{M} almost surely (with probability 1) will enter a leaf component.

Linear Temporal Logic. An LTL formula, over a finite set of atomic propositions AP is defined as $\varphi ::= a \in \text{AP} \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \text{X}\varphi \mid \text{G}\varphi \mid \text{F}\varphi \mid \varphi \text{U}\varphi$. Here X (Next), G (Globally/Always), F (Finally/Eventually) and U (Until) are temporal operators.

An ω -trace w is an infinite sequence of letters $w[0]w[1]w[2]\dots$ with $w[i] \in \Sigma = 2^{\text{AP}}$. We denote the infinite suffix $w[i]w[i+1]\dots$ by w_i . The satisfaction relation \models between ω -traces w and formulas φ is defined as follows:

$$\begin{aligned} w \models a &\iff a \in w[0], \\ w \models \neg\varphi &\iff w \not\models \varphi, \\ w \models \text{X}\varphi &\iff w_1 \models \varphi, \\ w \models \text{G}\varphi &\iff \forall k. w_k \models \varphi, \\ w \models \text{F}\varphi &\iff \exists k. w_k \models \varphi, \\ w \models \varphi \text{U}\psi &\iff \exists k. (w_k \models \psi \wedge \forall j < k, w_j \models \varphi) \end{aligned}$$

Further, $\forall w.w \models tt$ and $\forall w.w \not\models ff$. We denote by $[\varphi]$ the set of *infinite traces* satisfying the LTL formula φ .

Automata. A (nondeterministic) transition system (TS) is a tuple $\mathcal{T} = (Q, q_0, \delta)$, with a finite set of states Q , an initial state $q_0 \in Q$, and a transition function $\delta : Q \times \Sigma \rightarrow 2^Q$. We extend δ to sets by $\delta(S, \sigma) := \bigcup_{q \in S} \delta(q, \sigma)$. A *deterministic* TS satisfies $|\delta(q, \sigma)| \leq 1$ for each $q \in Q$ and $\sigma \in \Sigma$.

An automaton \mathcal{A} is defined as a tuple (\mathcal{T}, α) , where \mathcal{T} is a TS and α is an acceptance condition. We define $\Delta(\delta) = \{(q, \sigma, q') \mid q, q' \in Q, \sigma \in \Sigma, q' \in \delta(q, \sigma)\}$. For an infinite trace $w \in \Sigma^\omega$, a *run* of \mathcal{A} on w is an infinite sequence of transitions $\rho = (q_0, w[0], q_1)(q_1, w[1], q_2) \dots$ such that, for every $i \geq 0$, $q_{i+1} \in \delta(q_i, w[i])$. Let $\text{inf}(\rho)$ be the set of transitions that occur infinitely often in a run ρ . The acceptance condition $\alpha \subseteq \Delta(\delta)$ is a set of *accepting* (*rejecting*, resp.) transitions for Büchi (co-Büchi, resp.). A run ρ satisfies the Büchi (co-Büchi, resp.) acceptance condition α if $\text{inf}(\rho) \cap \alpha \neq \emptyset$ ($\text{inf}(\rho) \cap \alpha = \emptyset$, resp.).

A run is *accepting* if it satisfies the condition α ; a trace $w \in \Sigma^\omega$ is *accepted* by \mathcal{A} if there is an accepting run ρ of \mathcal{A} over w . We use three letter acronyms in $\{D, N\} \times \{B, C\} \times \{A\}$ to denote automata types where the first letter stands for the TS mode, the second for the acceptance type and the third for automaton. For example, DBA means deterministic Büchi automaton. We assume that all automata are *complete*, i.e., for each state $s \in Q$ and letter $\sigma \in \Sigma$, $|\delta(s, \sigma)| \geq 1$.

We denote by $\mathcal{L}(\mathcal{A})$ the ω -language recognised by an ω -automaton \mathcal{A} , i.e., the set of ω -traces accepted by \mathcal{A} .

3 Stochastic Planning Problem

Overview. In a stochastic planning problem, we model the interaction between the agent and the environment as an MDP \mathcal{M} and its task specification as an LTL formula φ . The goal is to find an optimal strategy μ for \mathcal{M} that maximises the chance that the infinite sequence of observed propositions satisfies φ . Formally, we want the probability of the ω -traces generated by \mathcal{M}^μ and belonging to $[\varphi]$, to be maximal among all possible choices of μ .

MDPs and LTL Task Specification. Every LTL formula φ can be translated into a (nondeterministic) Büchi automaton \mathcal{A} , accepting $[\varphi]$, i.e., $\mathcal{L}(\mathcal{A}) = [\varphi]$ (Vardi and Wolper 1986). Recall that, for a given strategy μ on an MDP \mathcal{M} , we can induce an MC \mathcal{M}^μ . We define the semantic satisfaction probability of the induced MC \mathcal{M}^μ for $\mathcal{L}(\mathcal{A})$ as $\mathbb{P}_{\mathcal{M}^\mu}(\mathcal{L}(\mathcal{A})) = \mathbb{P}\{\xi \in \Omega_\mu^{\mathcal{M}}(s_0) : L(\xi) \in \mathcal{L}(\mathcal{A})\}$, which is the probability of \mathcal{M}^μ generating an ω -trace in $[\varphi]$. For an MDP \mathcal{M} and an automaton \mathcal{A} , we define the *maximal semantic satisfaction probability* as $\text{Psem}(\mathcal{M}, \mathcal{A}) = \sup_\mu \mathbb{P}_{\mathcal{M}^\mu}(\mathcal{L}(\mathcal{A}))$. We look for an *optimal* strategy μ that achieves $\text{Psem}(\mathcal{M}, \mathcal{A})$ in the planning problem.

Product of MDPs and Büchi automata. If \mathcal{A} is a DBA, we can solve the problem of finding an optimal strategy μ by constructing the product MDP $\mathcal{M} \times \mathcal{A}$ and extract a memoryless strategy μ on $\mathcal{M} \times \mathcal{A}$ that reaches accepting components with maximal probability (Baier and Katoen 2008). If \mathcal{A} is nondeterministic, this is in general not possible, because we have to resolve the nondeterminism, which corresponds to future forecasting capabilities that we can assign neither to the agent, nor to the (probabilistic) environment. However, it has been observed (Hahn et al. 2020; Sickert et al. 2016; Hahn et al. 2015) that there is a class of automata called good-for-MDPs (GFM) for which we can assign the nondeterminism to the agent without loss of optimality. Intuitively, a nondeterministic automaton is GFM, if the nondeterminism can be resolved by the agent (we postpone the formal definition of GFM to later in this section). If an automaton is GFM, then we can still adopt a variant of the product construction shown below:

Formally, let $\mathcal{M} = (S, \text{Act}, \Sigma, P, s_0, L)$ be an MDP and $\mathcal{A} = (Q, q_0, \delta, \alpha)$ be an NBA. We define the product MDP $\mathcal{M} \times \mathcal{A} = (S^\times, \text{Act}^\times, \Sigma, P^\times, \langle s_0, q_0 \rangle, L^\times, \alpha^\times)$ augmented with the acceptance condition α^\times where

- $S^\times = S \times Q$ is the state space.
- $\text{Act}^\times = \text{Act} \times [k]$ is the action set where k is the *maximal* out-degree of \mathcal{A} for any state in Q and letter in Σ . We denote $[k] = \{1, \dots, k\}$ for any integer $k > 0$.
- $P^\times : S^\times \times \text{Act}^\times \times S^\times \rightarrow [0, 1]$ is the transition probability function such that $P^\times(\langle s, q \rangle, \langle a, i \rangle, \langle s', q' \rangle) = P(s, a, s')$ if $P(s, a, s') > 0$ and $q' \in \delta(q, L(s, a))$ where q' is the i -th state in the *ordered* set $\delta(q, L(s, a))$.
- $L^\times(\langle s, q \rangle, \langle a, i \rangle) = L(s, a)$ for state $\langle s, q \rangle \in S \times Q$ and action $\langle a, i \rangle \in \text{Act}^\times$, and
- For Büchi/co-Büchi, $\alpha^\times = \{(\langle s, q \rangle, \langle a, i \rangle, \langle s', q' \rangle) \mid P^\times(\langle s, q \rangle, \langle a, i \rangle, \langle s', q' \rangle) > 0, (q, L(s, a), q') \in \alpha\}$.

Instead of storing the selected successor state in the action name as in other literature, e.g. (Hahn et al. 2020), we index the choice by its position in a predefined order over successors. Formally, for each $q \in Q$ and $\sigma \in \Sigma$, we define a function $\text{id}_{\mathcal{A}, q, \sigma}$ mapping each $q' \in \delta(q, \sigma)$ to its unique position $i \in [k]$. This index i serves as the identifier for the agent's choice in our product construction. This encoding is equivalent to (Hahn et al. 2020) but has a slight advantage of using fewer actions. For full details, see the extended version.

Conditions for GFMness. In order to see if an NBA is GFM, we have to characterise the probabilities that come from the syntactic construction above and show that they are the same of those intrinsic in the original problem. Recall that the probability of the original problem is the following:

$$\text{Psem}(\mathcal{M}, \mathcal{A}) = \sup_\mu \mathbb{P}_{\mathcal{M}^\mu}(\mathcal{L}(\mathcal{A})).$$

The probability coming from the syntactic product construction above is characterised as:

$$\text{Psyn}(\mathcal{M}, \mathcal{A}) = \sup_\mu \mathbb{P}\{\xi \in \Omega_\mu^{\mathcal{M} \times \mathcal{A}}(\langle s_0, q_0 \rangle) : \xi \text{ accepting}\}.$$

This can be simplified to $\text{Psyn}(\mathcal{M}, \mathcal{A}) = \sup_\mu \mathbb{P}_{(\mathcal{M} \times \mathcal{A})^\mu}(\text{FX})$, where X is the set of states of the accepting MECs in $\mathcal{M} \times \mathcal{A}$ that contain accepting transitions in α^\times .

Clearly, $\text{Psyn}(\mathcal{M}, \mathcal{A}) \leq \text{Psem}(\mathcal{M}, \mathcal{A})$, because accepting runs ξ only occur on accepting words. Thus, a strategy μ chooses an accepting run on accepting words in the best case, but it is also possible for μ to make wrong decisions.

Obviously, $\text{Psyn}(\mathcal{M}, \mathcal{A}) = \text{Psem}(\mathcal{M}, \mathcal{A})$ if \mathcal{A} is deterministic. Following (Hahn et al. 2020), an NBA \mathcal{A} is said to be GFM if, for all finite MDPs \mathcal{M} , $\text{Psem}(\mathcal{M}, \mathcal{A}) = \text{Psyn}(\mathcal{M}, \mathcal{A})$. See extended version for a non-GFM case.

LTL to GFM automata. There are several algorithms to transform an LTL formula into an NBA that is GFM (Sickert et al. 2016; Hahn et al. 2020). The runtime to construct a GFM automaton from an LTL formula is doubly exponential in the size of the formula, which is the same cost as obtaining a deterministic automata. However, the practical advantage of going through GFM instead of deterministic automata is enormous and all state-of-the-art systems implement an LTL to GFM construction (Sickert et al. 2016; Hahn et al. 2020).

4 Probabilistic ω -Automata

Our state reduction technique is based on obtaining a small probabilistic (Büchi) automaton (PA) out of a GFM automaton. The key characteristic of PAs (Baier, Größer, and Bertrand 2012) is that, while being nondeterministic, their nondeterministic choices are resolved randomly. A PA $\mathcal{P} = (\Sigma, Q, \delta, q_0, \alpha)$ is a nondeterministic automaton equipped with a Büchi acceptance condition α and a randomised transition function $\delta : Q \times \Sigma \mapsto \text{Distr}(Q)$, where, from state q and letter σ , transition to q' is taken with probability $\delta(q, \sigma)(q')$. We often abuse notation by writing $\delta(q, \sigma)$ to denote its support. Each word $w \in \Sigma^\omega$ induces a probability measure $\mathbb{P}_\mathcal{P}^w$ on Q^ω in the usual way. The probability that \mathcal{P} accepts w , denoted by $\mathbb{P}_\mathcal{P}(w)$, is the probability measure of all accepting runs of w on \mathcal{P} , that is:

$$\mathbb{P}_\mathcal{P}(w) = \mathbb{P}_\mathcal{P}^w(\{\pi : \pi \text{ is an accepting run of } w\}).$$

A PA \mathcal{P} is a 0/1-PA if, for any word $w \in \Sigma^\omega$, we have either $\mathbb{P}_\mathcal{P}(w) = 1$ or $\mathbb{P}_\mathcal{P}(w) = 0$. For a 0/1-PA \mathcal{P} , with a slight abuse of notation, we say a word w is accepted by \mathcal{P} (w is in $\mathcal{L}(\mathcal{P})$) if $\mathbb{P}_\mathcal{P}(w) = 1$. 0/1-PAs are known to be *semantically deterministic* (Li et al. 2025). That is, given a state p and a letter σ of a 0/1-PA \mathcal{P} , for any two states $q, r \in \delta(p, \sigma)$, we have that $\mathcal{L}(\mathcal{P}^q) = \mathcal{L}(\mathcal{P}^r)$ where \mathcal{P}^s is the automaton by setting the initial state to s . For a detailed introduction to PAs please refer to (Baier, Größer, and Bertrand 2012).

Product of MDPs and 0/1-PA. In the following, we show for the *first* time that 0/1-PAs can also be perceived as another type of GFM automata. To this end, we first define the product MDP of an MDP \mathcal{M} and a 0/1-PA \mathcal{P} ².

Formally we define the product as follows. Let $\mathcal{M} = (S, \text{Act}, \Sigma, \text{P}, s_0, L)$ be an MDP and $\mathcal{P} = (Q, q_0, \delta, \alpha)$ be a 0/1-PA over Σ . We define the product MDP $\mathcal{M} \otimes \mathcal{P} = (S^\otimes, \text{Act}^\otimes, \Sigma, \text{P}^\otimes, \langle s_0, q_0 \rangle, L^\otimes, \alpha^\otimes)$ where

- $S^\otimes = S \times Q$ is the state space,
- $\text{Act}^\otimes = \text{Act}$ is the action set,
- $\text{P}^\otimes : S^\otimes \times \text{Act}^\otimes \times S^\otimes \rightarrow [0, 1]$ is the transition probability function such that $\text{P}^\otimes(\langle s, q \rangle, a, \langle s', q' \rangle) = \text{P}(s, a, s') \cdot \delta(q, \sigma, q')$ if $\text{P}(s, a, s') > 0$ and $q' \in \delta(q, \sigma)$ where $\sigma = L(s, a)$,
- $L^\otimes : S^\otimes \times \text{Act}^\otimes \rightarrow \Sigma$ where $L^\otimes(\langle s, q \rangle, a) = L(s, a)$ and
- $\alpha^\otimes = \{(\langle s, q \rangle, a, \langle s', q' \rangle) \mid (q, L(s, a), q') \in \alpha, a \in \text{Act}^\otimes, \text{P}^\otimes(\langle s, q \rangle, a, \langle s', q' \rangle) > 0\}$.

Intuitively, we still ask the agent to make decisions on which letter to choose for the 0/1-PA \mathcal{P} , but leave the choice of a successor to a random strategy. That is, the agent resolves the nondeterminism by choosing actions, and once the action a is selected, according to the definition of MDP $\mathcal{M} \otimes \mathcal{P}$, we also know the chosen letter $\sigma = L^\otimes(\langle s, q \rangle, a)$ as well. This is different from the classical product $\mathcal{M} \times \mathcal{A}$, where the agent not only decides the next letter, but also the next successor in \mathcal{A} , by selecting the action $\langle a, i \rangle \in \text{Act}^\times$.

For the product $\mathcal{M} \otimes \mathcal{P}$, we can define a similar maximal *syntactic* satisfaction probability:

$$\text{Psyn}(\mathcal{M}, \mathcal{P}) = \sup_{\mu} \mathbb{P}\{\xi \in \Omega_{\mu}^{\mathcal{M} \otimes \mathcal{P}}(\langle s_0, q_0 \rangle) : \xi \text{ accepting}\}.$$

Crucially, we can see that 0/1-PAs are another type of GFM Büchi automata in the following sense:

Theorem 2. *Let \mathcal{D} be a DBA and \mathcal{P} be an equivalent 0/1-PA such that $\mathcal{L}(\mathcal{P}) = \mathcal{L}(\mathcal{D})$. For a finite MDP \mathcal{M} , we have that $\text{Psyn}(\mathcal{M}, \mathcal{P}) = \text{Psyn}(\mathcal{M}, \mathcal{D}) = \text{Psem}(\mathcal{M}, \mathcal{D}) = \text{Psem}(\mathcal{M}, \mathcal{P})$.*

An equivalent 0/1-PA \mathcal{P} always exists for a DBA \mathcal{D} because a DBA can be easily transformed to an equivalent 0/1-PA by transitioning to the only successor with probability one. Since 0/1-PAs are GFM based on our product MDP definition, we can then apply standard strategy synthesis approach once the product MDP $\mathcal{M} \otimes \mathcal{P}$ is constructed.

5 Good-for-MDP State Reduction

In this section, we present our main contribution, a general state-space reduction for GFM automata. We assume to have a GFM automaton \mathcal{A}_{GFM} and through a polynomial *Redux* procedure, we return a 0/1-PA \mathcal{A}_{PA} , whose state space can be significantly reduced. *Redux* works in several stages, see Figure 1, which we detail below.

²We note that a source of 0/1-PAs can be a special type of automata whose nondeterminism can be resolved by randomness discussed in (Henzinger, Prakash, and Thejaswini 2025). This type of automata is also said to be GFM in a discussion without formal proofs in (Henzinger, Prakash, and Thejaswini 2025).




Figure 1: Overview of our state reduction pipeline

- First step (Poly): we perform a GFM-to-DBA transformation by embedding in the transitions of the GFM \mathcal{A}_{GFM} the choices $[k]$ used in the product construction of Section 3, making it a complete DBA \mathcal{A}_{DBA} .
- Second step (Const): we treat the DBA \mathcal{A}_{DBA} as a DCA, \mathcal{A}_{DCA} , by simply changing the acceptance condition.
- Third step (Poly): we apply a minimisation algorithm originally developed for good-for-games (GFG) NCAs (Radi and Kupferman 2022) to \mathcal{A}_{DCA} to obtain a minimal GFG-NCA $\mathcal{A}_{\text{GFG-Min}}$.
- Fourth step (Poly): we transform the GFG-NCA $\mathcal{A}_{\text{GFG-Min}}$ into a 0/1-PA \mathcal{A}_{PA} with the same number of states as $\mathcal{A}_{\text{GFG-Min}}$.

Here “Poly” and “Const” indicate that their corresponding steps perform in polynomial and constant time, respectively.

Step 1: GFM to DBA

We consider the product $\mathcal{M} \times \mathcal{A}_{\text{GFM}}$ of Section 3 and modify both the MDP \mathcal{M} and automaton \mathcal{A}_{GFM} by embedding the transition choices $[k]$. The resulting MDP \mathcal{M}' has its action set expanded from Act to $\text{Act} \times [k]$. Similarly, the resulting automaton \mathcal{A}_{DBA} will have the transitions re-labelled from Σ to $\Sigma \times [k]$. As a result the automaton \mathcal{A}_{DBA} is deterministic.

Formally, we define the modified MDP $\mathcal{M}' = (S', s'_0, \text{Act}', \text{P}', L')$ from $\mathcal{M} = (S, s_0, \text{Act}, \text{P}, L)$ as

- $S' = S, s'_0 = s_0, \text{Act}' = \text{Act} \times [k]$,
- $L' : S' \times \text{Act}' \rightarrow \Sigma'$ where $\Sigma' = \Sigma \times [k]$, and if $L(s, a) = \sigma$, then $L'(s, \langle a, i \rangle) = \langle \sigma, i \rangle$ for all $i \in [k]$, and
- $\text{P}'(s, \langle a, i \rangle, s') = \text{P}(s, a, s')$ for all $i \in [k]$ and $\text{P}(s, a, s') > 0$.

We define the DBA $\mathcal{A}_{\text{DBA}} = (Q_{\text{DBA}}, q_0, \delta_{\text{DBA}}, \alpha_{\text{DBA}})$ from $\mathcal{A}_{\text{GFM}} = (Q, q_0, \delta, \alpha)$ as:

- $Q_{\text{DBA}} = Q, \alpha_{\text{DBA}} = \{(q, \langle \sigma, i \rangle, q') \mid (q, \sigma, q') \in \alpha, i = \text{id}_{\mathcal{X}_{\mathcal{A}, q, \sigma}}(q')\}$, and
- $\delta_{\text{DBA}}(q, \langle \sigma, i \rangle) = q'$ if $q' \in \delta(q, \sigma)$ and $\text{id}_{\mathcal{X}_{\mathcal{A}, q, \sigma}}(q') = i$.

\mathcal{A}_{DBA} will then be made complete. Intuitively, we assign the nondeterminism of \mathcal{A}_{GFM} to the agent by extending the action set Act to $\text{Act} \times [k]$ in \mathcal{M}' , which then would allow us to obtain the DBA \mathcal{A}_{DBA} since the nondeterminism in \mathcal{A}_{GFM} will be taken care of by the agent. We observe that $\mathcal{M} \times \mathcal{A}_{\text{GFM}}$ and $\mathcal{M}' \times \mathcal{A}_{\text{DBA}}$ (obtained by applying the product construction in Section 3) have the same state space. Moreover, in $\mathcal{M}' \times \mathcal{A}_{\text{DBA}}$, the action set is $\text{Act} \times [k] \times [1]$, as the maximal out-degree of the DBA \mathcal{A}_{DBA} is 1. We can drop $[1]$, which is a constant in every transition and obtain:

Lemma 1. *The product $\mathcal{M} \times \mathcal{A}_{GFM}$ and $\mathcal{M}' \times \mathcal{A}_{DBA}$ are the same. Moreover, the optimal strategy μ for each state pair $(s, q) \in S \times Q$ in either of the products will obtain optimal satisfaction probabilities in both products.*

By Lemma 1, it is easy to observe that the syntactic probabilities of $\mathcal{M} \times \mathcal{A}_{GFM}$ and $\mathcal{M}' \times \mathcal{A}_{DBA}$ are the same. Since \mathcal{A}_{DBA} is deterministic and thus GFM, we have $\text{Psyn}(\mathcal{M}', \mathcal{A}_{DBA}) = \text{Psem}(\mathcal{M}', \mathcal{A}_{DBA})$. It then follows:

Theorem 3. $\text{Psem}(\mathcal{M}, \mathcal{A}_{GFM}) = \text{Psyn}(\mathcal{M}, \mathcal{A}_{GFM}) = \text{Psyn}(\mathcal{M}', \mathcal{A}_{DBA}) = \text{Psem}(\mathcal{M}', \mathcal{A}_{DBA})$.

The GFM to DBA conversion can be done in polynomial time. The construction of the modified MDP \mathcal{M}' and the DBA \mathcal{A}_{DBA} involves a direct translation of the original components. The number of states remains the same, while the size of the action set and the alphabet expands by a factor of k , which is a polynomial increase.

Step 2: DBA to DCA

Next, we syntactically convert the DBA \mathcal{A}_{DBA} into a DCA \mathcal{A}_{DCA} . To do so, we only need to modify the acceptance condition. \mathcal{A}_{DCA} has the same set of states, initial state, and transition function as \mathcal{A}_{DBA} . The co-Büchi acceptance condition requires that a run is accepting if it intersects with the set of rejecting transitions only a finite number of times. The set α_{DCA} of rejecting transitions for the DCA \mathcal{A}_{DCA} is the same as the set α_{DBA} of accepting transitions in \mathcal{A}_{DBA} , but only interpreted differently. Observe that the languages of \mathcal{A}_{DBA} and \mathcal{A}_{DCA} complement each other. \mathcal{A}_{DCA} is just an intermediate automaton on which we apply GFG-minimisation next.

Since this conversion does not require any computational modification to the automaton's transition structure, it can be done in constant runtime.

Step 3: GFG-Minimisation

\mathcal{A}_{DCA} is deterministic and thus a GFG automaton that we can apply minimisation algorithms (Radi and Kupferman 2022) on in this step. To formally understand this step, we take a detour to introduce GFG automata.

GFG automata (Henzinger and Piterman 2006) are automata in which nondeterminism can be resolved *deterministically* by assigning the choice to the protagonist in an adversarial (vs. stochastic MDPs) games. Formally, an automaton \mathcal{A} is said to be GFG if there exists a strategy $f : \Sigma^* \rightarrow Q$ such that for every accepting word $w = \sigma_0\sigma_1 \dots$, the run $f(\epsilon)f(\sigma_0) \dots f(\sigma_0 \dots \sigma_i) \dots$ is accepting. Intuitively, this indicates that there is a deterministic strategy to produce an accepting run in a GFG automaton even if the accepting word is given letter by letter. Interestingly, (transition-based) GFG co-Büchi automata can be minimised in *polynomial* time (Radi and Kupferman 2022). Since GFG automata have more restricted nondeterminism than GFM automata, GFG automata are also GFM (Klein et al. 2014). GFG automata are not adopted in state-of-the-art stochastic planning tools, such as PRISM (Kwiatkowska, Norman, and Parker 2011). The reason is that GFG Büchi (resp. co-Büchi) automata cannot recognise all languages expressed by LTL as they are only as expressive as their deterministic counterparts. Moreover, current approaches to construct GFG automata are not

as efficient as the ones for GFM. Notably, our work is able to use GFG minimisation on GFM automata that recognise all ω -regular properties. See extended version for more details.

In our construction we apply GFG co-Büchi minimisation (Radi and Kupferman 2022) on \mathcal{A}_{DCA} , yielding a minimal GFG-NCA $\mathcal{A}_{GFG-\text{Min}}$ accepting $\mathcal{L}(\mathcal{A}_{DCA})$.

Step 4: GFG-Min to 0/1-PA

(Li et al. 2025) has proven that the minimal GFG-NCA produced by (Radi and Kupferman 2022) can be turned into a language equivalent 0/1-PA by resolving its nondeterminism with random choices.

The translation from the minimal GFG-NCA $\mathcal{A}_{GFG-\text{Min}}$ to a 0/1-PA \mathcal{A}_{PA} is quite simple: we only need to resolve the nondeterminism by random choices. Formally, we treat $\mathcal{A}_{GFG-\text{Min}}$ as a Büchi automaton $\mathcal{A}_{NBA} = (Q_{NBA}, q_0, \delta_{NBA}, \alpha_{NBA})$ and build the PA $\mathcal{A}_{PA} = (Q_{PA}, q_0, \delta_{PA}, \alpha_{PA})$ as:

- $Q_{PA} = Q_{NBA}$, $\alpha_{PA} = \alpha_{NBA}$, and
- $\delta_{PA}(q, \sigma)(q') = \frac{1}{|\delta_{NBA}(q, \sigma)|}$ for each $q \in Q_{PA}$, $\sigma \in \Sigma'$ and $q' \in \delta_{NBA}(q, \sigma)$.

Following from (Li et al. 2025, Lemma 3), we get:

Lemma 2. $\mathcal{L}(\mathcal{A}_{PA}) = \mathcal{L}(\mathcal{A}_{DBA})$ and \mathcal{A}_{PA} is a 0/1-PA.

It immediately follows that, according to Theorem 2, we can use \mathcal{A}_{PA} to obtain an optimal strategy for \mathcal{M}' to achieve $\text{Psem}(\mathcal{M}', \mathcal{A}_{PA}) = \text{Psem}(\mathcal{M}', \mathcal{A}_{DBA})$ since $\mathcal{L}(\mathcal{A}_{PA}) = \mathcal{L}(\mathcal{A}_{DBA})$ holds. Together with Theorems 2 and 3, we then obtain our main result:

Theorem 4. *From an optimal strategy μ on $\mathcal{M}' \otimes \mathcal{A}_{PA}$, one can obtain an optimal strategy μ' for \mathcal{M} that achieves the maximal probability $\text{Psem}(\mathcal{M}, \mathcal{A}_{GFM})$.*

Summary. Details on optimal strategy synthesis, along with further optimisations, are in the extended version. We summarise our stochastic planning algorithm below:

- Construct a GFM (Büchi) automaton \mathcal{A}_{GFM} from φ .
- Create \mathcal{M}' from \mathcal{M} and \mathcal{A}_{DBA} (and \mathcal{A}_{DCA}) from \mathcal{A}_{GFM} .
- Apply GFG minimisation on \mathcal{A}_{DCA} and treat the minimised automaton $\mathcal{A}_{GFG-\text{Min}}$ as a Büchi automaton \mathcal{A}_{NBA} .
- Translate \mathcal{A}_{NBA} to a 0/1-PA \mathcal{A}_{PA} .
- Create $\mathcal{M}' \otimes \mathcal{A}_{PA}$.
- Identify the accepting MECs that have some accepting transitions.
- Synthesise an optimal strategy μ that maximises the reachability probability of accepting MECs as usual.

6 Direct GFM Construction for $\text{GF}\varphi$

$\text{GF}\varphi$ fragment. We now focus on the syntactic class of LTL formulas of the form $\text{GF}\varphi$ where φ is a *co-safety* formula with the following syntax:

$$\varphi := a \mid \neg a \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid X\varphi \mid F\varphi \mid \varphi U\varphi.$$

They express that a finite trace pattern repeats infinitely often, which we refer to as *repeated reachability* properties. Formally, $P \subseteq \Sigma^\omega$ is a repeated reachability property if

there exists a language of finite words $R \subseteq \Sigma^*$ such that, for every $w \in P$, infinitely many prefixes of w belong to $\Sigma^* \cdot R$. Then, we have that:

Lemma 3. *P is a repeated reachability property specifiable in LTL if, and only if, P is specifiable in $\text{GF}\varphi$, where φ is a co-safety formula.*

The $\text{GF}\varphi$ fragment is very common in verification (Holeček et al. 2004) and also in reinforcement learning (Jackermeier and Abate 2025). More generally, repeated reachability properties are within a commonly used subclass of so-called *recurrence properties* (Manna and Pnueli 1990; Chang, Manna, and Pnueli 1992).

Direct GFM Construction for $\text{GF}\varphi$. Consider a formula of the form $\text{GF}\varphi$, the subformula φ expresses a *finite trace property*, which can be accepted by a nondeterministic finite automaton (NFA)³. To make our construction more general, we will start from an NFA instead of the formula φ in the following. An NFA \mathcal{N} is a tuple (\mathcal{T}, F) where $\mathcal{T} = (Q, q_0, \delta)$ is a *nondeterministic TS* and $F \subseteq Q$ is a set of *final states*. Unlike NBAs, NFAs only operate on *finite traces* rather than ω -traces. A finite trace $u \in \Sigma^*$ is accepted by an NFA \mathcal{N} if one of its finite runs *terminates* at a final state.

We can construct for a co-safety property formula φ , an NFA $\mathcal{N}_\varphi = (Q, q_0, \delta, F)$ such that $\mathcal{L}_*(\mathcal{N}_\varphi) \cdot (tt)^\omega = [\varphi]$ where $\mathcal{L}_*(\mathcal{N})$ denotes the set of finite traces accepted by \mathcal{N} . Note that the alphabet here is $\Sigma = 2^{\text{AP}}$.

Now we can construct from \mathcal{N}_φ a GFM automaton $\mathcal{A} = (Q_{\mathcal{A}}, q_{\mathcal{A}}, \delta_{\mathcal{A}}, \alpha_{\mathcal{A}})$ for $\text{GF}\varphi$ where

- $Q_{\mathcal{A}} \subseteq ((Q \setminus F) \cup \{q_0\})$, $q_{\mathcal{A}} = q_0$,
- $\delta_{\mathcal{A}} : Q_{\mathcal{A}} \times \Sigma \rightarrow 2^{Q_{\mathcal{A}}}$ is defined such that for each $q \in Q_{\mathcal{A}}$ and $\sigma \in \Sigma$, we have (1) $q' \in \delta_{\mathcal{A}}(q, \sigma)$ if $q' \in \delta(q, \sigma)$ with $q' \notin F$, and (2) $q_0 \in \delta_{\mathcal{A}}(q, \sigma)$, and
- $\alpha_{\mathcal{A}} = \{(q, \sigma, q_0) \in \Delta(\delta_{\mathcal{A}}) \mid \exists q' \in \delta(q, \sigma) \wedge q' \in F\}$.

The intuition is that, for a repeated reachability property $\text{GF}\varphi$, we can forget the past finite trace at *any point* and start tracking whether the following finite trace satisfies φ . The Büchi acceptance condition will make sure that the reachability/co-safety formula φ will be satisfied infinitely often. Therefore, in the construction, we can always reset and go back to the initial state q_0 , which is why every state q has a successor q_0 on every letter $\sigma \in \Sigma$. Once the formula φ has been fulfilled, i.e., a final state q' has been reached from q over letter σ in \mathcal{N}_φ , we need to mark the transition (q, σ, q_0) in \mathcal{A} as accepting and start tracking finite traces satisfying φ again by moving back to the initial state q_0 .

It is also easy to show that a memoryless random strategy on \mathcal{A} can generate an accepting run almost surely over an ω -trace from $[\text{GF}\varphi]$. This is because accepting transitions can be reached with positive probabilities from anywhere in \mathcal{A} and it is not possible to skip accepting transitions forever in an infinite run with positive probabilities.

Let $|\varphi|$ be the number of modalities and connectives in φ . It then follows that:

³Our construction applies to any NFA, not just those from co-safety LTL (see extended version), covering repeated reachability properties not expressible in LTL, i.e., monadic first-order and finite LTL (De Giacomo and Vardi 2013).

Pattern	Owl	Owl-Red	Slim	Slim-Red
TDR[6]	64 (0.46)	64 (0.03)	65 (0.09)	34 (0.01)
TDR[7]	128 (0.46)	128 (0.15)	129 (0.12)	66 (0.02)
TDR[8]	256 (0.48)	256 (0.69)	257 (0.14)	130 (0.05)
LIB[4]	17 (0.44)	10 (0.01)	33 (0.11)	18 (0.13)
LIB[5]	21 (0.48)	12 (0.01)	65 (0.20)	34 (3.83)
LIB[6]	25 (0.62)	14 (0.01)	129 (0.69)	66 (128.34)
BRP[6]	69 (0.55)	17 (0.02)	317 (0.14)	65 (0.17)
BRP[7]	133 (0.65)	19 (0.04)	637 (0.21)	95 (0.71)
BRP[8]	261 (0.86)	21 (0.14)	1277 (0.33)	146 (3.49)
EHP	13 (0.43)	9 (0.01)	127 (0.14)	54 (0.31)
NU[4]	44 (1.16)	23 (0.02)	51 (0.22)	39 (0.02)
NU[5]	150 (1.41)	56 (0.17)	232 (0.82)	159 (0.76)
NU[6]	433 (5.92)	249 (27.75)	1425 (12.24)	753 (163.32)
LFR[6]	39 (0.77)	40 (0.27)	40 (0.44)	21 (0.19)
LFR[7]	72 (0.92)	73 (2.99)	73 (1.57)	34 (1.76)
LFR[8]	137 (1.45)	138 (39.11)	138 (10.74)	59 (20.84)

Table 1: Comparison of automata state spaces of Owl and Slim against their reduced versions, Owl-Red and Slim-Red. We report number of states (smallest in bold) and runtime (in seconds). Within Owl-Red and Slim-Red we only measure the time it took to reduce the automaton.

Theorem 5. (1) $\mathcal{L}(\mathcal{A}) = [\text{GF}\varphi]$, (2) \mathcal{A} is GFM and (3) \mathcal{A} has $2^{\mathcal{O}(|\varphi|)}$ states.

Note that as, for a co-safety/reachability property φ , the NFA \mathcal{N}_φ of φ has $2^{\mathcal{O}(|\varphi|)}$ states (De Giacomo and Vardi 2013), so does our GFM automaton \mathcal{A} for $\text{GF}\varphi$. Current GFM constructions such as (Sickert et al. 2016; Hahn et al. 2020) normally output a DBA for $\text{GF}\varphi$, which in general has $2^{2^{\mathcal{O}(|\varphi|)}}$ states. Therefore, our specialised GFM construction for repeated reachability property $\text{GF}\varphi$ can be *exponentially more efficient* than current approaches.

7 Experiments for GFM state reduction

In this section, we validate the performance of our GFM state-space reduction on LTL specifications extracted from literature, which we define below. The full benchmark set, with consistent state-space reductions on complex LTL specifications from over 8 sources, is in the extended version.

Trigger with Delayed Response (TDR). This pattern requires a followed by b after n steps, to hold infinitely often. We define this fragment as $\text{TDR}[n] = \text{GF}(a \wedge X^n b)$ where X^n denotes n nested applications of the next operator.

Liberouter (LIB). An example from the Liberouter verification project (Holeček et al. 2004) checks liveness of binary signals, ensuring at least one signal does not become permanently stuck. Formally, $\text{LIB}[n] = \text{GF}((a \wedge X \neg a) \vee (\neg a \wedge X a) \vee (b \wedge X \neg b) \vee (\neg b \wedge X b) \dots)$, with n numbers of signals.

Bounded Retransmission Protocol (BRP). Based on

(Baier et al. 2023), BRP models that whenever a message is sent, a corresponding acknowledgement must occur within a bounded number of steps. $\text{BRP}[n] = \mathbf{G}(\text{"msg_sent"} \rightarrow \mathbf{F}(\text{"ack_send"} \wedge \varphi_n))$, where the subformula φ_n ensures that "ack_rev" occurs within n steps after "ack_send". We define $\varphi_n = \text{"ack_rev"} \vee \mathbf{X}(\text{"ack_rev"} \vee \mathbf{X}(\dots \vee \mathbf{X}(\text{"ack_rev"})))$ with \mathbf{X} applied n times. A smaller n enforces stricter guarantees, while a larger n relaxes them.

Etenessami-Holzmann Patterns (EHP). To evaluate the effectiveness of Büchi automata optimisations, (Etenessami and Holzmann 2000) present deeply nested specifications. One formula is given by $\text{EHP} = a \mathbf{U}(b \wedge \mathbf{X}(c \wedge \mathbf{F}(d \wedge \mathbf{XF}(e \wedge \mathbf{XF}(f \wedge \mathbf{XF}g))))))$. It expresses that condition a must hold continuously until a specific sequence of events is triggered.

Nested Until Dependencies (NU). We define the nested-until pattern $\text{NU}[n] = \mathbf{G}(p_1 \rightarrow \varphi_n)$ where φ_k is recursively defined by $\varphi_k = p_k \mathbf{U} \varphi_{k+1}$ for $k < n$, and $\varphi_n = p_n \mathbf{U} p_{n+1}$. This pattern was used in the runtime verification benchmarks for SystemC models (Tabakov, Rozier, and Vardi 2012).

Layered Fairness and Reachability (LFR). As part of their evaluation, (Müller and Sickert 2017) introduced $\text{LFR}[n] = \mathbf{F}(b_1) \wedge (\mathbf{F}(b_2 \wedge \dots \mathbf{F}(b_n))) \wedge \mathbf{GF}(a_1 \wedge \mathbf{X}(a_2 \wedge \dots \mathbf{X}(a_n)))$. The left-hand conjunct captures a nested reachability, while the right-hand side encodes a fairness condition, requiring a_1, \dots, a_n to reoccur, with strict temporal progression.

Experiment Setup. In order to translate these LTL specifications to GFM automata, we use two state-of-the-art approaches. Owl (Kretínský, Meggendorfer, and Sickert 2018), which we denote by Owl, and the "slim" GFM construction of (Hahn et al. 2020), which we implemented ourselves and denote by Slim. We refer to reduced Owl automata as Owl-Red and to reduced "slim" automata as Slim-Red. We use SPOT (Duret-Lutz 2013) for parts of patterns, for the GFG minimisation and we have used LTL datasets available in the Owl repository. All experiments were run on an 8-core ARM chip, with 16GB of RAM.

Results. Table 1 compares automata states before and after our reduction. We report the LTL to GFM construction time (in seconds) for Owl and Slim, and reduction time (excluding initial construction) for Owl-Red and Slim-Red.

Across all LTL patterns, the smallest automata (marked in bold) are consistently reduced ones. Owl features advanced formula simplification and post-processing optimisations and therefore constructs smaller automata than the "slim" GFM construction. Since our "slim" GFM construction is not as optimised as Owl, we can consistently reduce the state space of all "slim" GFMs and for cases like TDR and LFR, we obtain smaller Slim-Red than Owl-Red.

Note that for LFR, Owl returns incomplete automata, while our reduction returns complete ones with an explicit sink (see extended version). For the remaining cases NU, EHP, BRP and LIB, the smallest automata are our Owl-Red. More results, are in the extended version.

8 Experiments for $\text{GF}\varphi$

The first two LTL specification patterns above, namely TDR and LIB, refer to formulas that have the syntactic $\text{GF}\varphi$

Pattern	Owl-Red	Slim-Red	GFM-GF
TDR[6]	64 (0.49)	34 (0.10)	7 (0.08)
TDR[7]	128 (0.61)	66 (0.14)	8 (0.08)
TDR[8]	256 (1.17)	130 (0.19)	9 (0.08)
TDR[9]	512 (4.27)	258 (0.48)	10 (0.10)
TDR[10]	1024 (23.67)	514 (0.91)	11 (0.09)
LIB[6]	14 (0.63)	66 (129.03)	13 (0.09)
LIB[7]	16 (1.47)	timeout	15 (0.12)
LIB[8]	18 (5.95)	timeout	17 (0.20)
LIB[9]	20 (37.78)	timeout	19 (0.23)

Table 2: State and runtime (in seconds) comparison of our direct construction GFM-GF against our reduced automata, Owl-Red and GFM-Red. The runtime for Owl-Red and GFM-Red includes construction and consecutive application of our reduction (in seconds). Timeout is 300 seconds.

structure and therefore, we can use them to evaluate our direct GFM automata construction. In particular, we show below that the direct construction gives us a large improvement with respect to the standard state-of-the-art GFM constructions (Kretínský, Meggendorfer, and Sickert 2018; Hahn et al. 2020), even after applying our GFM reduction.

Experiment Setup. Table 2 compares our direct GFM construction (denoted by GFM-GF) to the automata obtained by our GFM state-space reduction. Again, we denote number of states of reduced Owl and "slim" GFM automata (Owl-Red, Slim-Red) but in contrast to Table 1, the runtime (in seconds), now contains both LTL to GFM construction and subsequent application of our reduction.

Results. For the LTL pattern TDR, Table 2 clearly shows that our direct construction GFM-GF, significantly outperforms both Owl-Red and Slim-Red, both in state size and runtime. Even when comparing the state size and construction time of GFM-GF to Owl and Slim, we can build exponentially more succinct automata, in a fraction of the time. For the LIB pattern, GFM-GF builds automata with a similar amount of states than Owl-Red, but again, we only require a fraction of the time. LIB[7] and LIB[8] timed out during reduction, while LIB[9] timed out during automata construction. More results, are in the extended version.

9 Conclusion

We conclude the paper by observing that our results immediately have an impact on a wide range of applications that use GFM automata, such as obtaining smaller strategies for planning problems, improving the efficiency of probabilistic verification (Hahn et al. 2015; Sickert and Kretínský 2016), and reinforcement learning (Hahn et al. 2020; Jackermeier and Abate 2025). Furthermore, our results can be used to reduce the GFM automata obtained from variants of LTL, including LTLf+ and PPLTL+ recently proposed in (Aminof et al. 2025; De Giacomo et al. 2025).

Acknowledgements

We would like to thank the anonymous reviewers for their constructive feedback, that helped improve the paper. This work was supported in part by the CAS Project for Young Scientists in Basic Research (Grant No. YSBR-040), ISCAS Basic Research (Grant Nos. ISCAS-JCZD-202406, ISCAS-JCZD-202302), ISCAS New Cultivation Project ISCAS-PYFX-202201, the UKRI Erlangen AI Hub on Mathematical and Computational Foundations of AI (No. EP/Y028872/1) and by the EPSRC through grants EP/X03688X/1 and EP/X042596/1.

References

- Aminof, B.; De Giacomo, G.; Rubin, S.; and Vardi, M. Y. 2025. LTLf+ and PPLTL+: Extending LTLf and PPLTL to Infinite Traces. In *IJCAI*.
- Bacchus, F.; Boutilier, C.; and Grove, A. J. 1996. Rewarding Behaviors. In *AAAI/IAAI, Vol. 2*, 1160–1167. AAAI Press / The MIT Press.
- Bacchus, F.; Boutilier, C.; and Grove, A. J. 1997. Structured Solution Methods for Non-Markovian Decision Processes. In *AAAI/IAAI*, 112–117. AAAI Press / The MIT Press.
- Baier, C.; Größer, M.; and Bertrand, N. 2012. Probabilistic ω -automata. *J. ACM*, 59(1): 1:1–1:52.
- Baier, C.; and Katoen, J. 2008. *Principles of model checking*. MIT Press.
- Baier, C.; Kiefer, S.; Klein, J.; Müller, D.; and Worrell, J. 2023. Markov chains and unambiguous automata. *J. Comput. Syst. Sci.*, 136: 113–134.
- Brafman, R. I.; and Giacomo, G. D. 2024. Regular decision processes. *Artif. Intell.*, 331: 104113.
- Chang, E. Y.; Manna, Z.; and Pnueli, A. 1992. Characterization of Temporal Property Classes. In *ICALP*, volume 623 of *Lecture Notes in Computer Science*, 474–486. Springer.
- de Alfaro, L. 1997. *Formal verification of probabilistic systems*. Ph.D. thesis, Stanford University, USA.
- De Giacomo, G.; Li, Y.; Schewe, S.; Weinhuber, C.; and Yu, P. 2025. Solving MDPs with LTLf+ and PPLTL+ Temporal Objectives. In *IJCAI*.
- De Giacomo, G.; and Vardi, M. Y. 2013. Linear Temporal Logic and Linear Dynamic Logic on Finite Traces. In *IJCAI*, 854–860. IJCAI/AAAI.
- Duret-Lutz, A. 2013. Manipulating LTL Formulas Using Spot 1.0. In *ATVA*, volume 8172 of *Lecture Notes in Computer Science*, 442–445. Springer.
- Esparza, J.; and Kretínský, J. 2014. From LTL to Deterministic Automata: A Safrless Compositional Approach. In *CAV*, volume 8559 of *Lecture Notes in Computer Science*, 192–208. Springer.
- Etessami, K.; and Holzmann, G. J. 2000. Optimizing Büchi Automata. In *CONCUR*, volume 1877 of *Lecture Notes in Computer Science*, 153–167. Springer.
- Hahn, E. M.; Li, G.; Schewe, S.; Turrini, A.; and Zhang, L. 2015. Lazy Probabilistic Model Checking without Determinisation. In *CONCUR*, volume 42 of *LIPICs*, 354–367. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Hahn, E. M.; Perez, M.; Schewe, S.; Somenzi, F.; Trivedi, A.; and Wojtczak, D. 2020. Good-for-MDPs Automata for Probabilistic Analysis and Reinforcement Learning. In *TACAS (1)*, volume 12078 of *Lecture Notes in Computer Science*, 306–323. Springer.
- Henzinger, T. A.; and Piterman, N. 2006. Solving Games Without Determinization. In *CSL*, volume 4207 of *Lecture Notes in Computer Science*, 395–410. Springer.
- Henzinger, T. A.; Prakash, A.; and Thejaswini, K. S. 2025. Resolving Nondeterminism with Randomness. In Gawrychowski, P.; Mazowiecki, F.; and Skrzypczak, M., eds., *50th International Symposium on Mathematical Foundations of Computer Science, MFCS 2025, August 25-29, 2025, Warsaw, Poland*, volume 345 of *LIPICs*, 57:1–57:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Holeček, J.; Kratochvíla, T.; Řehák, V.; Šafránek, D.; Šimeček, P.; et al. 2004. Verification results in Liberouter project.
- Jackermeier, M.; and Abate, A. 2025. DeepLTL: Learning to Efficiently Satisfy Complex LTL Specifications for Multi-Task RL. In *International Conference on Learning Representations (ICLR)*.
- Klein, J.; Müller, D.; Baier, C.; and Klüppelholz, S. 2014. Are Good-for-Games Automata Good for Probabilistic Model Checking? In *LATA*, volume 8370 of *Lecture Notes in Computer Science*, 453–465. Springer.
- Kretínský, J.; Meggendorfer, T.; and Sickert, S. 2018. Owl: A Library for ω -Words, Automata, and LTL. In *ATVA*, volume 11138 of *Lecture Notes in Computer Science*, 543–550. Springer.
- Kwiatkowska, M. Z.; Norman, G.; and Parker, D. 2011. PRISM 4.0: Verification of Probabilistic Real-Time Systems. In *CAV*, volume 6806 of *Lecture Notes in Computer Science*, 585–591. Springer.
- Lacerda, B.; Parker, D.; and Hawes, N. 2014. Optimal and dynamic planning for Markov decision processes with co-safe LTL specifications. In *IROS*, 1511–1516. IEEE.
- Li, Y.; Paul, S.; Schewe, S.; and Tang, Q. 2025. Accelerating Markov Chain Model Checking: Good-for-Games Meets Unambiguous Automata. In Piskac, R.; and Rakamarić, Z., eds., *CAV*, volume 15932 of *Lecture Notes in Computer Science*, 276–298. Springer.
- Littman, M. L. 1997. Probabilistic Propositional Planning: Representations and Complexity. In *AAAI/IAAI*, 748–754. AAAI Press / The MIT Press.
- Littman, M. L.; Goldsmith, J.; and Mundhenk, M. 1998. The Computational Complexity of Probabilistic Planning. *J. Artif. Intell. Res.*, 9: 1–36.
- Manna, Z.; and Pnueli, A. 1990. A Hierarchy of Temporal Properties. In *PODC*, 377–410. ACM.
- Meyer, P. J.; Sickert, S.; and Luttenberger, M. 2018. Strix: Explicit Reactive Synthesis Strikes Back! In *CAV (1)*, volume 10981 of *Lecture Notes in Computer Science*, 578–586. Springer.
- Müller, D.; and Sickert, S. 2017. LTL to Deterministic Emerson-Lei Automata. In *GandALF*, volume 256 of *EPTCS*, 180–194.

- Pnueli, A. 1977. The Temporal Logic of Programs. In *FOCS*, 46–57. IEEE Computer Society.
- Radi, B. A.; and Kupferman, O. 2022. Minimization and Canonization of GFG Transition-Based Automata. *Log. Methods Comput. Sci.*, 18(3).
- Safra, S. 1988. On the Complexity of omega-Automata. In *FOCS*, 319–327. IEEE Computer Society.
- Schillinger, P.; Bürger, M.; and Dimarogonas, D. V. 2019. Hierarchical LTL-task mdps for multi-agent coordination through auctioning and learning. *The international journal of robotics research*.
- Sickert, S.; Esparza, J.; Jaax, S.; and Kretínský, J. 2016. Limit-Deterministic Büchi Automata for Linear Temporal Logic. In *CAV (2)*, volume 9780 of *Lecture Notes in Computer Science*, 312–332. Springer.
- Sickert, S.; and Kretínský, J. 2016. MoChiBA: Probabilistic LTL Model Checking Using Limit-Deterministic Büchi Automata. In *ATVA*, volume 9938 of *Lecture Notes in Computer Science*, 130–137.
- Tabakov, D.; Rozier, K. Y.; and Vardi, M. Y. 2012. Optimized temporal monitors for SystemC. *Formal Methods Syst. Des.*, 41(3): 236–268.
- Thiébaux, S.; Gretton, C.; Slaney, J. K.; Price, D.; and Kabanza, F. 2006. Decision-Theoretic Planning with non-Markovian Rewards. *J. Artif. Intell. Res.*, 25: 17–74.
- Vardi, M. Y.; and Wolper, P. 1986. An Automata-Theoretic Approach to Automatic Program Verification (Preliminary Report). In *LICS*, 332–344. IEEE Computer Society.
- Yang, C.; Littman, M. L.; and Carbin, M. 2022. On the (In)Tractability of Reinforcement Learning for LTL Objectives. In *IJCAI*, 3650–3658. ijcai.org.
- Yu, P.; Li, Y.; Parker, D.; and Kwiatkowska, M. 2025. Planning with Linear Temporal Logic Specifications: Handling Quantifiable and Unquantifiable Uncertainty. *CoRR*, abs/2502.19603.