Quasipolynomial Computation of Nested Fixpoints

Daniel Hausmann (✉) and Lutz Schröder (✉) *

Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany
{daniel.hausmann,lutz.schroeder}@fau.de

Abstract. It is well-known that the winning region of a parity game with n nodes and k priorities can be computed as a k -nested fixpoint of a suitable function; straightforward computation of this nested fixpoint requires $\mathcal{O}(n^{\frac{k}{2}})$ iterations of the function. Calude et al.’s recent quasipolynomial-time parity game solving algorithm essentially shows how to compute the same fixpoint in only quasipolynomially many iterations by reducing parity games to quasipolynomially sized safety games. Universal graphs have been used to modularize this transformation of parity games to equivalent safety games that are obtained by combining the original game with a universal graph. We show that this approach naturally generalizes to the computation of solutions of systems of *any* fixpoint equations over finite lattices; hence, the solution of fixpoint equation systems can be computed by quasipolynomially many iterations of the equations. We present applications to modal fixpoint logics and games beyond relational semantics. For instance, the model checking problems for the energy μ -calculus, finite latticed μ -calculi, and the graded and the (two-valued) probabilistic μ -calculus – with numbers coded in binary – can be solved via nested fixpoints of functions that differ substantially from the function for parity games but still can be computed in quasipolynomial time; our result hence implies that model checking for these μ -calculi is in QP. Moreover, we improve the exponent in known exponential bounds on satisfiability checking.

Keywords: Fixpoint theory, model checking, satisfiability checking, parity games, energy parity games, μ -calculus

1 Introduction

Fixpoints are pervasive in computer science, governing large portions of recursion theory, concurrency theory, logic, and game theory. One famous example are parity games, which are central, e.g., to networks and infinite processes [5], tree automata [54], and μ -calculus model checking [21]. Winning regions in parity games can be expressed as nested fixpoints of particular set functions (e.g. [8, 20]). In recent breakthrough work on the solution of parity games in quasipolynomial

* Work forms part of the DFG-funded project CoMoC (SCHR 1118/15-1, MI 717/7-1).

time, Calude et al. [9] essentially show how to compute this particular fixpoint in quasipolynomial time, that is, in time $2^{\mathcal{O}((\log n)^c)}$ for some constant c . Subsequently, it has been shown [16, 17, 34] that universal graphs (that is, even graphs into which every even graph of a certain size embeds by a graph morphism) can be used to transform parity games to equivalent safety games obtained by pairing the original game with a universal graph; the size of these safety games is determined by the size of the employed universal graphs and it has been shown [16, 17] that there are universal graphs of quasipolynomial size. This yields a uniform algorithm for solving parity games to which all currently known quasipolynomial algorithms for parity games have been shown to instantiate using appropriately defined universal graphs [16, 17].

Briefly, our contribution in the present work is to show that the method of using universal graphs to solve parity games generalizes to the computation of nested fixpoints of arbitrary functions over finite lattices. That is, given functions $f_i : L^{k+1} \rightarrow L$, $0 \leq i \leq k$ on a finite lattice L , we give an algorithm that uses universal graphs to compute the solutions of systems of equations

$$X_i =_{\eta_i} f_i(X_0, \dots, X_k) \quad 0 \leq i \leq k$$

where $\eta_i = \text{GFP}$ (greatest fixpoint) or $\eta_i = \text{LFP}$ (least fixpoint). Since there are universal graphs of quasipolynomial size, the algorithm requires only quasipolynomially many iterations of the functions f_i and hence runs in quasipolynomial time, provided that all f_i are computable in quasipolynomial time. While it seems plausible that this time bound may also be obtained by translating equation systems to equivalent standard parity games by emulating Turing machines to encode the functions f_i as Boolean circuits (leading to many additional states but avoiding exponential blowup during the process), we emphasize that the main point of our result is not so much the ensuing time bound but rather the insight that universal graphs and hence many algorithms for parity games can be used on a much more general level which yields a precise (and relatively low) quasipolynomial bound on the number of function calls that are required to obtain solutions of fixpoint equation systems.

In more detail, the method of Calude et al. can be described as annotating nodes of a parity game with histories of quasipolynomial size and then solving this annotated game, but with a safety winning condition instead of the much more involved parity winning condition. It has been shown that these histories can be seen as nodes in universal graphs, in a more general reduction of parity games to safety games in which nodes from the parity game are annotated with nodes from a universal graph. This method has also been described as pairing *separating automata* with safety games [17]. It has been shown [16, 17] that there are exponentially sized universal graphs (essentially yielding the basis for e.g. the fixpoint iteration algorithm [8] or the small progress measures algorithm [33]) and quasipolynomially sized universal graphs (corresponding, e.g., to the succinct progress measure algorithm [34], or to the recent quasipolynomial variant of Zielonka's algorithm [46]).

Hasuo et al. [28], and more generally, Baldan et al. [4] show that nested fixpoints in highly general settings can be computed by a technique based on

progress measures, implicitly using exponentially sized universal graphs, obtaining an exponential bound on the number of iterations. Our technique is based on showing that one can make explicit use of universal graphs, correspondingly obtaining a quasipolynomial upper bound on the number of iterations. In both cases, computation of the nested fixpoint is reduced to a single (least or greatest depending on exact formulation) fixpoint of a function that extends the given set function to keep track of the exponential and quasipolynomial histories, respectively, in analogy to the previous reduction of parity games to safety games. Our central result can then be phrased as saying that the method of transforming parity conditions to safety conditions using universal graphs generalizes from solving parity games to solving systems of equations that use arbitrary functions over finite lattices. We use *fixpoint games* [4, 53] to obtain the crucial result that the solutions of equation systems have history-free witnesses, in analogy to history-freeness of winning strategies in parity games. These fixpoint games have exponential size but we show how to extract *polynomial-size* witnesses for winning strategies of Eloise, and use these witnesses to show that any node won by Eloise is also won in the safety game obtained by a universal graph. For the backwards direction, we show that a witness for satisfaction of the safety condition regarding the universal graph induces a winning strategy in the fixpoint game. This proves that universal graphs can be used to compute nested fixpoints of arbitrary functions over finite lattices and hence yields the quasipolynomial upper bound for computation of nested fixpoints. Moreover, we present a progress measure algorithm that uses the nodes of a quasipolynomial universal graph to measure progress and that can be used to efficiently compute nested fixpoints of arbitrary functions over finite lattices.

As an immediate application of these results, we improve known deterministic algorithms for solving *energy parity games* [10], that is, parity games in which edges have additional integer weights and for which the winning condition is a combined parity condition and a (quantitative) positivity condition on the sum of the accumulated weights. Our results also show that the model checking problem for the associated *energy μ -calculus* [2] is in QP. In a similar fashion, we obtain quasipolynomial algorithms for model checking in latticed μ -calculi [7] in which the truth values of formulae are computed over arbitrary finite lattices, and for solving associated latticed parity games [36].

Furthermore, our results improve generic upper complexity bounds on model checking and satisfiability checking in the *coalgebraic μ -calculus* [14], which serves as a generic framework for fixpoint logics beyond relational semantics. Well-known instances of the coalgebraic μ -calculus include the alternating-time μ -calculus [1], the graded μ -calculus [38], the (two-valued) probabilistic μ -calculus [14, 41], and the monotone μ -calculus [22] (the ambient fixpoint logic of concurrent dynamic logic CPDL [48] and Parikh’s game logic [45]). This level of generality is achieved by abstracting system types as set functors and systems as coalgebras for the given functor following the paradigm of universal coalgebra [49]. It was previously shown [30] that the model checking problem for coalgebraic μ -calculi reduces to the computation of a nested fixpoint. This

fixpoint may be seen as a coalgebraic generalization of a parity game winning region but can be literally phrased in terms of small standard parity games (implying quasipolynomial run time) only in restricted cases. Our results show that the relevant nested fixpoint can be computed in quasipolynomial time in all cases of interest. Notably, we thus obtain as new specific upper bounds that even under binary coding of numbers, the model checking problems of both the graded μ -calculus and the probabilistic μ -calculus are in QP, even when the syntax is extended to allow for (monotone) polynomial inequalities.

Similarly, the satisfiability problem of the coalgebraic μ -calculus has been reduced to a computation of a nested fixpoint [31], and our present results imply a marked improvement in the exponent of the associated exponential time bound. Specifically, the nesting depth of the relevant fixpoint is exponentially smaller than the basis of the lattice. Our results imply that this fixpoint is computable in polynomial time so that the complexity of satisfiability checking in coalgebraic μ -calculi drops from $2^{\mathcal{O}(n^2 k^2 \log n)}$ to $2^{\mathcal{O}(nk \log n)}$ for formulae of size n and with alternation depth k .

Related Work The quasipolynomial bound on parity game solving has in the meantime been realized by a number of alternative algorithms. For instance, Jurdzinski and Lazic [34] use succinct progress measures to improve to quasilinear (instead of quasipolynomial) space; Fearnley et al. [24] similarly achieve quasilinear space. Lehtinen [40] and Boker and Lehtinen [6] present a quasipolynomial algorithm using register games. Parys [46] improves Zielonka’s algorithm [54] to run in quasipolynomial time. In particular the last algorithm is of interest as an additional candidate for generalization to nested fixpoints, due to the known good performance of Zielonka’s algorithm in practice. Daviaud et al. [19] generalize quasipolynomial-time parity game solving by providing a pseudo-quasipolynomial algorithm for mean-payoff parity games. On the other hand, Czerwinski et al. [17] give a quasipolynomial lower bound on universal trees, implying a barrier for prospective polynomial-time parity game solving algorithms. Chatterjee et al. [11] describe a quasipolynomial time set-based symbolic algorithm for parity game solving that is parametric in a *lift* function that determines how ranks of nodes depend on the ranks of their successors, and thereby unifies the complexity and correctness analysis of various parity game algorithms. Although part of the parity game structure is encapsulated in a set operator $CPre$, the development is tied to standard parity games, e.g. in the definition of the *best* function, which picks minimal or maximal ranks of successors depending on whether a node belongs to Abelard or Eloise.

Early work on the computation of unrestricted nested fixpoints has shown that greatest fixpoints require less effort in the fixpoint iteration algorithm, which can hence be optimized to compute nested fixpoints with just $\mathcal{O}(n^{\frac{k}{2}})$ calls of the functions at hand [42, 52], improving the previously known (straightforward) bound $\mathcal{O}(n^k)$; here, n denotes the size of the basis of the lattice and k the number of fixpoint operators. Recent progress in the field has established the above-mentioned approaches using progress measures [28] and fixpoint games [4] in general settings, both with a view to applications in coalgebraic model checking

like in the present paper. In comparison to the present work, the respective bounds on the required number of function iterations in the above unrestricted approaches all are exponential.

A preprint of our present results, specifically the quasipolynomial upper bound on function iteration in fixpoint computation, has been available as an arXiv preprint for some time [29]. Subsequent to this preprint, Arnold, Niwinski and Parys [3] have improved the actual run time by reducing the overhead incurred per iteration (and they give a form of quasipolynomial lower bound for universal-tree-based algorithms), working (like [29]) in the less general setting of directly nested fixpoints over powerset lattices; we show in Section 6 how such an improvement can be incorporated also in our lattice-based algorithm.

2 Notation and Preliminaries

Let U and V be sets, and let $R \subseteq U \times U$ be a binary relation on U . For $u \in U$, we then put $R(u) := \{v \in U \mid (u, v) \in R\}$. We put $[k] = \{0, \dots, k\}$ for $k \in \mathbb{N}$. *Labelled graphs* $G = (W, R)$ consist of a set W together with a relation $R \subseteq W \times A \times W$ where A is some set of labels; typically, we use $A = [k]$ for some $k \in \mathbb{N}$. An R -path in a labelled graph is a finite or infinite sequence $v_0, a_0, v_1, a_1, v_2, \dots$ (ending in a node from W if finite) such that $(v_i, a_i, v_{i+1}) \in R$ for all i . For $v \in W$ and $a \in A$, we put $R_a(v) = \{w \in W \mid (v, a, w) \in R\}$ and sometimes write $|G|$ to refer to $|W|$. As usual, we write U^* and U^ω for the sets of finite sequences or infinite sequences, respectively, of elements of U . The *domain* $\text{dom}(f)$ of a partial function $f : U \rightarrow V$ is the set of elements on which f is defined. As usual, the (*forward*) *image* of $A' \subseteq A$ under a function $f : A \rightarrow B$ is $f[A'] = \{b \in B \mid \exists a \in A'. f(a) = b\}$ and the *preimage* $f^{-1}[B']$ of $B' \subseteq B$ under f is defined by $f^{-1}[B'] = \{a \in A \mid \exists b \in B'. f(a) = b\}$. *Projections* $\pi_j : A_1 \times \dots \times A_m \rightarrow A_j$ for $1 \leq j \leq m$ are given by $\pi_i(a_1, \dots, a_m) = a_i$. We often regard (finite) sequences $\tau = u_0, u_1, \dots \in U^* \cup U^\omega$ of elements of U as partial functions of type $\mathbb{N} \rightarrow U$ and then write $\tau(i)$ to denote the element u_i , for $i \in \text{dom}(\tau)$. For $\tau \in U^* \cup U^\omega$, we define the set $\text{Inf}(\tau) = \{u \in U \mid \forall i \geq 0. \exists j > i. \tau(j) = u\}$ of elements that occur infinitely often in τ (so $\text{Inf}(\tau) = \emptyset$ for $\tau \in U^*$). An infinite R -path $v_0, p_0, v_1, p_1, \dots$ in a labelled graph $G = (W, R)$ with labels from $[k]$ is *even* if $\max(\text{Inf}(p_0, p_1, \dots))$ is even, and G is *even* if every infinite R -path in G is even. We write $\mathcal{P}(U)$ for the powerset of U , and U^m for the m -fold Cartesian product $U \times \dots \times U$.

Finite Lattices and Fixpoints A *finite lattice* $(L, \sqsubseteq)$ (often written just as L) consists of a non-empty finite set L together with a partial order $\sqsubseteq$ on L , such that there is, for all subsets $X \subseteq L$, a join $\bigsqcup X$ and a meet $\bigsqcap X$. The least and greatest elements of L are defined as $\top = \bigsqcup \emptyset$ and $\perp = \bigsqcap \emptyset$, respectively. A set $B_L \subseteq L$ such that $l = \bigsqcup \{b \in B_L \mid b \sqsubseteq l\}$ for all $l \in L$ is a *basis* of L . Given a finite lattice L , a function $g : L^k \rightarrow L$ is *monotone* if $g(V_1, \dots, V_k) \sqsubseteq g(W_1, \dots, W_k)$ whenever $V_i \sqsubseteq W_i$ for all $1 \leq i \leq k$. For monotone $f : L \rightarrow L$, we put

$$\text{GFP } f = \bigsqcup \{V \sqsubseteq L \mid V \sqsubseteq f(V)\} \quad \text{LFP } f = \bigsqcap \{V \sqsubseteq L \mid f(V) \sqsubseteq V\},$$

which, by the Knaster-Tarski fixpoint theorem, are the greatest and the least fixpoint of f , respectively. Furthermore, we define $f^0(V) = V$ and $f^{m+1}(V) = f(f^m(V))$ for $m \geq 0$, $V \sqsubseteq L$; since L is finite, we have $\text{GFP } f = f^n(\top)$ and $\text{LFP } f = f^n(\perp)$ by Kleene's fixpoint theorem. Given a finite set U and a natural number n , $(n^U, \sqsubseteq)$ is a finite lattice, where $n^U = \{f : U \rightarrow [n-1]\}$ denotes the function space from U to $[n-1]$ and $f \sqsubseteq g$ if and only if for all $u \in U$, $f(u) \leq g(u)$. For $n = 2$, we obtain the powerset lattice $(2^U, \sqsubseteq)$, also denoted by $\mathcal{P}(U)$, with least and greatest elements $\emptyset$ and U , respectively, and basis $\{\{u\} \mid u \in U\}$.

Parity games A *parity game* (V, E, Ω) consists of a set of *nodes* V , a left-total relation $E \subseteq V \times V$ of *moves* encoding the rules of the game, and a *priority function* $\Omega : V \rightarrow \mathbb{N}$, which assigns *priorities* $\Omega(v) \in \mathbb{N}$ to nodes $v \in V$. Moreover, each node belongs to exactly one of the two players **Eloise** or **Abelard**, where we denote the set of **Eloise's** nodes by $V_\exists$ and that of **Abelard's** nodes by $V_\forall$. A *play* $\rho \in V^\omega$ is an infinite sequence of nodes that follows the rules of the game, that is, such that for all $i \geq 0$, we have $(\rho(i), \rho(i+1)) \in E$. We say that an infinite play $\rho = v_0, v_1, \dots$ is *even* if the largest priority that occurs infinitely often in it (i.e. $\max(\text{Inf}(\Omega \circ \rho))$) is even, and *odd* otherwise, and call this property the *parity* of ρ . Player **Eloise** *wins* exactly the even plays and player **Abelard** *wins* all other plays. A (*history-free*) **Eloise-strategy** $s : V_\exists \rightharpoonup V$ is a partial function that assigns single moves $s(x)$ to **Eloise-nodes** $x \in \text{dom}(s)$. Given an **Eloise-strategy** s , a play ρ is an *s-play* if for all $i \in \text{dom}(\rho)$ such that $\rho(i) \in V_\exists$, we have $\rho(i+1) = s(\rho(i))$. An **Eloise-strategy** *wins* a node $v \in V$ if **Eloise** wins all *s*-plays that start at v . We have a dual notion of **Abelard-strategies**; *solving* a parity game consists in computing the *winning regions* $\text{win}_\exists$ and $\text{win}_\forall$ of the two players, that is, the sets of states that they respectively win by some strategy.

It is known that solving parity games is in $\text{NP} \cap \text{coNP}$ (and, more specifically, in $\text{UP} \cap \text{co-UP}$). Recently it has also been shown [9] that for parity games with n nodes and k priorities, $\text{win}_\exists$ and $\text{win}_\forall$ can be computed in quasipolynomial time $\mathcal{O}(n^{\log k+6})$. Another crucial property of parity games is that they are *history-free determined* [26], that is, that every node in a parity game is won by exactly one of the two players and then there is a history-free strategy for the respective player that wins the node.

3 Systems of Fixpoint Equations

We now introduce our central notion, that is, systems of fixpoint equations over a finite lattice. Throughout, we fix a finite lattice $(L, \sqsubseteq)$ and a basis B_L of L such that $\perp \notin B_L$, and $k+1$ monotone functions $f_i : L^{k+1} \rightarrow L$, $0 \leq i \leq k$.

Definition 3.1. A *system of equations* consists of $k+1$ equations of the form

$$X_i =_{\eta_i} f_i(X_0, \dots, X_k)$$

where $\eta_i \in \{\text{LFP}, \text{GFP}\}$, briefly referred to as f . For a partial valuation $\sigma : [k] \rightarrow L$, we inductively define

$$\llbracket X_i \rrbracket^\sigma = \eta_i X_i.f_i^\sigma,$$

where the function f_i^σ is given by

$$f_i^\sigma(A) = f_i(\llbracket X_0 \rrbracket^{\sigma'}, \dots, \llbracket X_{i-1} \rrbracket^{\sigma'}, A, \text{ev}(\sigma', i+1), \dots, \text{ev}(\sigma', k))$$

for $A \in L$, where $(\sigma[i \mapsto A])(j) = \sigma(j)$ for $j \neq i$ and $(\sigma[i \mapsto A])(i) = A$, $\sigma' = \sigma[i \mapsto A]$ and where $\text{ev}(\sigma, j) = \sigma(j)$ if $j \in \text{dom}(\sigma)$ and $\text{ev}(\sigma, j) = \llbracket X_j \rrbracket^\sigma$ otherwise (the latter clause handles *free variables*). Then, the *solution* of the system of equations is $\llbracket X_k \rrbracket^\epsilon$ where $\epsilon : [k] \rightarrow L$ denotes the empty valuation (i.e. $\text{dom}(\epsilon) = \emptyset$). Similarly, we can obtain solutions for the other components as $\llbracket X_i \rrbracket^\epsilon$ for $0 \leq i < k$; we drop the valuation index if no confusion arises, and sometimes write $\llbracket X_i \rrbracket_f$ to make the equation system f explicit. We denote by E^{f_0} the solution $\llbracket X_k \rrbracket$ for the *canonical system of equations* of the particular shape

$$X_i =_{\eta_i} X_{i-1} \quad X_0 =_{\text{GFP}} f_0(X_0, \dots, X_k),$$

where $0 < i \leq k$, $\eta_i = \text{LFP}$ for odd i and $\eta_i = \text{GFP}$ for even i .

Example 3.2. (1) *Parity games and the modal μ -calculus*: Let (V, E, Ω) be a parity game with priorities 0 to k , take $L = \mathcal{P}(V)$, and consider the canonical system of fixpoint equations $E^{f_\exists}$ for the function $f_\exists : \mathcal{P}(V)^{k+1} \rightarrow \mathcal{P}(V)$ given by

$$f_\exists(V_0, \dots, V_k) = \{v \in V_\exists \mid E(v) \cap V_{\Omega(v)} \neq \emptyset\} \cup \{v \in V_\forall \mid E(v) \subseteq V_{\Omega(v)}\}$$

for $(V_0, \dots, V_k) \in \mathcal{P}(V)^{k+1}$. It is well known that $\text{win}_\exists = E^{f_\exists}$, i.e. parity games can be solved by solving fixpoint equation systems. Intuitively, $v \in f_\exists(V_0, \dots, V_k)$ iff Eloise can enforce that some node in $V_{\Omega(v)}$ is reached in the next step. The nested fixpoint expressed by $E^{f_\exists}$ (in which least (greatest) fixpoints correspond to odd (even) priorities) is constructed in such a way that Eloise only has to rely infinitely often on an argument V_i for odd i if she can also ensure that some argument V_j for $j > i$ is used infinitely often.

Model checking for the *modal μ -calculus* [35] and solving parity games are linear-time equivalent problems. Formulae of the μ -calculus are evaluated over Kripke frames (U, R) with set of states U and transition relation R . Formulae ϕ of the μ -calculus can be directly represented as equation systems over the lattice $\mathcal{P}(U)$ by recursively translating ϕ to equations, mapping subformulae $\mu X_i. \psi(X_0, \dots, X_k)$ and $\nu X_j. \chi(X_0, \dots, X_k)$ to equations

$$X_i =_\mu \psi(X_0, \dots, X_k) \quad X_j =_\nu \chi(X_0, \dots, X_k),$$

and interpreting the modalities $\Diamond$ and $\Box$ by functions

$$f_\Diamond(X) = \{u \in U \mid R(u) \cap X \neq \emptyset\} \quad f_\Box(X) = \{u \in U \mid R(u) \subseteq X\}$$

The solution of the resulting system of equations then is the truth set of the formula ϕ , that is, model checking for the model μ -calculus reduces to solving fixpoint equation systems. Furthermore, satisfiability checking for the modal μ -calculus can be reduced to solving so-called *satisfiability games* [25], that is, parity games that are played over the set of states of a determinized parity automaton. These satisfiability games can be expressed as systems of fixpoint equations, where the functions track transitions in the determinized automaton.

(2) *Energy parity games and the energy μ -calculus*: Energy parity games [10] are two-player games played over weighted game arenas (V, E, w, Ω) , where $w : E \rightarrow \mathbb{Z}$ assigns integer weights to edges. The winning condition is the combination of a parity condition with a (quantitative) positivity condition on the sum of the accumulated weights. It has been shown [2, 10], that $b = n \cdot d \cdot W$ is a sufficient upper bound on energy level accumulations in energy parity games with n nodes, k priorities and maximum absolute weight W . We define a function $f_{\exists}^e : ((b+1)^V)^{k+1} \rightarrow (b+1)^V$ over the finite lattice $(b+1)^V$ (whose elements are functions from V to the set $\{0, \dots, b+1\}$) by putting

$$(f_{\exists}^e(V_0, \dots, V_k))(v) = \begin{cases} \min(\text{en}(v, V_{\Omega(v)})) & \text{if } v \in V_{\exists} \\ \max(\text{en}(v, V_{\Omega(v)})) & \text{if } v \in V_{\forall}, \end{cases}$$

for $(V_0, \dots, V_k) \in ((b+1)^V)^{k+1}$ and $v \in V$, using $\text{en}(v, \sigma)$ as abbreviation for

$$\begin{aligned} \text{en}(v, \sigma) = & \{n \in \{0, \dots, b\} \mid \exists u \in E(v). n = \max\{0, \sigma(u) - w(v, u)\}\} \cup \\ & \{b+1 \mid \exists u \in E(v). \sigma(u) - w(v, u) > b \text{ or } \sigma(u) > b\}, \end{aligned}$$

where $\sigma : V \rightarrow \{0, \dots, b+1\}$. Then it follows from the results of [2] that player Eloise wins a node v in the energy parity game with minimal initial credit $c < b+1$ if $(E f_{\exists}^e)(v) = c$, that is, if the solution of the canonical equation system over $f_{\exists}^e$ maps v to a value c that is at most b .

The *energy μ -calculus* [2] is the fixpoint logic that corresponds to energy parity games. Its formulae are evaluated over *weighted game structures* and involve operators $\Diamond_E \phi$ and $\Box_E \phi$ that are evaluated depending on the energy function $\llbracket \phi \rrbracket : V \rightarrow \{0, \dots, b+1\}$ that is obtained by first evaluating the argument formula ϕ . The semantics of the diamond operator then is an energy function that assigns, to each state v , the least energy value $c \in \{0, \dots, b+1\}$ such that there is a move from v to some node u such that the credit c suffices to take the move from v to u and retain an energy level of at least $\llbracket \phi \rrbracket(u)$. Formulae can be translated to equation systems over the finite lattice $(b+1)^V$, where the functions for modal operators are defined according to their semantics as presented in [2]. Solving these equation systems then amounts to model checking energy μ -calculus formulae over weighted game structures.

(3) *Latticed μ -calculi*: In latticed μ -calculi [7], formulae are evaluated over complete lattices L rather than the powerset lattice; for finite lattices L , formulae of latticed μ -calculi hence can be translated to fixpoint equation systems over L , so that model checking reduces to solving equation systems. An associated latticed

variant of games has been introduced in [36] and for finite lattices L , solving latticed parity games over L reduces to solving equation systems over L .

(4) *The coalgebraic μ -calculus and coalgebraic parity games:* The coalgebraic μ -calculus [14] supports generalized modal branching types by using *predicate liftings* to interpret formulae over T -coalgebras, that is, over structures whose transition type is specified by an endofunctor T on the category of sets. For instance the functors $T = \mathcal{P}$, $T = \mathcal{D}$ and $T = \mathcal{G}$ map sets X to their powerset $\mathcal{P}(X)$, the set of probability distributions $\mathcal{D}(X) = \{f : X \rightarrow [0, \dots, 1]\}$ over X , and to the set of multisets $\mathcal{G}(X) = \{f : X \rightarrow \mathbb{N}\}$ over X , respectively. The corresponding T -coalgebras then are Kripke frames (for $T = \mathcal{P}$), Markov chains (for $T = \mathcal{D}$) and graded transition systems (for $T = \mathcal{G}$), respectively. Instances of the coalgebraic μ -calculus comprise, e.g. the two-valued probabilistic μ -calculus [14, 41] with modalities $\Diamond_p \phi$ for $p \in [0, \dots, 1]$, expressing ‘the next state satisfies ϕ with probability more than p ’; the graded μ -calculus [38] with modalities $\Diamond_g \phi$ for $g \in \mathbb{N}$, expressing ‘there are more than ϕ successor states that satisfy ϕ ’; or the alternating-time μ -calculus [1] that is interpreted over concurrent game frames and uses modalities $\langle D \rangle \phi$ for finite $D \subseteq \mathbb{N}$ (encoding a *coalition*) that express that ‘coalition D has a joint strategy to enforce ϕ ’.

It has been shown in previous work [30] that model checking for coalgebraic μ -calculi against coalgebras with state space U reduces to solving a canonical fixpoint equation system over the powerset lattice $\mathcal{P}(U)$, where the involved function interprets modal operators using predicate liftings, as described in [14, 30]. This canonical equation system can alternatively be seen as the winning region of Eloise in *coalgebraic parity games*, a highly general variant of parity games where the game structure is a coalgebra and nodes are annotated with modalities. Examples include *two-valued probabilistic parity games* and *graded parity games* in which nodes and edges are annotated with probabilities or grades, respectively. In order to win a node v , player Eloise then has to have a strategy that picks a *set* of moves to nodes that in turn are *all* won by Eloise, and such that the joint probability (joint grade) of the picked moves is greater than the probability (grade) that is assigned to v . It is known that solving coalgebraic parity games reduces to solving fixpoint equation systems [30].

Furthermore, the satisfiability problem of the coalgebraic μ -calculus has been reduced to solving canonical fixpoint equations systems over lattices $\mathcal{P}(U)$, where U is the state set of a determinized parity automaton and where the innermost equation checks for joint one-step satisfiability of sets of coalgebraic modalities [31]. By interpreting coalgebraic formulae over finite lattices d^U rather than over powerset lattices, one obtains the *finite-valued coalgebraic μ -calculus* (with values $\{0, \dots, d\}$), which has the *finite-valued probabilistic μ -calculus* (e.g. [43]) as an instance. Model checking for the finite-valued probabilistic μ -calculus hence reduces to solving equation systems over the finite lattice $d^{|U|}$, where $\{0, \dots, d\}$ encodes a finite set of probabilities.

4 Fixpoint Games and History-free Witnesses

We instantiate the existing notion of fixpoint games [4, 53], which characterize solutions of equation systems, to our setting (that is, to finite lattices), and then use these games as a technical tool to establish our crucial notion of history-freeness for systems of fixpoint equations.

Definition 4.1 (Fixpoint games). Let $X_i =_{\eta_i} f_i(X_0, \dots, X_k)$, $0 \leq i \leq k$, be a system of fixpoint equations. The associated *fixpoint game* is a parity game (V, E, Ω) with set of nodes $V = (B_L \times [k]) \cup L^{k+1}$, where nodes from $B_L \times [k]$ belong to player **Eloise** and nodes from L^{k+1} belong to player **Abelard**. For nodes $(u, i) \in B_L \times [k]$, we put

$$E(u, i) = \{(U_0, \dots, U_k) \in L^{k+1} \mid u \sqsubseteq f_i(U_0, \dots, U_k)\},$$

and for nodes $(U_0, \dots, U_k) \in L^{k+1}$, we put

$$E(U_0, \dots, U_k) = \{(u, i) \in B_L \times [k] \mid u \sqsubseteq U_i\}.$$

The *alternation depth* $\text{ad}(i)$ of an equation $X_i =_{\eta_i} f_i(X_0, \dots, X_k)$ is defined as ad_i^μ if $\eta_i = \mu$ and as ad_i^ν if $\eta_i = \nu$, where $\text{ad}_i^\mu, \text{ad}_i^\nu$ are recursively defined by

$$\text{ad}_i^\mu = \begin{cases} \text{ad}_{i-1}^\mu & i > 0, \eta_{i-1} = \mu \\ \text{ad}_{i-1}^\nu + 1 & i > 0, \eta_{i-1} = \nu \\ 1 & i = 0 \end{cases} \quad \text{ad}_i^\nu = \begin{cases} \text{ad}_{i-1}^\mu + 1 & i > 0, \eta_{i-1} = \mu \\ \text{ad}_{i-1}^\nu & i > 0, \eta_{i-1} = \nu \\ 0 & i = 0 \end{cases}$$

for $0 \leq i \leq k$. The priority function $\Omega : V \rightarrow [\text{ad}(k)]$ then is defined by $\Omega(u, i) = \text{ad}(i)$ and $\Omega(U_0, \dots, U_k) = 0$.

Remark 4.2. In [4], an alternative priority function $\Omega' : V \rightarrow [2k + 1]$ with

$$\Omega'(u, i) = \begin{cases} 2i & \text{if } \eta_i = \text{GFP} \\ 2i + 1 & \text{if } \eta_i = \text{LFP} \end{cases}$$

and $\Omega'(U_0, \dots, U_k) = 0$ is used. Since $\text{ad}(i)$ is even if and only if η_i is even, and moreover $\text{ad}(i) \leq \text{ad}(j)$ for $i \leq j$, and $i < j$ whenever $\text{ad}(i) < \text{ad}(j)$, it is easy to see that Ω and Ω' in fact assign identical parities to all plays. In the following, we will use the more economic parity function Ω so that fixpoint games have only $d := \text{ad}(k) \leq k$ priorities.

We import the associated characterization theorem [4, Theorem 4.8]:

Theorem 4.3 ([4]). *We have $u \sqsubseteq \llbracket X_i \rrbracket_f$ if and only if Eloise wins the node (u, i) in the fixpoint game for the given system f of equations.*

Remark 4.4. While this shows that parity game solving can be used to solve equation systems, the size of fixpoint games is exponential in $|B_L|$, so they do not directly yield a quasipolynomial algorithm for solving equation systems.

Next we define our notion of history-freeness for systems of fixpoint equations.

Definition 4.5 (History-free witness). A *history-free witness* for $u \sqsubseteq \llbracket X_i \rrbracket_f$ is an even labelled graph (W, R) with labels from $[d]$ such that $W \subseteq B_L \times [d]$, $(u, i) \in W$, and for all $(v, p) \in W$, we have $v \sqsubseteq f_p(U_0, \dots, U_k)$ where $U_j = \bigsqcup \pi_1[R_{\text{ad}(j)}(v, p)]$ for $0 \leq j \leq k$, noting that $R_{\text{ad}(j)}(v, p) \subseteq W$ so that $\pi_1[R_{\text{ad}(j)}(v, p)] \subseteq B_L$ and $U_j \in L$.

In analogy to history-free strategies for parity games, history-free witnesses assign tuples $(R_1(v, p), \dots, R_d(v, p))$ of sets $R_j(v, p) \subseteq W$ to pairs $(v, p) \in W$ without relying on a history of previously visited pairs. We have $|W| \leq (d+1)|B_L|$ and $|R| \leq (d+1)|W|^2$, that is, the size of history-free witnesses is polynomial in $|B_L|$. Crucially, history-free witnesses always exist:

Lemma 4.6. *For all $u \in B_L$ and $i \in [k]$, we have*

$$u \sqsubseteq \llbracket X_i \rrbracket_f \text{ if and only if there is a history-free witness for } u \sqsubseteq \llbracket X_i \rrbracket_f.$$

Proof. In one direction, we have $u \sqsubseteq \llbracket X_i \rrbracket_f$ so that Eloise wins the node (u, i) in the according fixpoint game by Lemma 4.3. Let s be a corresponding *history-free* winning strategy (such strategies always exists, see e.g. [26]). We inductively construct a witness for $u \sqsubseteq \llbracket X_i \rrbracket_f$, starting at (u, i) . When at $(v, p) \in B_L \times [k]$ with $s(v, p) = (U_0, \dots, U_k)$, we put $R_i(v, p) = \bigcup_{j|\text{ad}(j)=i} (U_j \times \{j\})$ for $0 \leq i \leq d$ and hence have $\text{ad}(j) = i$ for all $((v, p), i, (u, j)) \in R$. Since s is a winning strategy, the resulting graph (W, R) is a history-free witness for $u \sqsubseteq \llbracket X_i \rrbracket_f$ by construction; in particular, (W, R) is even. For the converse direction, the witness for $u \sqsubseteq \llbracket X_i \rrbracket_f$ directly yields a winning Eloise-strategy for the node (u, i) in the associated fixpoint game. This implies $u \sqsubseteq \llbracket X_i \rrbracket_f$ by Lemma 4.3. $\square$

5 Solving Equation Systems using Universal Graphs

We go on to prove our main result. To this end, we fix a system f of fixpoint equations $f_i : L^{k+1} \rightarrow L$, $0 \leq i \leq k$, and put $n := |B_L|$ and $d := \text{ad}(k)$ for the remainder of the paper.

Definition 5.1 (Universal graphs [16, 17]). Let $G = (W, R)$ and $G' = (W', R')$ be labelled graphs with labels from $[d]$. A *homomorphism of labelled graphs* from G to G' is a function $\Phi : W \rightarrow W'$ such that for all $(v, p, w) \in R$, we have $(\Phi(v), p, \Phi(w)) \in R'$. An $(n, d+1)$ -*universal graph* S is an even graph with labels from $[d]$ such that for all even graphs G with labels from $[d]$ and with $|G| \leq n$, there is a homomorphism from G to S .

We fix an $(n(d+1), (d+1))$ -universal graph $S = (Z, K)$, noting that there are $(n(d+1), (d+1))$ -universal graphs (obtained from universal trees) of size quasipolynomial in n and d [17]. We now combine the system f with the universal graph S to turn the parity conditions associated to general systems of fixpoint equations into a safety condition, associated to a single greatest fixpoint equation.

Definition 5.2 (Chained-product fixpoint). We define a function

$$g: \mathcal{P}(B_L \times [k] \times Z) \rightarrow \mathcal{P}(B_L \times [k] \times Z)$$

$$U \mapsto \{(v, p, q) \in B_L \times [k] \times Z \mid v \sqsubseteq f_p(P_0^{U,q}, \dots, P_k^{U,q})\}$$

where

$$P_i^{U,q} = \bigsqcup \{u \in B_L \mid \exists s \in K_{\text{ad}(i)}(q). (u, i, s) \in U\}.$$

We refer to $Y_0 =_{\text{GFP}} g(Y_0)$ as the *chained-product fixpoint (equation)* of f and S .

We now show our central result: apart from the annotation with states from the universal graph, the chained-product fixpoint g is the solution of the system f .

Theorem 5.3. *For all $u \in B_L$ and $0 \leq i \leq k$, we have*

$$u \sqsubseteq \llbracket X_i \rrbracket_f \text{ if and only if there is } q \in Z \text{ such that } (u, i, q) \in \llbracket Y_0 \rrbracket_g.$$

Proof. For the forward direction, let $u \sqsubseteq \llbracket X_i \rrbracket_f$. By Lemma 4.6, there is a history-free witness $G = (W, R)$ for $u \sqsubseteq \llbracket X_i \rrbracket_f$. Since S is a $(n(d+1), d+1)$ -universal graph and since G is a witness and hence an even labelled graph of suitable size $|G| \leq n(d+1)$, there is a graph homomorphism Φ from G to S . Starting at $(u, i, \Phi(u, i), 0)$, we inductively construct a witness for containment of $(u, i, \Phi(u, i))$ in $\llbracket Y_0 \rrbracket_g$. When at $(v_1, p_1, \Phi(v_1, p_1), 0)$ with $(v_1, p_1) \in W$, we put

$$R'_0(v_1, p_1, \Phi(v_1, p_1), 0) = \{(v_2, p_2, \Phi(v_2, p_2), 0) \in B_L \times [d] \times Z \times [0] \mid$$

$$(v_2, p_2) \in R_{\text{ad}(p_2)}(v_1, p_1), \Phi(v_2, p_2) \in K_{\text{ad}(p_2)}(\Phi(v_1, p_1))\}$$

and continue the inductive construction with all these $(v_2, p_2, \Phi(v_2, p_2), 0)$, having $(v_2, p_2) \in W$. The resulting structure $G' = (W', R')$ indeed is a witness for containment of (u, i, q) in $\llbracket Y_0 \rrbracket_g$: G' is even by construction. Moreover, we need to show that for $(v_1, p_1, \Phi(v_1, p_1), 0) \in W'$, we have $(v_1, p_1, \Phi(v_1, p_1), 0) \in g(\pi_1[R'_0(v_1, p_1, \Phi(v_1, p_1), 0)])$, i.e. $v_1 \sqsubseteq f_{p_1}(P_0^{U, \Phi(v_1, p_1)}, \dots, P_k^{U, \Phi(v_1, p_1)})$ where $U = \pi_1[R'_0(v_1, p_1, \Phi(v_1, p_1), 0)]$. Since G is a witness and $(v_1, p_1) \in W$ by construction of W' , we have $v_1 \sqsubseteq f_{p_1}(U_0, \dots, U_k)$ where $U_j = \bigsqcup (\pi_j[R_{\text{ad}(i)}(v_1, p_1)])$. By monotonicity of f_{p_1} , it thus suffices to show that $U_j \sqsubseteq P_j^{U, \Phi(v_1, p_1)}$ for $0 \leq j \leq k$; by definition of $P_j^{U, \Phi(v_1, p_1)}$ this follows if

$$\pi_1[R_{\text{ad}(j)}(v_1, p_1)] \subseteq \{u \in B_L \mid \exists s \in K_{\text{ad}(j)}(\Phi(v_1, p_1)). (u, j, s) \in W\},$$

where $W = \pi_1[R'_0(v_1, p_1, q_1, 0)]$. So let $w \in B_L$ such that $w \in \pi_1[R_{\text{ad}(j)}(v_1, p_1)]$. Since R is a witness that is constructed as in the proof of Lemma 4.6, we have $i = \text{ad}(i')$ for all $((v', p'), i, (w', i')) \in R$. Thus $(w, j) \in R_{\text{ad}(j)}(v_1, p_1)$ for some j such that $\text{ad}(j) = i$, that is, $((v_1, p_1), \text{ad}(j), (w, j)) \in R$, hence $(\Phi(v_1, p_1), \text{ad}(j), \Phi(w, j)) \in K$ because Φ is a graph homomorphism. By definition of R'_0 we have $(w, j, \Phi(w, j), 0) \in R'_0(v_1, p_1, \Phi(v_1, p_1), 0)$ so that $(w, j, \Phi(w, j)) \in \pi_1[R'_0(v_1, p_1, \Phi(v_1, p_1), 0)]$. We are done since $\Phi(w, j) \in K_{\text{ad}(j)}(\Phi(v_1, p_1))$.

For the converse implication, let $(u_0, p_0, q_0) \in \llbracket Y_0 \rrbracket_g$ for some $q_0 \in Z$. Let $G = (W, R)$ be a history-free witness for this fact. By Lemma 4.3, it suffices to provide a strategy in the fixpoint game for the system f with which Eloise wins the node (u_0, p_0) . We inductively construct a *history-dependent* strategy s as follows: For $i \geq 0$, we abbreviate $U_i = R_0(u_i, p_i, q_i, 0)$. We put $s(u_0, p_0) = (P_0^{U_0, q_0}, \dots, P_k^{U_0, q_0})$. For the inductive step, let

$$\tau = (u_0, p_0), (P_0^{U_0, q_0}, \dots, P_k^{U_0, q_0}), \dots, (P_0^{U_{n-1}, q_{n-1}}, \dots, P_k^{U_{n-1}, q_{n-1}}), (u_n, p_n)$$

be a partial play of the fixpoint game that follows the strategy that has been constructed so far. Then we have an R -path $(u_0, p_0, q_0, 0), (u_1, p_1, q_1, 0), \dots, (u_n, p_n, q_n, 0)$, where, for $0 \leq i < n$, we have $(q_i, p_{i+1}, q_{i+1}) \in K$ since $u_{i+1} \sqsubseteq P_{p_{i+1}}^{U_i, q_i}$ by the inductive construction. Put $s(\tau) = (P_0^{U_n, q_n}, \dots, P_k^{U_n, q_n})$. Since G is a witness, the strategy uses only moves that are available to Eloise (i.e. ones with $u_n \sqsubseteq f_{p_n}(s(\tau))$). Also, s is a winning strategy as can be seen by looking at the K -paths that are induced by complete plays τ that follow s , as described (for partial plays) above. Since S is a universal graph and hence even, every such K -path is even and the sequence of priorities in τ is just the sequence of priorities of one of these K -paths. $\square$

Remark 5.4. Since the set $\llbracket Y_0 \rrbracket_g$ is the greatest fixpoint of g , it can be computed by simple approximation from above, that is, as $g^m(B_L \times [k] \times Z)$ where $m = |B_L \times [k] \times Z|$. However, each iteration of the function g may require up to $|Z|$ evaluations of an equation. In the next section, we will show how this additional iteration factor in the computation of $\llbracket Y_0 \rrbracket_g$ can be avoided.

6 A Progress Measure Algorithm

We next introduce a lifting algorithm that computes the set $\llbracket Y_0 \rrbracket_g$ efficiently, following the paradigm of the progress measure approach for parity games (e.g. [33, 34]). Our progress measures will map pairs $(u, i) \in B_L \times [k]$ to nodes in a universal graph that is equipped with a simulation order, that is, a total order that is suitable for measuring progress.

Definition 6.1 (Simulation order). For natural numbers i, i' , we put $i \succeq i'$ if and only if either i is even and $i = i'$, or both i and i' are odd and $i \geq i'$. A total order $\leq$ on Z is a *simulation order* if for all $q, q' \in Z$,

$$q \leq q' \text{ implies that for all } 0 \leq i \leq k \text{ and } s \in K_i(q), \text{ there are} \\ i' \succeq i \text{ and } s' \in K_{i'}(q') \text{ such that } s \leq s'.$$

Lemma 6.2. *There is an $(n(d+1), d+1)$ -universal graph (Z, K) of size quasipolynomial in n and d , and over which a simulation order $\leq$ exists.*

Proof (Sketch). It has been shown [17, Theorem 2.2] (originally, in different terminology, [34]) that there are (l, h) -universal trees (a concept similar to, but

slightly more concrete than universal graphs) with set of leaves T such that $|T| \leq 2l \binom{l+h+1}{h}$. Leaves in universal trees are identified by *navigation paths*, that is, sequences of branching directions, so that the leaves are linearly ordered by the lexicographic order $\leq$ on navigation paths (which orders leaves from the left to the right). As described in [16], one can obtain a universal graph (T, K) over T in which transitions $(q, i, q') \in K$ for odd i (the crucial case) move to the left, that is, q' is a leaf that is to the left of q in the universal tree (so that $q' < q$), ensuring universality. As it turns out, the lexicographic ordering on T is a simulation order. Adapting this construction to our setting, we put $l = n(d+1)$ and $h = d+1$ and obtain a $(n(d+1), d+1)$ -universal graph (along with a simulation order $\leq$) of size at most $2n(d+1) \binom{\log(n(d+1))+d+2}{d+1}$ which is quasipolynomial in n and d . $\square$

We fix an $(n(d+1), d+1)$ -universal graph (Z, K) and a simulation order $\leq$ on Z for the remainder of the paper (these exist by the above lemma).

Definition 6.3 (Progress measure, lifting function). We let $q_{\min} \in Z$ denote the least node w.r.t. $\leq$ and fix a distinguished top element $\star \notin Z$, and extend $\geq$ to $Z \cup \{\star\}$ by putting $\star \geq q$ for all $q \in Z$. A *measure* is a map $\mu: B_L \times [k] \rightarrow Z \cup \{\star\}$, i.e. assigns nodes in the universal graph or $\star$ to pairs $(v, p) \in B_L \times [k]$. A measure μ is a *progress measure* if whenever $\mu(v, p) \neq \star$, then $v \sqsubseteq f_p(U_0^{\mu, q}, \dots, U_k^{\mu, q})$ where $q = \mu(v, p)$ and

$$U_i^{\mu, q} = \bigsqcup \{u \in B_L \mid \exists s \in K_{\text{ad}(i)}(q). \mu(u, i) \leq s\}.$$

We define a function $\text{Lift} : (B_L \times [k] \rightarrow Z \cup \{\star\}) \rightarrow (B_L \times [k] \rightarrow Z \cup \{\star\})$ on measures by

$$(\text{Lift}(\mu))(v, p) = \min\{q \in Z \mid v \sqsubseteq f_p(U_0^{\mu, q}, \dots, U_k^{\mu, q})\}$$

where $\min(Z')$ denotes the least element of Z' w.r.t. $\leq$, for $\emptyset \neq Z' \subseteq Z$; also we put $\min(\emptyset) = \star$.

The lifting algorithm then starts with the least measure $\mathbf{m}_{\min}$ that maps all pairs $(v, p) \in B_L \times [k]$ to the minimal node (i.e. $\mathbf{m}_{\min}(v, p) = q_{\min}$) and repeatedly updates the current measure using Lift until the measure stabilizes.

Lifting algorithm

- (1) Initialize: Put $\mu := \mathbf{m}_{\min}$.
- (2) If $\text{Lift}(\mu) \neq \mu$, then put $\mu := \text{Lift}(\mu)$ and go to 2. Otherwise go to 3.
- (3) Return the set $\mathbb{E} = \{(v, p) \in B_L \times [k] \mid \mu(v, p) \neq \star\}$.

Lemma 6.4 (Correctness). *For all $v \in B_L$ and $0 \leq p \leq k$, we have*

$$(v, p) \in \mathbb{E} \text{ if and only if } v \in \llbracket X_p \rrbracket_f.$$

Proof (Sketch). Let μ denote the progress measure that the algorithm computes. For one direction of the proof, let $(v, p) \in \mathbb{E}$. By Lemma 4.6 it suffices to construct a witness for $v \in \llbracket X_p \rrbracket_f$. We extract such a witness $(\mathbb{E}, R)$ from the progress measure μ , relying on the properties of the simulation order $\leq$ that is used to measure the progress of μ to ensure that any infinite sequence of measures that μ assigns to some R -path induces an infinite (and hence even) path in the employed universal graph. This shows that $(\mathbb{E}, R)$ indeed is an even graph and hence a witness. For the converse direction, let $v \in \llbracket X_p \rrbracket_f$ so that there is, by Theorem 5.3, some $q \in Z$ such that $(v, p, q) \in \llbracket Y_0 \rrbracket_g$. For (u, i) such that there is $q' \in Z$ such that $(u, i, q') \in \llbracket Y_0 \rrbracket_g$, let $q_{(u,i)} \in Z$ denote the minimal such node w.r.t. $\leq$. It now suffices that $\mu(u, i) \leq q_{(u,i)}$ for all such (u, i) , which is shown by induction on the number of iterations of the lifting algorithm. $\square$

Corollary 6.5. *Solutions of systems of fixpoint equations can be computed with quasipolynomially many evaluations of equations.*

Proof. Given an $(n(d+1), d+1)$ -universal graph (Z, K) and a simulation order on Z , the lifting algorithm terminates and returns the solution of f after at most $n(d+1) \cdot |Z|$ many iterations. This is the case since each iteration (except the final iteration) increases the measure for at least one of the $n(d+1)$ nodes and the measure of each node can be increased at most $|Z|$ times. Using the universal graph and the simulation order from the proof of Lemma 6.2, we have $|Z| \leq 2n(d+1)^{\binom{\log(n(d+1))+d+2}{d+1}}$ so that the algorithm terminates after at most $2(n(d+1))^2 \binom{\log(n(d+1))+d+2}{d+1} \in \mathcal{O}((n(d+1))^{\log(d+1)})$ iterations of the function Lift. Each iteration can be implemented to run with at most $n(d+1)$ evaluations of an equation. $\square$

Corollary 6.6. *The number of function calls required for the solution of systems of fixpoint equations with $d \leq \log n$ is bounded by a polynomial in n and d .*

Proof. Following the insight of Theorem 2.8 in [9], Theorem 2.2. in [17] implies that if $d < \log n$, then there is an $(n(d+1), d+1)$ -universal tree of size polynomial in n and d . In the same way as in the proof of Lemma 6.2, one obtains a universal graph of polynomial size and a simulation order on it. $\square$

Example 6.7. Applying Corollary 6.5 and Corollary 6.6 to Example 3.2, we obtain the following results:

- (1) The model checking problems for the energy μ -calculus and finite latticed μ -calculi are in QP. For energy parity games with sufficient upper bound b on energy level accumulations, we obtain a progress measure algorithm that terminates after a number of iterations that is quasipolynomial in b .
- (2) Under mild assumptions on the modalities (see [30]), the model checking problem for the coalgebraic μ -calculus is in QP; in particular, this yields QP model checking algorithms for the graded μ -calculus and the two-valued probabilistic μ -calculus (equivalently: QP progress measure algorithms for solving graded and two-valued probabilistic parity games).

(3) Under mild assumptions on the modalities (see [31]), we obtain a novel upper bound $2^{\mathcal{O}(nd \log n)}$ for the satisfiability problems of coalgebraic μ -calculi, in particular including the monotone μ -calculus, the alternating-time μ -calculus, the graded μ -calculus and the (two-valued) probabilistic μ -calculus, even when the latter two are extended with (monotone) polynomial inequalities. This improves on the best previous bounds in all cases.

7 Conclusion

We have shown how to use universal graphs to compute solutions of systems of fixpoint equations $X_i = \eta_i \cdot f_i(X_0, \dots, X_k)$ (with the η_i marking least or greatest fixpoints) that use functions $f_i : L^{k+1} \rightarrow L$ (over a finite lattice L with basis B_L) and involve up to $k + 1$ -fold nesting of fixpoints. Our progress measure algorithm needs quasipolynomially many evaluations of equations, and runs in time $\mathcal{O}(q \cdot t(f))$, where q is a quasipolynomial in $|B_L|$ and the alternation depth of the equation system, and where $t(f)$ is an upper bound on the time it takes to compute f_i for all i .

As a consequence of our results, the upper time bounds for the evaluation of various general parity conditions improve. Example domains beyond solving parity games to which our algorithm can be instantiated comprise model checking for latticed μ -calculi and solving latticed parity games [7, 36], solving energy parity games and model checking for the energy μ -calculus [2, 10], and model checking and satisfiability checking for the coalgebraic μ -calculus [14]. The resulting model checking algorithms for latticed μ -calculi and the energy μ -calculus run in time quasipolynomial in the provided basis of the respective lattice. In terms of concrete instances of the coalgebraic μ -calculus, we obtain, e.g., quasipolynomial-time model checking for the graded [38] and the probabilistic μ -calculus [14, 41] as new results (corresponding results for, e.g., the alternating-time μ -calculus [1] and the monotone μ -calculus [22] follow as well but have already been obtained in our previous work [30]), as well as improved upper bounds for satisfiability checking in the graded μ -calculus, the probabilistic μ -calculus, the monotone μ -calculus, and the alternating-time μ -calculus. We foresee further applications, e.g. in the computation of fair bisimulations and fair equivalence [32, 37] beyond relational systems, e.g. for probabilistic systems.

As in the case of parity games, a natural open question that remains is whether solutions of fixpoint equations can be computed in polynomial time (which would of course imply that parity games can be solved in polynomial time). A more immediate perspective for further investigation is to generalize the recent quasipolynomial variant [46] of Zielonka’s algorithm [54] for solving parity games to solving systems of fixpoint equations, with a view to improving efficiency in practice.

References

1. Alur, R., Henzinger, T., Kupferman, O.: Alternating-time temporal logic. J. ACM **49**, 672–713 (2002), <https://doi.org/10.1145/585265.585270>

2. Amram, G., Maoz, S., Pistiner, O., Ringert, J.O.: Energy mu-calculus: Symbolic fixed-point algorithms for omega-regular energy games. CoRR **abs/2005.00641** (2020), <https://arxiv.org/abs/2005.00641>
3. Arnold, A., Niwinski, D., Parys, P.: A quasi-polynomial black-box algorithm for fixed point evaluation. In: Computer Science Logic, CSL 2021. LIPIcs, vol. 183, pp. 9:1–9:23. Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2021), <https://doi.org/10.4230/LIPIcs.CSL.2021.9>
4. Baldan, P., König, B., Mika-Michalski, C., Padoan, T.: Fixpoint games on continuous lattices. In: Principles of Programming Languages, POPL 2021. Proceedings of the ACM on Programming Languages, vol. 3, pp. 26:1–26:29. ACM (2019), <https://doi.org/10.1145/3290339>
5. Bodlaender, H., Dinneen, M., Khoussainov, B.: On game-theoretic models of networks. In: Algorithms and Computation, ISAAC 2001. LNCS, vol. 2223, pp. 550–561. Springer (2001), https://doi.org/10.1007/3-540-45678-3_47
6. Boker, U., Lehtinen, K.: On the way to alternating weak automata. In: Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2018. LIPIcs, vol. 122, pp. 21:1–21:22. Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2018), <https://doi.org/10.4230/LIPIcs.FSTTCS.2018.21>
7. Bruns, G., Godefroid, P.: Model checking with multi-valued logics. In: Automata, Languages and Programming, ICALP 2004. LNCS, vol. 3142, pp. 281–293. Springer (2004), https://doi.org/10.1007/978-3-540-27836-8_26
8. Bruse, F., Falk, M., Lange, M.: The fixpoint-iteration algorithm for parity games. In: Games, Automata, Logics and Formal Verification, GandALF 2014. EPTCS, vol. 161, pp. 116–130. Open Publishing Association (2014), <https://doi.org/10.4204/EPTCS.161.12>
9. Calude, C., Jain, S., Khoussainov, B., Li, W., Stephan, F.: Deciding parity games in quasipolynomial time. In: Theory of Computing, STOC 2017. pp. 252–263. ACM (2017), <https://doi.org/10.1145/3055399.3055409>
10. Chatterjee, K., Doyen, L.: Energy parity games. Theor. Comput. Sci. **458**, 49–60 (2012). <https://doi.org/10.1016/j.tcs.2012.07.038>
11. Chatterjee, K., Dvorák, W., Henzinger, M., Svozil, A.: Quasipolynomial set-based symbolic algorithms for parity games. In: Logic for Programming, Artificial Intelligence and Reasoning, LPAR 2018. EPIc, vol. 57, pp. 233–253. EasyChair (2018), <https://doi.org/10.29007/5z5k>
12. Chatterjee, K., Jurdziński, M., Henzinger, T.A.: Simple stochastic parity games. In: Baaz, M., Makowsky, J.A. (eds.) Computer Science Logic. pp. 100–113. Springer (2003)
13. Chellas, B.F.: Modal Logic. Cambridge University Press (1980)
14. Cîrstea, C., Kupke, C., Pattinson, D.: EXPTIME tableaux for the coalgebraic μ -calculus. Log. Meth. Comput. Sci. **7** (2011), [https://doi.org/10.2168/LMCS-7\(3:3\)2011](https://doi.org/10.2168/LMCS-7(3:3)2011)
15. Cîrstea, C., Kurz, A., Pattinson, D., Schröder, L., Venema, Y.: Modal logics are coalgebraic. Comput. J. **54**, 31–41 (2011)
16. Colcombet, T., Fijalkow, N.: Universal graphs and good for games automata: New tools for infinite duration games. In: Foundations of Software Science and Computation Structures, FOSSACS 2019. LNCS, vol. 11425, pp. 1–26. Springer (2019), https://doi.org/10.1007/978-3-030-17127-8_1
17. Czerwinski, W., Daviaud, L., Fijalkow, N., Jurdzinski, M., Lazic, R., Parys, P.: Universal trees grow inside separating automata: Quasi-polynomial lower bounds for parity games. In: Symposium on Discrete Algorithms, SODA 2019. pp. 2333–2349. SIAM (2019), <https://doi.org/10.1137/1.9781611975482.142>

18. D’Agostino, G., Visser, A.: Finality regained: A coalgebraic study of Scott-sets and multisets. *Arch. Math. Logic* **41**, 267–298 (2002)
19. Daviaud, L., Jurdzinski, M., Lazic, R.: A pseudo-quasi-polynomial algorithm for mean-payoff parity games. In: *Logic in Computer Science, LICS 2018*. pp. 325–334. ACM (2018), <https://doi.org/10.1145/3209108.3209162>
20. Dawar, A., Grädel, E.: The descriptive complexity of parity games. In: *Computer Science Logic, CSL 2008. LNCS*, vol. 5213, pp. 354–368. Springer (2008), https://doi.org/10.1007/978-3-540-87531-4_26
21. Emerson, E.A., Jutla, C., Sistla, A.P.: On model checking for the μ -calculus and its fragments. *Theor. Comput. Sci.* **258**, 491–522 (2001), [https://doi.org/10.1016/S0304-3975\(00\)00034-7](https://doi.org/10.1016/S0304-3975(00)00034-7)
22. Enqvist, S., Seifan, F., Venema, Y.: Monadic second-order logic and bisimulation invariance for coalgebras. In: *Logic in Computer Science, LICS 2015*. pp. 353–365. IEEE (2015), <https://doi.org/10.1109/LICS.2015.41>
23. Fagin, R., Halpern, J.: Reasoning about knowledge and probability. *J. ACM* **41**(2), 340–367 (1994). <https://doi.org/10.1145/174652.174658>, <https://doi.org/10.1145/174652.174658>
24. Fearnley, J., Jain, S., de Keijzer, B., Schewe, S., Stephan, F., Wojtczak, D.: An ordered approach to solving parity games in quasi-polynomial time and quasi-linear space. *STTT* **21**(3), 325–349 (2019), <https://doi.org/10.1007/s10009-019-00509-3>
25. Friedmann, O., Lange, M.: Deciding the unguarded modal μ -calculus. *J. Appl. Non-Classical Log.* **23**, 353–371 (2013), <https://doi.org/10.1080/11663081.2013.861181>
26. Grädel, E., Thomas, W., Wilke, T. (eds.): *Automata, Logics, and Infinite Games: A Guide to Current Research*, LNCS, vol. 2500. Springer (2002), <https://doi.org/10.1007/3-540-36387-4>
27. Hahn, E.M., Schewe, S., Turrini, A., Zhang, L.: A simple algorithm for solving qualitative probabilistic parity games. In: Chaudhuri, S., Farzan, A. (eds.) *Computer Aided Verification*. pp. 291–311. Springer (2016)
28. Hasuo, I., Shimizu, S., Cirstea, C.: Lattice-theoretic progress measures and coalgebraic model checking. In: *Principles of Programming Languages, POPL 2016*. pp. 718–732. ACM (2016), <https://doi.org/10.1145/2837614.2837673>
29. Hausmann, D., Schröder, L.: Computing nested fixpoints in quasipolynomial time. *CoRR* **abs/1907.07020** (2019), <http://arxiv.org/abs/1907.07020>
30. Hausmann, D., Schröder, L.: Game-based local model checking for the coalgebraic μ -calculus. In: *Concurrency Theory, CONCUR 2019. LIPIcs*, vol. 140, pp. 35:1–35:16. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik (2019), <https://doi.org/10.4230/LIPIcs.CONCUR.2019.35>
31. Hausmann, D., Schröder, L.: Optimal satisfiability checking for arithmetic μ -calculi. In: *Foundations of Software Science and Computation Structures, FOSACS 2019. LNCS*, vol. 11425, pp. 277–294. Springer (2019), https://doi.org/10.1007/978-3-030-17127-8_16
32. Henzinger, T., Rajamani, S.: Fair bisimulation. In: *Tools and Algorithms for Construction and Analysis of Systems, TACAS 2000. LNCS*, vol. 1785, pp. 299–314. Springer (2000), https://doi.org/10.1007/3-540-46419-0_21
33. Jurdzinski, M.: Small progress measures for solving parity games. In: *Symposium on Theoretical Aspects of Computer Science, STACS 2000. LNCS*, vol. 1770, pp. 290–301. Springer (2000), https://doi.org/10.1007/3-540-46541-3_24
34. Jurdzinski, M., Lazic, R.: Succinct progress measures for solving parity games. In: *Logic in Computer Science, LICS 2017*. pp. 1–9. IEEE Computer Society (2017), <https://doi.org/10.1109/LICS.2017.8005092>

35. Kozen, D.: Results on the propositional μ -calculus. *Theor. Comput. Sci.* **27**, 333–354 (1983), [https://doi.org/10.1016/0304-3975\(82\)90125-6](https://doi.org/10.1016/0304-3975(82)90125-6)
36. Kupferman, O., Lustig, Y.: Latticed simulation relations and games. In: *Automated Technology for Verification and Analysis, ATVA 2007*. LNCS, vol. 4762, pp. 316–330. Springer (2007), https://doi.org/10.1007/978-3-540-75596-8_23
37. Kupferman, O., Piterman, N., Vardi, M.: Fair equivalence relations. In: *Verification: Theory and Practice*. LNCS, vol. 2772, pp. 702–732. Springer (2003), https://doi.org/10.1007/978-3-540-39910-0_30
38. Kupferman, O., Sattler, U., Vardi, M.: The complexity of the graded μ -calculus. In: *Automated Deduction, CADE 2002*. LNCS, vol. 2392, pp. 423–437. Springer (2002), https://doi.org/10.1007/3-540-45620-1_34
39. Kupke, C., Pattinson, D., Schröder, L.: Reasoning with global assumptions in arithmetic modal logics. In: *Fundamentals of Computation Theory, FCT 2015*. LNCS, vol. 9210, pp. 367–380. Springer (2015)
40. Lehtinen, K.: A modal μ perspective on solving parity games in quasi-polynomial time. In: *Logic in Computer Science, LICS 2018*. pp. 639–648. ACM (2018), <https://doi.org/10.1145/3209108.3209115>
41. Liu, W., Song, L., Wang, J., Zhang, L.: A simple probabilistic extension of modal mu-calculus. In: *International Joint Conference on Artificial Intelligence, IJCAI 2015*. pp. 882–888. AAAI Press (2015), <http://ijcai.org/proceedings/2015>
42. Long, D.E., Browne, A., Clarke, E.M., Jha, S., Marrero, W.R.: An improved algorithm for the evaluation of fixpoint expressions. In: *Computer Aided Verification, CAV 1994*. LNCS, vol. 818, pp. 338–350. Springer (1994), https://doi.org/10.1007/3-540-58179-0_66
43. Mio, M.: On the equivalence of game and denotational semantics for the probabilistic μ -calculus. *Log. Methods Comput. Sci.* **8** (2012), [https://doi.org/10.2168/LMCS-8\(2:7\)2012](https://doi.org/10.2168/LMCS-8(2:7)2012)
44. Niwinski, D.: On fixed-point clones (extended abstract). In: *Automata, Languages and Programming, ICALP 1986*. LNCS, vol. 226, pp. 464–473. Springer (1986)
45. Parikh, R.: The logic of games and its applications. *Ann. Discr. Math.* **24**, 111–140 (1985), [https://doi.org/10.1016/S0304-0208\(08\)73078-0](https://doi.org/10.1016/S0304-0208(08)73078-0)
46. Parys, P.: Parity games: Zielonka’s algorithm in quasi-polynomial time. In: *Mathematical Foundations of Computer Science, MFCS 2019*. LIPIcs, vol. 138, pp. 10:1–10:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2019), <https://doi.org/10.4230/LIPIcs.MFCS.2019.10>
47. Pauly, M.: A modal logic for coalitional power in games. *J. Logic Comput.* **12**, 149–166 (2002)
48. Peleg, D.: Concurrent dynamic logic. *J. ACM* **34**, 450–479 (1987), <https://doi.org/10.1145/23005.23008>
49. Rutten, J.: Universal coalgebra: A theory of systems. *Theor. Comput. Sci.* **249**, 3–80 (2000), [https://doi.org/10.1016/S0304-3975\(00\)00056-6](https://doi.org/10.1016/S0304-3975(00)00056-6)
50. Schröder, L., Pattinson, D.: Strong completeness of coalgebraic modal logics. In: *Theoretical Aspects of Computer Science, STACS 09*. pp. 673–684. Schloss Dagstuhl – Leibniz-Zentrum für Informatik; Dagstuhl, Germany (2009)
51. Schröder, L., Venema, Y.: Completeness of flat coalgebraic fixpoint logics. *ACM Trans. Comput. Log.* **19**(1), 4:1–4:34 (2018)
52. Seidl, H.: Fast and Simple Nested Fixpoints. *Inf. Process. Lett.* **59**, 303–308 (1996), [https://doi.org/10.1016/0020-0190\(96\)00130-5](https://doi.org/10.1016/0020-0190(96)00130-5)
53. Venema, Y.: Lectures on the modal μ -calculus. Lecture notes, Institute for Logic, Language and Computation, Universiteit van Amsterdam (2008), <https://staff.fnwi.uva.nl/y.venema/teaching/ml/notes/20201212-mu.pdf>

54. Zielonka, W.: Infinite games on finitely coloured graphs with applications to automata on infinite trees. Theor. Comput. Sci. **200**(1-2), 135–183 (1998), [https://doi.org/10.1016/S0304-3975\(98\)00009-7](https://doi.org/10.1016/S0304-3975(98)00009-7)

A Appendix

A.1 Omitted Proofs

Full proof of Lemma 6.2.

Proof. It has been shown in [17], Theorem 2.2. that there are (l, h) -universal trees with set of leaves T such that $|T| \leq 2l \binom{\log l + h + 1}{h}$; following [16], we show how to obtain a universal graph (Z, K) of the same size. The leaves in universal trees can be identified by *navigation paths*, that is, tuples $(m_k, m_{k-2}, \dots, m_1)$ of natural numbers that describe the position of the leaf within the tree. We put $Z = T$ and define the relation K by putting

$$K_i(m_k, m_{k-2}, \dots, m_1) = \begin{cases} \{(m_k, m_{k-2}, \dots, m_i - 1) \mid m_i > 0\} & \text{if } i \text{ is odd} \\ \{(m_k, m_{k-2}, \dots, m_{i-1})\} & \text{if } i \text{ is even} \end{cases}$$

for $0 \leq i \leq k$, $(m_k, m_{k-2}, \dots, m_i) \in Z$. Here we use $(m_k, m_{k-2}, \dots, m_p)$ (for odd $p \geq 1$) to denote the rightmost leaf that is reachable from position $(m_k, m_{k-2}, \dots, m_p)$ in the universal tree. By (l, h) -universality of the used universal tree, (K, Z) is a (l, h) -universal graph. Let $\leq$ be the lexicographic order on the leaves of the universal tree; then $\leq$ orders the leaves linearly from left to right. We show that $\leq$ is a simulation order on (Z, K) . Let $q = (m_k, m_{k-2}, \dots, m_1) \in Z$ and $q' = (m'_k, m'_{k-2}, \dots, m'_1) \in Z$ such that $q \leq q'$. Furthermore let $0 \leq i \leq k$ and $s \in K_i(q)$. We have to show that there are $i' \succeq i$ and $s' \in K_{i'}(q')$ such that $s \leq s'$. If $q = q'$, then we are done. If $q \neq q'$ and there is $s' \in K_i(q')$, then $s' = (m'_k, m'_{k-2}, \dots, m'_i - 1)$ (if i is odd) or $s' = (m'_k, m'_{k-2}, \dots, m'_{i-1})$ (if i is even). In the first case, we have $s = (m_k, m_{k-2}, \dots, m_i - 1)$ and in the second case, we have $s = (m_k, m_{k-2}, \dots, m_{i-1})$. In both cases we have $i \succeq i$ and $s \leq s'$ follows from $q \leq q'$. If $q \neq q'$ and $K_i(q') = \emptyset$, then i is odd (and $m'_i = 0$) and there is an odd number $i' > i$ such that $m_{i'} < m'_{i'}$, since $q < q'$. Since $m'_{i'} > m_{i'} \geq 0$, there is $s' = (m_k, m_{k-2}, \dots, m_{i'} - 1) \in K_{i'}(q') \neq \emptyset$. It remains to show $i' \succeq i$ and $s = (m_k, m_{k-2}, \dots, m_{i-1}) \leq (m_k, m_{k-2}, \dots, m_{i'} - 1) = s'$. The former follows from $i' > i$ and the fact that both i and i' are odd, and the latter follows from $i' > i$ together with $m_{i'} < m'_{i'}$ by the definition of lexicographic ordering.

Putting $l = n(k + 1)$ and $h = k + 1$, we hence obtain an $(n(k + 1), k + 1)$ -universal graph of (quasipolynomial) size at most $2n(k + 1) \binom{\log(n(k+1)) + k + 2}{k+1}$ and a simulation order $\geq$ on this graph. $\square$

Full proof of Lemma 6.4.

Proof. For one direction of the proof, we let μ be the measure that is obtained at the end of the execution of the lifting algorithm. Then we have $\text{Lift}(\mu) = \mu$ so

that μ is a progress measure. Let $(v, p) \in \mathbb{E}$. By Lemma 4.6 it suffices to exhibit a history-free witness (W, R) for $v \in \llbracket X_p \rrbracket_f$. We put $W = \mathbb{E}$, having $(v, p) \in W$, and define the relation $R \subseteq W \times [d] \times W$ by

$$R_{i'}(u, i) = \{(w, i') \in W \mid \exists s \in K_{\text{ad}(i')}(\mu(u, i)). \mu(w, i') \leq s\}$$

for $(u, i) \in W$ and $0 \leq i' \leq d$, having $\star \neq \mu(u, i) \in Z$ by definition of W . To show that (W, R) is a history-free witness, we first show that for all $(v, p) \in W$, we have $v \sqsubseteq f_p(U_0, \dots, U_k)$ where

$$U_i = \bigsqcup \pi_1[R_{\text{ad}(i)}(v, p)]$$

for $1 \leq i \leq k$. Since μ is a progress measure, we have

$$v \sqsubseteq f_p(U_0^{\mu, \mu(v, p)}, \dots, U_k^{\mu, \mu(v, p)}).$$

Our claim $v \sqsubseteq f_p(U_0, \dots, U_k)$ follows since we have, for all $1 \leq i \leq k$, that by the definition of $R_{\text{ad}(i)}$,

$$\begin{aligned} U_i &= \bigsqcup \pi_1[R_{\text{ad}(i)}(v, p)] = \bigsqcup \pi_1[\{(w, i) \in W \mid \exists s \in K_{\text{ad}(i)}(\mu(v, p)). \mu(w, i) \leq s\}] \\ &= \bigsqcup \{w \in B_L \mid \exists s \in K_{\text{ad}(i)}(\mu(v, p)). \mu(w, i) \leq s\} \\ &= U_i^{\mu, \mu(v, p)}. \end{aligned}$$

It remains to show that (W, R) is an even graph. To see this, let

$$\pi = (u_0, p_0), p_1, (u_1, p_1), p_2, (u_2, p_2) \dots$$

be an infinite path in (W, R) , that is, let $((u_i, p_i), p_{i+1}, (u_{i+1}, p_{i+1})) \in R$ for all $i \geq 0$. We have to show that π is even. We construct an infinite path

$$\tau = s_0, q_1, s_1, q_2, s_2 \dots$$

in the universal graph (Z, K) as follows: We put $s_0 = \mu(u_0, p_0)$, and define q_{i+1} and s_{i+1} inductively from s_i for all $i \geq 0$, ensuring the invariant

$$\mu(u_i, p_i) \leq s_i. \tag{A.1}$$

First we note that (A.1) holds trivially at $i = 0$. Next we define q_{i+1} and s_{i+1} for $i \geq 0$, assuming that (A.1) holds at i . By definition of R , it follows from $((u_i, p_i), p_{i+1}, (u_{i+1}, p_{i+1})) \in R$ that there is $s \in K_{\text{ad}(p_{i+1})}(\mu(u_i, p_i))$ such that $\mu(u_{i+1}, p_{i+1}) \leq s$. By the assumption on $\geq$, $s \in K_{\text{ad}(p_{i+1})}(\mu(u_i, p_i))$ and the invariant $s_i \geq \mu(u_i, p_i)$ imply that there is some p and some $s' \in K_p(s_i)$ such that $p \succeq \text{ad}(p_{i+1})$ and $s' \geq s$. Put $q_{i+1} := p$ and $s_{i+1} := s'$; then $s_{i+1} \in K_{q_{i+1}}(s_i)$ and $\mu(u_{i+1}, p_{i+1}) \leq s \leq s_{i+1}$, so the invariant (A.1) holds at $i + 1$. For later use, we point out that we have $q_i \succeq \text{ad}(p_i)$ for all $i > 0$ by construction of τ .

The path τ , being an infinite path in the universal graph (Z, K) , is even. Since we have $q_i \succeq \text{ad}(p_i)$ for all $i > 0$, p_i is even if and only if q_i is even (and

then $q_i = \text{ad}(p_i)$, and p_i is odd if and only if q_i is odd (and then $q_i \geq \text{ad}(p_i)$). Let q be the highest priority that occurs infinitely often in τ . Then q is even and there are infinitely many positions i such that $q_i = q = \text{ad}(p_i)$ and there is some position i' such that for all positions $i'' > i'$, $\text{ad}(p_{i''}) \leq q_{i''} \leq q$. Hence π is an even path, as required.

For the converse direction, let $u \in \llbracket X_i \rrbracket_f$. We note that the lifting algorithm computes a progress measure as the least fixpoint of the function Lift in the lattice of measures, pointwise ordered by the order $\leq$ on nodes of the universal graph. Let μ be the measure that is obtained at the end of the execution of the lifting algorithm. Then we have $\text{Lift}(\mu) = \mu$ and $\mu = \text{Lift}^n(\mathbf{m}_{\min})$ for some natural number n (recall that $\mathbf{m}_{\min}$ denotes the minimal measure). We have to show that $(u, i) \in \mathbb{E}$, i.e. $\mu(u, i) \neq \star$. By Theorem 5.3, there is $q \in Z$ such that $(u, i, q) \in \llbracket Y_0 \rrbracket_g$. For all (v, p) such that there is $q' \in Z$ such that $(v, p, q') \in \llbracket Y_0 \rrbracket_g$, we let $q_{(v,p)} \in Z$ denote the least (w.r.t. $\leq$) node of the universal graph with this property. It suffices to show that in this case,

$$(\text{Lift}^n(\mathbf{m}_{\min}))(v, p) \leq q_{(v,p)}.$$

We proceed by induction over n . If $n = 0$, then $\text{Lift}^n(\mathbf{m}_{\min}) = \mathbf{m}_{\min}$ and $\mathbf{m}_{\min}(v, p) = q_{\min} \leq q_{(v,p)}$. If $n > 0$, then we have to show the inequality in

$$\begin{aligned} (\text{Lift}^n(\mathbf{m}_{\min}))(v, p) &= \text{Lift}(\text{Lift}^{n-1}(\mathbf{m}_{\min}))(v, p) \\ &= \min\{q \in Z \mid v \sqsubseteq f_p(U_0^{\mu_{n-1}, q}, \dots, U_k^{\mu_{n-1}, q})\} \leq q_{(v,p)} \end{aligned}$$

where $\mu_{n-1} = \text{Lift}^{n-1}(\mathbf{m}_{\min})$. This follows once we show that

$$q_{(v,p)} \in \{q \in Z \mid v \sqsubseteq f_p(U_0^{\mu_{n-1}, q}, \dots, U_k^{\mu_{n-1}, q})\},$$

i.e.

$$v \sqsubseteq f_p(U_0^{\mu_{n-1}, q_{(v,p)}}, \dots, U_k^{\mu_{n-1}, q_{(v,p)}}). \quad (\text{A.2})$$

Now since $\llbracket Y_0 \rrbracket_g$ is a fixpoint of g , we have

$$\llbracket Y_0 \rrbracket_g = g(\llbracket Y_0 \rrbracket_g) = \{(v, p, q) \in B_L \times [k] \times Z \mid v \sqsubseteq f_p(P_0^{\llbracket Y_0 \rrbracket_g, q}, \dots, P_k^{\llbracket Y_0 \rrbracket_g, q})\}.$$

So given that $(v, p, q_{(v,p)}) \in \llbracket Y_0 \rrbracket_g$, we have

$$v \sqsubseteq f_p(P_0^{\llbracket Y_0 \rrbracket_g, q_{(v,p)}}, \dots, P_k^{\llbracket Y_0 \rrbracket_g, q_{(v,p)}}),$$

which implies our goal (A.2) by monotonicity of f_p once we show that $P_i^{\llbracket Y_0 \rrbracket_g, q_{(v,p)}} \sqsubseteq U_i^{\mu_{n-1}, q_{(v,p)}}$ for $0 \leq i \leq k$. Both sides of this inequality are defined as joins of sets; it suffices to show the set inclusion between these sets. So let $u' \in B_L$ such that there is $q' \in K_{\text{ad}(i)}(q_{(v,p)})$ such that $(u', i, q') \in \llbracket Y_0 \rrbracket_g$ ($P_i^{\llbracket Y_0 \rrbracket_g, q_{(v,p)}}$ is the join of all such u'). We claim that u' is in the set

$$\{u'' \in B_L \mid \exists s \in K_{\text{ad}(i)}(q_{(v,p)}). (\text{Lift}^{n-1}(\mathbf{m}_{\min}))(u'', i) \leq s\}$$

(whose join is $U_i^{\mu_{n-1}, q(v,p)}$ since $\mu_{n-1} = \text{Lift}^{n-1}(\mathbf{m}_{\min})$). Indeed, pick $s = q'$ so that $s \in K_{\text{ad}(i)}(q(v,p))$. It remains to show that $(\text{Lift}^{n-1}(\mathbf{m}_{\min}))(u', i) \leq s$. By the inductive hypothesis, $(\text{Lift}^{n-1}(\mathbf{m}_{\min}))(u', i) \leq q_{(u', i)}$, so it suffices to show $q_{(u', i)} \leq s$. This follows from the fact that $s = q'$ and $(u', i, q') \in \llbracket Y_0 \rrbracket_g$, since $q_{(u', i)}$ is by definition the least node (w.r.t. $\leq$) such that $(u', i, q_{(u', i)}) \in \llbracket Y_0 \rrbracket_g$. $\square$

A.2 Details on Applications to Coalgebraic μ -Calculi

Example A.1. We consider probabilistic parity games, which make use of systems of fixpoint equations that deviate considerably from (and apparently do not reduce easily to) the ones for standard parity games. Probabilistic parity games are parity games in which both moves and nodes are annotated with probabilities (these games are not to be confused with the $2\frac{1}{2}$ -player *stochastic parity games* that are considered in [12, 27]). They arise naturally as model checking games for the (two-valued) probabilistic μ -calculus (see Example A.2.(2)); we postpone a more formal and detailed treatment to Section A.2 below, where we discuss the more general coalgebraic μ -calculus (covering e.g. probabilistic, graded and the alternating-time μ -calculi as instances) and its model checking problem (corresponding to solving e.g. probabilistic, graded and alternating-time parity games).

A *probabilistic parity game* (V, D, Ω, σ) consists of a set V of nodes, a set of probabilistic moves, given by a function D which assigns probability distributions $D(v)$ over V (with $\sum_{w \in V} (D(v))(w) = 1$) to nodes $v \in V$, a priority function $\Omega : V \rightarrow \mathbb{N}$ and a *probability assignment* $\sigma : V \rightarrow [0, 1]$. The intuition of $(D(v))(w) = p$ is that the move from v node to node w has probability p . A *play* $\rho = v_0, v_1, \dots$ in a probabilistic parity game is a sequence of nodes such that for all $i \geq 0$, we have $(D(v_i))(v_{i+1}) > 0$ and the winning conditions on plays are the same as in standard parity games. A (*history-free*) *strategy* is a partial function $s : V \rightarrow \mathcal{P}(V)$ such that for all $v \in \text{dom}(s)$, $\sum_{w \in s(v)} (D(v))(w) > \sigma(v)$, that is, strategies pick *sets* of moves whose joint probability is larger than the probability assignment of the respective node. Crucially, strategies in probabilistic parity games involve branching, unlike strategies in standard parity games which pick single moves for each node on which they are defined. An *s-play* is a play $\rho = v_0, v_1, \dots$ such that for all $i \geq 0$, $v_i \in \text{dom}(s)$ and $v_{i+1} \in s(v_i)$. Player Eloise wins a node v if there is a strategy s such that Eloise wins all *s*-plays that start at v . We point out that the (somewhat concealed) two-player nature of probabilistic parity games manifests in the fact that Eloise has to pick, in each turn, *some* suitable set of moves, whereupon Abelard can challenge *any* of these moves (that is, *all* *s*-plays need to be even in order for s to be a winning strategy for Eloise). Player Eloise hence wins a node v if and only if there is a set $W \subseteq V$ containing v and a graph $G = (W, R \subseteq W \times W)$ such that

- each R -edge has D -probability greater than 0,
- for each $w \in W$, the R -successors of w have a joint D -probability of more than $\sigma(w)$ and

- for each infinite R -path that starts at v , the highest priority that is visited infinitely often by the path is even.

Consider, for instance, the probabilistic parity game depicted below with $V = \{0, 1, 2, 3\}$, $\Omega(i) = i$ for $i \in V$ and, e.g. $(D(0))(1) = 0.2$, $(D(3))(1) = 1$ and $(D(3))(3) = 0$; let us fix the probability assignment σ by putting $\sigma(0) = 0.7$, $\sigma(1) = 0.3$, $\sigma(2) = 0.1$ and $\sigma(3) = 0$.

To win e.g. the node 0, Eloise has to have a strategy that selects a set of nodes that have a joint probability (of being reached from 0 in one step) greater than $\sigma(0) = 0.7$ and that are in turn all won by the strategy. In this example, player Eloise wins the nodes 0 and 2 with the strategy s defined by $s(0) = \{0, 2\}$ and $s(2) = \{2\}$: this function indeed is a valid strategy since it uses only moves with nonzero probabilities and also respects the probability assignment σ as we have $\sum_{j \in s(0)} (D(0))(j) = 0.5 + 0.3 > \sigma(0) = 0.7$ and $\sum_{j \in s(2)} (D(2))(j) = 0.4 > \sigma(2) = 0.1$. Also, every s -play that starts at node 0 is of the form 0^ω or 0^*2^ω and hence even and every s -play that starts at node 2 is of the form 2^ω and hence even. On the other hand, there is no strategy with which Eloise can win the nodes 1 or 3 since for any strategy t with $1 \in \text{dom}(t)$, we have $\sum_{w \in t(1)} (D(1))(w) > \sigma(1) = 0.3$ and hence $1 \in t(1)$; but then there is an odd t -play of the form 1^ω so that t is not a winning strategy for Eloise. Also, for any candidate winning strategy u with $3 \in \text{dom}(u)$, we have $u(3) = \{1\}$ and hence $1 \in \text{dom}(u)$ which shows that u is not a winning strategy for Eloise. Hence we have $\text{win}_\exists = \{0, 2\}$ and $\text{win}_\forall = \{1, 3\}$.

The winning regions in probabilistic parity games are again just nested fix-points, where the functions however deviate significantly from the functions for standard parity games. Player Eloise has to pick sets of moves now, so it does not suffice to consider existential or universal branching, like in the function $f_\exists$ for standard parity games. We define $f_{\exists p} : \mathcal{P}(V)^{k+1} \rightarrow \mathcal{P}(V)$, for $(V_0, \dots, V_k) \in \mathcal{P}(V)^{k+1}$, by putting

$$f_{\exists p}(V_0, \dots, V_k) = \{v \in V \mid \exists 0 \leq i \leq k. \Omega(v) = i, \\ \sum_{w \in V_i} (D(v))(w) > \sigma(v)\}$$

Then we have $\text{win}_\exists = \mathbf{E}^{f_{\exists p}}$ (formally, this is a consequence of Lemma A.4, below). The winning region of Abelard is characterized in a dual manner. Note that nodes v with $\sigma(v) = p$ correspond to model checking modal operators $\langle p \rangle$ which require that their argument is satisfied with probability more than p in the next step (see Example A.2.(2)). The full probabilistic μ -calculus also has dual operators $[p]$ which state that their argument holds with probability at least $1 - p$ in the next step; for brevity, we refrain from modelling these operators in this example.

While the above example apparently already goes beyond the setting of [11] in which standard parity games with existential and universal branching are hardwired throughout (e.g. in the set operator $CPre$ and the function $best$), we note that our results cover systems of fixpoint equations for arbitrary functions over finite lattices, which need not be ‘game-like’ at all, that is, they need not be parametrized by any graph structure or priority and player assignment.

We next show how to apply our results to model checking and satisfiability checking for generalized μ -calculi in the setting of coalgebraic logic, covering, for instance, graded [38], probabilistic [15, 41], and alternating-time [1] μ -calculi. It has been shown in previous work [30] that model checking for coalgebraic μ -calculi reduces to computing winning regions in a generalized variant of parity games where the game arenas are coalgebras instead of Kripke frames. We proceed to recall basic definitions and examples in universal coalgebra [49] and the coalgebraic μ -calculus [14] and then continue to show that our main result yields new quasipolynomial-time upper bounds for the model checking problem and improves the known exponential-time upper bound for the satisfiability problem [31] of the coalgebraic μ -calculus. These generic results instantiate to new upper bounds in all concrete cases except the standard relational μ -calculus.

The abstraction principle underlying universal coalgebra is to encapsulate system types as functors, for our present purposes on the category of sets. Such a functor $T : \mathbf{Set} \rightarrow \mathbf{Set}$, which we fix in the following, maps every set X to a set TX , and every map $f : X \rightarrow Y$ to a map $Tf : TX \rightarrow TY$, preserving identities and composition. We think of TX as a type of structured collections over X ; a basic example is the covariant powerset functor $\mathcal{P}$, which assigns to each set its powerset and acts on maps by taking forward image. Systems of the intended type are then cast as T -coalgebras (C, ξ) (or just ξ) consisting of a set C of *states* and a *transition map* $\xi : C \rightarrow TC$, thought of as assigning to each state $x \in C$ a structured collection $\xi(x) \in TC$ of successors. E.g. a $\mathcal{P}$ -coalgebra $\xi : C \rightarrow \mathcal{P}C$ assigns to each state a set of successors, i.e. is a transition system.

Following the paradigm of *coalgebraic logic* [15], we fix a set Λ of modal operators; we interpret each $\heartsuit \in \Lambda$ as *predicate lifting* $\llbracket \heartsuit \rrbracket$ for T , i.e. a natural transformation

$$\llbracket \heartsuit \rrbracket_X : 2^X \rightarrow 2^{TX}.$$

Here, the index X ranges over all sets; 2^X denotes the set of maps $X \rightarrow 2$ into the two-element set $2 = \{\perp, \top\}$, isomorphic to the powerset of X (i.e. 2^X is the *contravariant powerset functor*; we generally keep the conversion between 2^X and $\mathcal{P}(X)$ implicit); and naturality means that $\llbracket \heartsuit \rrbracket_X(f^{-1}[A]) = (Tf)^{-1}[\llbracket \heartsuit \rrbracket_Y(A)]$ for $f : X \rightarrow Y$ and $A \in 2^Y$. Thus, the predicate lifting $\llbracket \heartsuit \rrbracket$ indeed lifts predicates on a base set X to predicates on the set TX . Standard examples for $T = \mathcal{P}$ are the predicate liftings for the $\Box$ and $\Diamond$ modalities, given by

$$\begin{aligned} \llbracket \Box \rrbracket_X(A) &= \{B \in \mathcal{P}X \mid B \subseteq A\} & \text{and} \\ \llbracket \Diamond \rrbracket_X(A) &= \{B \in \mathcal{P}X \mid B \cap A \neq \emptyset\} \end{aligned}$$

for $A \in PX$. Since we mean to form fixpoint logics, we need to require that every $\llbracket \heartsuit \rrbracket$ is *monotone*, that is, $A \subseteq B \subseteq X$ implies $\llbracket \heartsuit \rrbracket_X(A) \subseteq \llbracket \heartsuit \rrbracket_X(B)$. To support negation, we assume moreover that Λ is closed under *duals*, i.e. for each $\heartsuit \in \Lambda$ we have $\overline{\heartsuit} \in \Lambda$ such that $\llbracket \overline{\heartsuit} \rrbracket_X(A) = TX \setminus \llbracket \heartsuit \rrbracket_X(X \setminus A)$, chosen so that $\overline{\overline{\heartsuit}} = \heartsuit$ (e.g. $\overline{\Box} = \Diamond$, $\overline{\Diamond} = \Box$).

Given a set $\mathbf{Var}$ of *fixpoint variables*, the set of *formulae* $\phi, \psi, \dots$ of the coalgebraic μ -calculus is then defined by the grammar

$$\psi, \phi := \top \mid \perp \mid \psi \vee \phi \mid \psi \wedge \phi \mid \heartsuit \psi \mid X \mid \eta X. \psi$$

$$(\heartsuit \in \Lambda, X \in \mathbf{Var}, \eta \in \{\mu, \nu\}).$$

Given a T -coalgebra $\xi : C \rightarrow TC$ and a valuation $\sigma : \mathbf{Var} \rightarrow \mathcal{PC}$, the *extension*

$$\llbracket \phi \rrbracket_\sigma \subseteq C$$

of a formula ϕ is defined recursively by $\llbracket X \rrbracket_\sigma = \sigma(X)$; the expected clauses for the propositional operators ($\llbracket \top \rrbracket_\sigma = C$; $\llbracket \perp \rrbracket_\sigma = \emptyset$; $\llbracket \phi \wedge \psi \rrbracket_\sigma = \llbracket \phi \rrbracket_\sigma \cap \llbracket \psi \rrbracket_\sigma$; $\llbracket \phi \vee \psi \rrbracket_\sigma = \llbracket \phi \rrbracket_\sigma \cup \llbracket \psi \rrbracket_\sigma$); and

$$\begin{aligned} \llbracket \heartsuit \psi \rrbracket_\sigma &= \xi^{-1}[\llbracket \heartsuit \rrbracket(\llbracket \psi \rrbracket_\sigma)] \\ \llbracket \mu X. \psi \rrbracket_\sigma &= \text{LFP} \llbracket \psi \rrbracket_\sigma^X \\ \llbracket \nu X. \psi \rrbracket_\sigma &= \text{GFP} \llbracket \psi \rrbracket_\sigma^X \end{aligned}$$

where the (monotone) map $\llbracket \psi \rrbracket_\sigma^X : \mathcal{PC} \rightarrow \mathcal{PC}$ is defined by $\llbracket \psi \rrbracket_\sigma^X(A) = \llbracket \psi \rrbracket_{\sigma[X \mapsto A]}$ for $A \subseteq C$, with $(\sigma[X \mapsto A])(X) = A$ and $(\sigma[X \mapsto A])(Y) = \sigma(Y)$ for $X \neq Y$.

The *alternation depth* $\text{ad}(\eta X. \psi)$ of a fixpoint $\eta X. \psi$ is the depth of alternating nesting of such fixpoints in ψ that depend on X ; we assign *odd* numbers to least fixpoints and *even* numbers to greatest fixpoints. E.g. for $\psi = \nu X. \phi$ and $\phi = \mu Y. (p \wedge \heartsuit X) \vee \heartsuit Y$, we have $\text{ad}(\psi) = 2$, $\text{ad}(\phi) = 1$. For a detailed definition of alternation depth, see e.g. [44].

Example A.2. As indicated above, the standard relational μ -calculus [35] is one example of a coalgebraic μ -calculus, with propositional atoms treated as nullary modalities. Further important examples are as follows [14, 50, 51].

(1) The *graded μ -calculus* [38] has modalities $\langle b \rangle$, $[b]$, indexed over $b \in \mathbb{N}$, read ‘in more than b successors’ and ‘in all but at most b successors’, respectively. These can be interpreted over relational structures but it is more natural and technically more convenient to use *multigraphs* [18], i.e. transition systems with edge weights (*multiplicities*) in $\mathbb{N} \cup \{\infty\}$, which are coalgebras for the multiset functor $\mathcal{B}$ given by $\mathcal{B}X = (X \rightarrow (\mathbb{N} \cup \{\infty\}))$. Over $\mathcal{B}$, we interpret $\langle b \rangle$ and $[b]$ by the mutually dual predicate liftings

$$\begin{aligned} \llbracket \langle b \rangle \rrbracket_X(A) &= \{\beta \in \mathcal{B}X \mid \sum_{x \in A} \beta(x) > b\} \\ \llbracket [b] \rrbracket_X(A) &= \{\beta \in \mathcal{B}X \mid \sum_{x \in X \setminus A} \beta(x) \leq b\}. \end{aligned}$$

E.g. the formula $\nu X. (\phi \wedge \Diamond_1 X)$ says that the current state is the root of an infinite tree with branching degree at least 2 (counting multiplicities) on which ϕ holds everywhere.

(2) The (two-valued) *probabilistic μ -calculus* [14, 41] is interpreted over Markov chains, which are coalgebras for the *discrete distribution functor* $\mathcal{D}$ where $\mathcal{D}X = \{\beta : X \rightarrow [0, 1] \mid \sum_{x \in X} \beta(x) = 1\}$ is the set of discrete probability distributions on X , represented, e.g., as probability mass functions $\beta : X \rightarrow [0, 1]$. We abuse β to denote also the induced probability distribution, writing $\beta(A) = \sum_{x \in A} \beta(x)$ for $A \subseteq X$. The logic has modalities $[p]$, $\langle p \rangle$ indexed over $p \in [0, 1] \cap \mathbb{Q}$, interpreted over $\mathcal{D}$ by

$$\begin{aligned} \llbracket \langle p \rangle \rrbracket_X(A) &= \{\beta \in \mathcal{D}X \mid \beta(A) > p\} \\ \llbracket [p] \rrbracket_X(A) &= \{\beta \in \mathcal{D}X \mid \beta(X \setminus A) \leq p\}. \end{aligned}$$

This example (as well as the previous one) can be extended to admit (monotone) polynomial inequalities among probabilities (or multiplicities, respectively) instead of only comparison with constants, allowing, e.g., for expressing probabilistic independence [23, 31, 39]. In more detail, we can introduce n -ary modalities $L_{p,b}$, $M_{p,b}$ indexed over polynomials $p \in \mathbb{Q}_{\geq 0}[x_1, \dots, x_n]$ and rational numbers $b \geq 0$, with $L_{p,b}$ interpreted by the predicate lifting

$$\begin{aligned} \llbracket L_{p,b} \rrbracket_X(A_1, \dots, A_n) &= \{\beta \in \mathcal{D}X \mid \\ &\quad p(\beta(A_1), \dots, \beta(A_n)) > b\} \end{aligned}$$

and $M_{p,b}$ by the corresponding dual predicate lifting. E.g. the formula

$$\nu X. \mu Y. L_{x_1 x_2, 0.8}(p \wedge X, q \vee Y)$$

says roughly that if we independently sample two successors of the current state, then with probability at least 0.8, the first successor state will satisfy p , and then X again (continuing indefinitely), and the second successor state will remain on a path where it satisfies Y again until it eventually reaches q .

(3) *Monotone μ -calculus*: The *monotone neighbourhood functor* $\mathcal{M}$ maps a set X to the set

$$\mathcal{M}X = \{\mathfrak{A} \in 2^{(2^X)} \mid \mathfrak{A} \text{ upwards closed}\}$$

of set systems over X that are upwards closed under subset inclusion (i.e. $A \in \mathfrak{A}$ and $A \subseteq B$ imply $B \in \mathfrak{A}$). Coalgebras for $\mathcal{M}$ are *monotone neighbourhood frames* in the sense of Scott-Montague semantics [13]. We take $\Delta = \{\Box, \Diamond\}$ and interpret $\Box$ over $\mathcal{M}$ by the predicate lifting

$$\begin{aligned} \llbracket \Box \rrbracket_X(A) &= \{\mathfrak{A} \in \mathcal{M}X \mid A \in \mathfrak{A}\} \\ &= \{\mathfrak{A} \in \mathcal{M}X \mid \exists B \in \mathfrak{A}. B \subseteq A\}, \end{aligned}$$

and $\Diamond$ by the corresponding dual lifting, $\llbracket \Diamond \rrbracket_X(A) = \{\mathfrak{A} \in \mathcal{M}X \mid (X \setminus A) \notin \mathfrak{A}\} = \{\mathfrak{A} \in \mathcal{M}X \mid \forall B \in \mathfrak{A}. B \cap A \neq \emptyset\}$. The arising coalgebraic μ -calculus is known

as the *monotone μ -calculus* [22]. When we add propositional atoms and actions, and replace $\mathcal{M}$ with its subfunctor $\mathcal{M}_s$ defined by $\mathcal{M}_s X = \{\mathfrak{A} \in \mathcal{M}X \mid \emptyset \notin \mathfrak{A} \ni X\}$, whose coalgebras are *serial* monotone neighbourhood frames, we arrive at the ambient fixpoint logic of *concurrent dynamic logic* [48] and Parikh's *game logic* [45]. In game logic, actions are understood as atomic games of Angel vs. Demon, and we read $\Box_a \phi$ as 'Angel has strategy to enforce ϕ in game a '. Game logic is then mainly concerned with composite games, formed by the control operators of dynamic logic and additional ones; the semantics can be encoded into fixpoint definitions. For instance, the formula $\nu X. p \wedge \Box_a X$ says that Angel can enforce p in the composite game where a is played repeatedly, with Demon deciding when to stop.

(4) *Alternating-time μ -calculus*: Fix a set $N = \{1, \dots, n\}$ of *agents*. Using alternative notation from *coalition logic* [47], we present the *alternating-time μ -calculus* (AMC) [1] by modalities $[D]$, $\langle D \rangle$ indexed over *coalitions* $D \subseteq N$, read ' D can enforce' and ' D cannot prevent', respectively. We define a functor $\mathcal{G}$ by

$$\mathcal{G}X = \{(k_1, \dots, k_n, f) \mid k_1, \dots, k_n \in \mathbb{N} \setminus \{0\}, \\ f : (\prod_{i \in N} [k_i]) \rightarrow X\}$$

where we write $[k] = \{1, \dots, k\}$ in this example. We understand $(k_1, \dots, k_n, f) \in \mathcal{G}X$ as a one-step concurrent game with k_i available moves for agent $i \in N$, and outcomes in X determined by the *outcome function* f from a joint choice of moves by all the agents. For $D \subseteq N$, we write $S_D = \prod_{i \in D} [k_i]$. Given joint choices $s_D \in S_D$, $s_{\overline{D}} \in S_{\overline{D}}$ of moves for D and $\overline{D} = N \setminus D$ respectively, we write $(s_D, s_{\overline{D}}) \in s_N$ for the joint move of all agents induced in the evident way. In this notation, we interpret the modalities $[D]$ over $\mathcal{G}$ by the predicate lifting

$$[[D]]_X(A) = \{(k_1, \dots, k_n, f) \in \mathcal{G}X \mid \\ \exists s_D \in S_D. \forall s_{\overline{D}} \in S_{\overline{D}}. f(s_D, s_{\overline{D}}) \in A\},$$

and the modalities $\langle D \rangle$ by dualization. This captures exactly the semantics of the AMC: $\mathcal{G}$ -coalgebras are precisely *concurrent game structures* [1], i.e. assign a one-step concurrent game to each state, and $[D]\phi$ says that the agents in D have a joint move such that however the agents in $\overline{D}$ move, the next state will satisfy ϕ . E.g. $\nu Y. \mu X. (p \wedge [D]Y) \vee [D]X$ says that coalition D can enforce that p is satisfied infinitely often.

We now fix a target formula χ that does not contain free fixpoint variables, assuming w.l.o.g. that χ is *clean*, i.e. that every fixpoint variable is bound by at most one fixpoint operator in χ . For a variable $x \in \text{Var}$ that is bound in χ , we then write $\theta(x)$ to denote the formula $\eta X. \psi$ that is a subformula of χ . Let $\text{Cl}(\chi)$ be the *closure* (that is, the set of subformulae) of χ . We have $|\text{Cl}(\chi)| \leq |\chi|$, where $|\chi|$ denotes the number of operators or variables in χ .

We proceed to recall how model checking in the coalgebraic μ -calculus is reduced to computing a nested fixpoint of a particular function [30]:

Definition A.3 (Coalgebraic model checking function). Let $\xi : C \rightarrow TC$ be a coalgebra, and $U = \text{Cl}(\chi) \times C$. The (coalgebraic) model checking function $\alpha_{\text{mc}} : \mathcal{P}(U)^{k+1} \rightarrow \mathcal{P}(U)$ is given by putting, for $\mathbf{U} = (U_1, \dots, U_{k+1}) \in \mathcal{P}(U)^{k+1}$,

$$\begin{aligned} \alpha_{\text{mc}}(\mathbf{U}) = & \{(\top, x) \mid (\top, x) \in U\} \cup \\ & \{(\heartsuit\psi, x) \in U \mid \xi(x) \in \llbracket \heartsuit \rrbracket \{y \mid (\psi, y) \in U_1\}\} \cup \\ & \{(\psi \vee \phi, x) \in U \mid \{(\psi, x), (\phi, x)\} \cap U_1 \neq \emptyset\} \cup \\ & \{(\psi \wedge \phi, x) \in U \mid \{(\psi, x), (\phi, x)\} \subseteq U_1\} \cup \\ & \{(\eta X. \psi, x) \in U \mid (\psi, x) \in U_1\} \cup \\ & \{(X, x) \mid (\theta(X), x) \in U_{\text{ad}(\theta(X))+1}\}. \end{aligned}$$

Lemma A.4 (Coalgebraic model checking [30]). Let χ be a formula of alternation depth k , $\xi : C \rightarrow TC$ a coalgebra, and $x \in C$ a state. Then we have

$$(\chi, x) \in \mathbf{A}^{\alpha_{\text{mc}}} \text{ if and only if } x \in \llbracket \chi \rrbracket.$$

The one-step satisfaction problem consists in deciding whether $t \in \llbracket \heartsuit \rrbracket(W)$, for given $t \in TC$, $\heartsuit \in \Lambda$ and $W \subseteq C$. The time $t(\alpha_{\text{mc}})$ it takes to compute the model checking function α_{mc} hence depends on the time it takes to solve the one-step satisfaction problem for the modal operators at hand. By Corollary 6.5, we obtain

Corollary A.5. Model checking for coalgebraic μ -calculus formulae of alternation depth k against coalgebras $\xi : C \rightarrow TC$ can be done in time $t_2 \cdot t_1$ where $t_2 = \max(t(\alpha_{\text{mc}}), t_1)$, $t_1 = 2n^3(k+1)^2 \binom{\log(n(k+1))+k+2}{k+1}$, and $n = |\text{Cl}(\chi)| \cdot |C|$.

Example A.6 (Quasipolynomial-time model checking for graded and probabilistic μ -calculi). In [30, Examples 3.2 and 3.3], it was shown that in the graded and probabilistic cases, the one-step satisfaction problem can be solved in time $\mathcal{O}(\text{size}(\chi) \cdot |C|)$ and $\mathcal{O}(\text{size}(C)^2 \cdot |C|^3)$, respectively; here, $\text{size}(\chi)$ denotes the representation size of the formula χ and $\text{size}(C)$ denotes the representation size of the coalgebra C . We hence obtain the following quasipolynomial upper time bounds for the model checking problems of the respective μ -calculi, both with numbers coded in binary (where $t_1 = 2n^3(k+1)^2 \binom{\log(n(k+1))+k+2}{k+1}$ and $n = |\text{Cl}(\chi)| \cdot |C|$):

- for the graded μ -calculus: $\mathcal{O}(t_1 \cdot t_2)$, where $t_2 = \max(\text{size}(\chi) \cdot |C|, t_1)$;
- for the probabilistic μ -calculus: $\mathcal{O}(t_1 \cdot t_2)$, where $t_2 = \max(\text{size}(C)^2 \cdot |C|^3, t_1)$.

Similar bounds, with slightly larger t_2 , are obtained for the respective extensions with polynomial inequalities. To the best of our knowledge, these bounds are new. We similarly obtain quasipolynomial bounds for model checking the monotone μ -calculus and the alternating-time μ -calculus. In these cases, the time bounds are already in [30], via an encoding into standard parity games; but we emphasize again that the point of our main result (Corollary 6.5) is not so much the time bound but rather the quasipolynomial bound on the number of iterations – in this case, we obtain that the fixpoint can be computed with quasipolynomially many calls to the one-step satisfaction problem (which

at least for the alternating-time case seems also algorithmically preferable to an encoding in parity games with many additional states).

We now consider satisfiability checking for the coalgebraic μ -calculus, which also reduces to the computation of a nested fixpoints of a certain function [31]. We recall the essential notions that are required to define this function; see [31] for details of the construction. We fix a target formula χ of size n and alternation depth k , to be checked for satisfiability. One then has a deterministic parity automaton that accepts precisely the *good branches* in tableaux representing prospective models of χ , i.e. the ones not containing infinite deferrals of least fixpoints (which represent eventualities). We work with parity automata in which priorities are assigned to the *transitions* (rather than the states); our automaton thus has the form $(D_\chi, \Sigma, \delta, \beta)$ where D_χ is the set of states; Σ is the alphabet (designed to allow identifying manipulations of formulae happening in the transitions); δ is the transition function; and β assigns priorities to transitions. Since the automaton is deterministic, we can take β to be a function $D_\chi \times \Sigma \rightarrow \mathbb{N}$. Recall that such a parity automaton accepts an infinite word w if and only if it has a run for w in which the highest priority that occurs infinitely often is even. We have $|D_\chi| \in \mathcal{O}(2^{\mathcal{O}(nk \log n)})$, and nodes $v \in D_\chi$ are labelled with sets $l(v)$ of formulae. We denote the set of nodes whose labels contain some propositional formula by **prestates** and the set of nodes whose labels contain only modal formulae by **states**; for $v \in \text{prestates}$, ψ_v is a fixed propositional formula from the label of v . The transition function δ tracks sets of formulae according to the logical manipulations described by a given letter from Σ . Besides letters identifying propositional transformations, Σ contains sets of modal formulae describing modal steps; we write **selections** $\subseteq \Sigma$ for the set of these letters.

Definition A.7 (Coalgebraic satisfiability checking function [31]). For sets $U \subseteq D_\chi$ and $\mathbf{U} = (U_1, \dots, U_{2nk}) \in \mathcal{P}(U)^{2nk}$, we put

$$\begin{aligned} \alpha_{\text{sat}}(\mathbf{U}) = & \{v \in \text{prestates} \mid \\ & \exists b \in \{0, 1\}. \delta(v, (\psi_v, b)) \in U_{\beta(v, (\psi_v, b))}\} \cup \\ & \{v \in \text{states} \mid T(\bigcup_{1 \leq i \leq 2nk} U_i(v)) \cap \llbracket l(v) \rrbracket_1 \neq \emptyset\} \end{aligned}$$

where $\beta(v, (\psi_v, b))$ abbreviates $\beta(v, (\psi_v, b), \delta(v, (\psi_v, b)))$ and where

$$\begin{aligned} U_i(v) &= \{l(u) \mid u \in X_i, \exists \kappa \in \text{selections}. \\ & \delta(v, \kappa) = u, \beta(v, \kappa, u) = i\}. \end{aligned}$$

The *one-step satisfiability problem* is to decide whether

$$T(\bigcup_{1 \leq i \leq 2nk} U_i(v)) \cap \llbracket l(v) \rrbracket_1 \neq \emptyset$$

for given U, v . Hence checking whether some $v \in \text{prestates}$ is contained in $\alpha_{\text{sat}}(\mathbf{U})$ for given $\mathbf{U}$ is an instance of the one-step satisfiability problem.

Lemma A.8 (Fixpoint characterization of satisfiability [31]). In the above notation,

$$v_0 \in A^{\alpha_{\text{sat}}} \text{ if and only if } \chi \text{ is satisfiable.}$$

Corollary A.9. *If the one-step satisfiability problem of a coalgebraic logic can be solved in time $2^{\mathcal{O}(nk \log n)}$, then the satisfiability problem of the μ -calculus over this logic can be solved in time $2^{\mathcal{O}(nk \log n)}$ as well.*

Proof. By the previous Lemma, it suffices to show that $A^{\alpha_{\text{sat}}}$ can be computed in time $2^{\mathcal{O}(nk \log n)}$. Since we have $2nk < \log |D_\chi|$, $A^{\alpha_{\text{sat}}}$ can – by Corollary 6.6 – be computed in time $\mathcal{O}(t(\alpha_{\text{sat}}) \cdot (|D_\chi|(2nk + 1))^8)$, where $t(\alpha_{\text{sat}})$ denotes the maximum of $(|D_\chi|(2nk + 1))^8$ and the time it takes to compute α_{sat} ; by assumption, α_{sat} can be computed in time $2^{\mathcal{O}(nk \log n)}$ so that we have $t(\alpha_{\text{sat}}) \in 2^{\mathcal{O}(nk \log n)}$. $\square$

Example A.10. It has been shown (e.g. in [31]) that the one-step satisfiability problems of all logics from Example A.2 can be solved in time $2^{\mathcal{O}(nk \log n)}$. Hence we obtain an upper bound $2^{\mathcal{O}(nk \log n)}$ for the satisfiability problems of all these logics, in particular including the monotone μ -calculus, the alternating-time μ -calculus, the graded μ -calculus and the (two-valued) probabilistic μ -calculus, even when the latter two are extended with (monotone) polynomial inequalities. This improves on the best previous bounds in all cases.