

SEMIRING PROVENANCE FOR BÜCHI GAMES: STRATEGY ANALYSIS WITH ABSORPTIVE POLYNOMIALS

ERICH GRÄDEL , NIELS LÜCKING, AND MATTHIAS NAAF 

RWTH Aachen University, Germany
e-mail address: graedel@logic.rwth-aachen.de, niels.luecking@rwth-aachen.de, naaf@logic.rwth-aachen.de

ABSTRACT. This paper presents a case study for the application of semiring semantics for fixed-point formulae to the analysis of strategies in Büchi games. Semiring semantics generalizes the classical Boolean semantics by permitting multiple truth values from certain semirings. Evaluating the fixed-point formula that defines the winning region in a given game in an appropriate semiring of polynomials provides not only the Boolean information on *who* wins, but also tells us *how* they win and *which strategies* they might use. This is well-understood for reachability games, where the winning region is definable as a least fixed point. The case of Büchi games is of special interest, not only due to their practical importance, but also because it is the simplest case where the fixed-point definition involves a genuine alternation of a greatest and a least fixed point.

We show that, in a precise sense, semiring semantics provide information about all *absorption-dominant* strategies – strategies that win with minimal effort, and we discuss how these relate to positional and the more general *persistent* strategies. This information enables applications such as game synthesis or determining minimal modifications to the game needed to change its outcome. Lastly, we discuss limitations of our approach and present questions that cannot be immediately answered by semiring semantics.

1. INTRODUCTION

Two-player games on finite graphs which admit infinite plays are of fundamental importance in many areas of logic and computer science, especially in the formal analysis of reactive systems, where they model the non-terminating interaction between a system and its environment. In such a game, the *objective* or *winning condition* of the player who represents the system specifies the desired set of behaviours of the system. The most basic classes of such objectives are *reachability* and *safety* objectives defined by a set of states (positions) that the player should reach, or avoid. We can assume, without loss of generality, that even though infinite plays are possible in a game with reachability or safety objectives, they are all won by the same player.

Games with genuine and non-trivial winning conditions for infinite plays are harder to analyse; they include games with arbitrary ω -regular objectives, such as liveness, Muller, Streett-Rabin, or parity objectives, and many others. The goal of this paper is to provide

Key words and phrases: Büchi games, strategy analysis, semiring semantics, provenance analysis, fixed-point logic.

This is a revised and expanded version of the conference paper [GLN21], presented at GandALF 2021.

a case study of a recent method for strategy analysis, based on semiring semantics, and we would like to explore its potential for providing detailed information about strategies in genuinely infinite games. One of the simplest class of games with a non-trivial winning condition for the infinite plays are games with Büchi objectives, which require that a specific target set F of states is reached infinitely often during the play (see e.g. [GTW02] for background). Büchi games, as well as some of their straightforward generalisations, have many applications in formal methods, and efficient algorithms for solving them have been studied thoroughly (see e.g. [CDHL16, CH14, CHP08]). They are also of interest from the points of view of topology and logic, because they are among the simplest games where the set of winning plays is neither open nor closed, and where logical definition of the winning region requires a genuine alternation of a greatest and a least fixed point (see Section 5).

Strategies in infinite games can be very complicated because, in principle, they may depend on the entire history of a play. Thus, there exist uncountably many different strategies, even on a finite game graph. Fortunately, in many cases and in particular for Büchi games, simple strategies are sufficient to win. A fundamental result in this context is the positional determinacy of parity games (of which Büchi games are a special case), saying that from each position, one of the two players has a *positional winning strategy*, i.e. a winning strategy that only depends on the current position and not on the history of the play. A positional strategy can be viewed as a subgraph of the game graph, and can therefore be represented in a compact way. As a consequence, the algorithmic analysis of Büchi games has concentrated almost exclusively on the positional strategies. Here we extend this point of view somewhat and take also other kinds of simple strategies into account. Specifically, we are interested in *absorption-dominant* winning strategies [GT20] which are strategies without redundant moves; this means that taking away anything, in the sense of demanding that some specific move is played less often, makes the strategy non-winning. Another way to distinguish positional strategies from absorption-dominant ones concerns their minimisation properties: while positional strategies minimize the *set* of moves that they use, absorption-dominant strategies take multiplicities into account and minimize the *multiset* of moves. A further interesting class are the *persistent* strategies [MT02], which are positional in each individual play but not necessarily across distinct plays. We shall study the relationship between these different classes of simple strategies, and prove that every positional strategy is absorption-dominant and every absorption-dominant strategy is persistent, and that these inclusions are strict.

The specific method for strategy analysis that we want to apply to Büchi games in this paper is based on the logical definability of the winning positions by a formula in the fixed-point logic LFP, and on the semiring semantics for LFP developed in [DGNT21]. In the classical Boolean semantics, a model \mathfrak{A} of a formula φ assigns to each (instantiated) literal a Boolean value. \mathbb{K} -interpretations π , for a suitable semiring \mathbb{K} , generalize this by assigning to each such literal a semiring value from \mathbb{K} . We then interpret 0 as *false* and all other semiring values as *nuances of true* that provide additional information, depending on the semiring: For example, the Boolean semiring $\mathbb{B} = (\{0, 1\}, \vee, \wedge, 0, 1)$ corresponds to Boolean semantics, the Viterbi-semiring $\mathbb{V} = ([0, 1], \max, \cdot, 0, 1)$ can model *confidence* scores, the tropical semiring $\mathbb{T} = (\mathbb{R}_+^\infty, \min, +, \infty, 0)$ is used for cost analysis, and min-max-semirings (A, \max, \min, a, b) for a totally ordered set $(A, <)$ can model different access levels. Most importantly, semirings of polynomials, such as $\mathbb{N}[X]$, allow us to *track* certain literals by mapping them to different indeterminates. The overall value of the formula is then a polynomial that describes precisely what combinations of literals prove the truth of the

formula. Semiring semantics has been studied for various logics [BOPP20, DG19, DG20, DGNT21, GT17a], following the successful development of semiring provenance in database theory and related fields (see e.g. [DMRT14, GP10, GKT07, GT17b, OP18, RMZ⁺20, Sen17]). While semiring provenance analysis for database queries had largely been confined to positive query languages such as conjunctive queries, positive relational algebra, and Datalog, the generalisation to logics such as first-order logic FO and least fixed-point logic LFP – featuring full negation and unrestricted interaction between least and greatest fixed points – poses non-trivial mathematical challenges and requires new algebraic constructions. Specifically, it has turned out that appropriate semirings for LFP should be absorptive and fully continuous [DGNT21]. Fortunately, this is the case for most of the important application semirings such as \mathbb{V} , \mathbb{T} or min-max-semirings, but not for the natural semiring \mathbb{N} , or the general provenance semirings of polynomials or formal power series, $\mathbb{N}[X]$ and $\mathbb{N}^\infty[X]$. Instead, we rely here on an absorptive version of $\mathbb{N}[X]$ (and $\mathbb{N}^\infty[X]$), the semirings $\mathbb{S}^\infty[X]$ of *generalized absorptive polynomials*, which we discuss in Section 4.1. These are the *universal* absorptive, fully-continuous semirings, in the sense that every mapping $h: X \rightarrow \mathbb{K}$ into an absorptive, fully-continuous semiring \mathbb{K} uniquely extends to a fully-continuous semiring homomorphism $h: \mathbb{S}^\infty[X] \rightarrow \mathbb{K}$ that preserves LFP semantics. From valuations of fixed-point formulae in such semirings we thus can derive detailed insights into why the formula holds – and by applying this to the fixed-point definition of winning positions in Büchi games we obtain compact descriptions of winning strategies, in particular of all positional strategies and all absorption-dominant ones.

After an analysis of simple winning strategies in Büchi games, and a short introduction to semiring semantics for fixed-point logic, we shall study the semiring valuations of the particular LFP-formula $\text{win}_0(x)$ that defines the winning region for Player 0 in Büchi games. Given that the objective of Player 0 is to ensure that the play hits the target set F infinitely often, we may informally describe their winning region as the *largest* set Y of positions from which they can enforce a (further) visit to $Y \cap F$ after $k \geq 1$ moves. On the other side the set of positions from which Player 0 can enforce a visit to a target set is the *smallest* set of positions that either are already in the target set, or from which Player 0 can enforce the play to come closer to it. Thus, the winning region of Player 0 can be described as a greatest fixed point inside of which there is a least fixed point, and it is well-known that this fixed-point alternation in the treatment of Büchi objectives cannot be avoided, see e.g. [BW18].

We shall prove a Sum-of-Strategies Theorem, saying that for any position v in a Büchi game, the valuation of the LFP-formula $\text{win}_0(v)$ in an absorptive, fully-continuous semiring coincides with the sum of the valuations of all absorption-dominant winning strategies from v . Besides being of theoretical interest, this result allows to study a number of interesting questions concerning the available winning strategies in a given Büchi game:

Strategy tracking. Introducing indeterminates for all edges in a fixed Büchi game \mathcal{G} , the semiring value $\pi_{\text{strat}}[\text{win}_0(v)]$ for a position v is a polynomial whose monomials are concise descriptions of all absorption-dominant strategies. From these monomials we can derive whether Player 0 wins from v (if there are any monomials) and which edges are used by each absorption-dominant strategy, and how often they appear in the strategy tree. In particular, we can immediately identify and count positional strategies from the polynomial. Going further, we can answer questions such as: can Player 0 still win if we remove edge e , or several edges at once? Can they still win if edge e may only be used finitely often in each play?

Repairing a game. Instead of analysing strategies in a fixed game, we may also reason about modifications or synthesis of (parts of) the game. For example, assuming Player 0 loses from v , what are minimal modifications to the game that would let Player 0 win from v ? To answer such questions we have to take into account also negative information (i.e., absent edges in the graph), so as to find a minimal repair consisting of both moves to delete and moves to add. Algebraically, this requires to extend our semirings by dual-indeterminates, which leads to quotient semirings $\mathbb{S}^\infty[X, \overline{X}]$ by a construction that has been used before in [GT17a, XZAT18, DGNT21] to deal with semiring semantics for negation. We illustrate with the example of minimal repairs that we can indeed derive the desired information from valuations in such semirings.

Cost computation. A typical application of semiring provenance in databases is cost analysis: assuming that atomic facts are not for free but come with a cost (a non-negative real number), then the minimal cost of evaluating a query is described by a provenance valuation in the tropical semiring $\mathbb{T} = (\mathbb{R}_+^\infty, \min, +, \infty, 0)$. In a game, we may ask the analogous question of what is the minimal cost of a winning strategy assuming that moves come with a cost. For reachability and safety games that admit only finite plays, such an analysis works in a reasonably straightforward way by means of an appropriate sum-of-strategies theorem [GT20] (which is much simpler than the one for Büchi games). However, as we shall show, this does not generalize in a nice way to Büchi games, and this seems to be a general limitation of the method of semiring valuations.

2. BÜCHI GAMES AND STRATEGIES

A Büchi game is given by a tuple $\mathcal{G} = (V, V_0, V_1, E, F)$ where V is a set of positions (here assumed to be finite), with a disjoint decomposition $V = V_0 \dot{\cup} V_1$ into positions of Player 0 and positions of Player 1. The relation $E \subseteq V \times V$ specifies the possible moves, and the target set $F \subseteq V$ describes the winning condition. We denote the set of immediate successors of a position v by $vE := \{w \mid vw \in E\}$ and require that $vE \neq \emptyset$ for all v . A play from an initial position v_0 is an infinite path $v_0v_1v_2\dots$ through \mathcal{G} where the successor $v_{i+1} \in v_iE$ is chosen by Player 0 if $v_i \in V_0$ and by Player 1 if $v_i \in V_1$. A play $v_0v_1v_2\dots$ is won by Player 0 if $v_i \in F$ for infinitely many $i < \omega$, otherwise it is won by Player 1. The winning region of Player σ is the set of those positions $v \in V$ such that Player σ has a winning strategy from v , i.e. a strategy that guarantees them a win, no matter what the opponent does.

A strategy for Player σ in $\mathcal{G} = (V, V_0, V_1, E, F)$ can be represented in different ways, for instance as a function $f: V^*V_\sigma \rightarrow V$ that assigns a next position to each partial play ending in a position of Player σ , or simply $f: V_\sigma \rightarrow V$ if the strategy is positional. Here we follow an alternative approach and represent strategies as trees, comprised of all plays that are consistent with the strategy (see, e.g., [GT20]). For simplicity, we only consider strategies of Player 0, so unless mentioned otherwise, *strategy* always refers to a strategy for Player 0.

Definition 2.1. Given a Büchi game $\mathcal{G} = (V, V_0, V_1, E, F)$, the *tree unraveling* from v_0 is the tree $\mathsf{T}(\mathcal{G}, v_0)$ whose nodes are all finite paths ρ from v_0 in \mathcal{G} and whose edges are $\rho \rightarrow \rho v$ for $v \in V$. We often denote a node of $\mathsf{T}(\mathcal{G}, v_0)$ as ρv to indicate a finite path ending in $v \in V$. We write $|\rho|$ for the length of ρ and $\rho \sqsubseteq \rho'$ if ρ is a (not necessarily strict) prefix of ρ' .

Strategies can then be defined as subtrees of the tree unraveling, which allows for a more visual way to reason about strategies. An important detail is that the strategy tree only contains positions (and thus choices for these positions) that are reachable when following

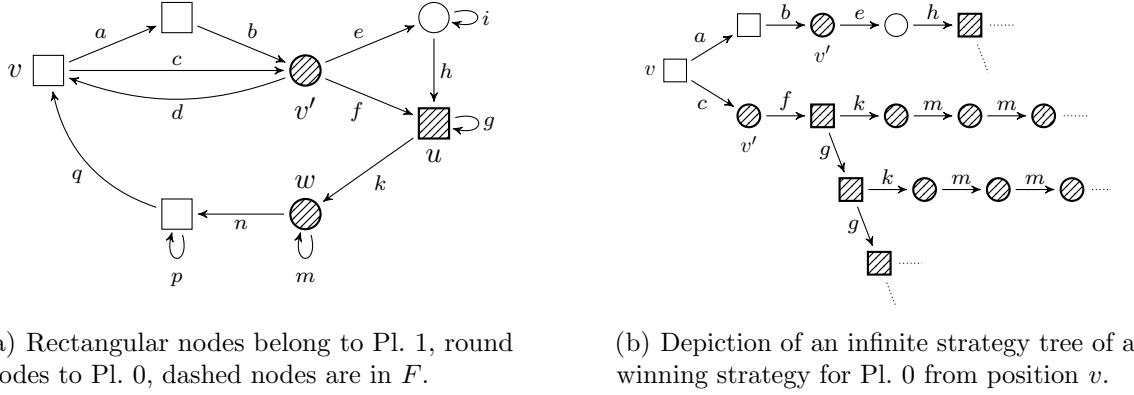


FIGURE 1. Running example of a Büchi game and a winning strategy.

the strategy. Moreover, we only consider finite Büchi games and hence the tree unravelling and all strategies are finitely branching.

Definition 2.2. A *strategy* \mathcal{S} (of Player 0) from v_0 in \mathcal{G} is a subtree of $\mathsf{T}(\mathcal{G}, v_0)$ induced by a node set W satisfying the following conditions:

- if $\rho v \in W$, then also $\rho \in W$ (prefix closure);
- if $\rho v \in W$ and $v \in V_0$, then there is a unique $v' \in vE$ with $\rho v v' \in W$ (unique choice);
- if $\rho v \in W$ and $v \in V_1$, then $\rho v v' \in W$ for all $v' \in vE$ (all moves of the opponent).

The strategy is winning if all plays contained in \mathcal{S} are winning.

We commonly write $\rho \in \mathcal{S}$ instead of $\rho \in W$, and we often refer to paths of the form $\rho v \in \mathcal{S}$ as *occurrences* of v in \mathcal{S} . When we depict strategies graphically, we represent finite paths ρv just by their last position v to ease readability (notice that in the tree unravelling ρ can be reconstructed from v by following the path to the root). See Fig. 1b for an example. For $v \in V_0$, we further write $\mathcal{S}(\rho v) = w$ if $\rho v w$ is the (unique) successor of ρv in \mathcal{S} . If \mathcal{S} is positional, we may also write $\mathcal{S}(v)$ to denote the unique successor of v chosen by \mathcal{S} . We write $\text{Strat}_{\mathcal{G}}(v)$ and $\text{WinStrat}_{\mathcal{G}}(v)$ to denote the set of all (winning) strategies of Player 0 from position $v \in \mathcal{G}$, and we drop \mathcal{G} if the game is clear from the context.

Example 2.3. An example of a Büchi game is depicted in Fig. 1. Player 0 has essentially three different positional winning strategies from v , by either choosing edge d , or edges e, h, m or f, m . Notice that for the first strategy, we did not specify moves for all positions in V_0 as these positions cannot be reached when edge d is played; this is the main reason why we represent strategies as trees. Figure 1b depicts such a tree representation of a strategy. This strategy is a typical example of a winning strategy that is not positional, but still minimal if we take edge multiplicities (see Definition 3.1) into account. \square

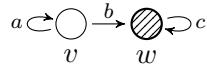
3. STRATEGIES WITH MINIMAL EFFORT

*La perfection est atteinte, non pas lorsqu'il n'y a plus rien à ajouter, mais lorsqu'il n'y a plus rien à retirer.*¹ — ANTOINE DE SAINT-EXUPÉRY

As a measure for the complexity or effort of a strategy, we consider the set of edges a strategy \mathcal{S} uses and how often each of these edges appears in the strategy tree. Under this measure, the simplest strategies are the ones that do not play redundant edges – hence no moves are left to take away.

Definition 3.1. Given an edge $e = vw \in E$ in a Büchi game \mathcal{G} and a strategy \mathcal{S} in \mathcal{G} , we denote by $\#_e(\mathcal{S}) = |\{\rho v \in \mathcal{S} \mid \rho v \rightarrow \rho vw \text{ is an edge in } \mathcal{S}\}| \in \mathbb{N} \cup \{\infty\}$ the number of times (possibly infinite) the edge e occurs in \mathcal{S} . With each strategy \mathcal{S} we associate its *edge profile*, the vector $\#_E(\mathcal{S}) = (\#_e(\mathcal{S}))_{e \in E}$.

Example 3.2. Consider the following Büchi game:



Player 0 wins by first looping n times at position v (for any fixed $n \in \mathbb{N}$) and then moving to w , corresponding to the edge profile $(n, 1, \infty)$. Clearly, looping at v is a redundant move, so we regard the strategy with $n = 0$ as the simplest one (that wins with the least effort). \square

To formalize the intuition of redundant moves, we define an order \succeq on strategies called *absorption*. This is defined in such a way that the \succeq -maximal strategies are the simplest ones that avoid redundant moves whenever possible.

Definition 3.3. Let $\mathcal{S}_1, \mathcal{S}_2$ be two strategies in a Büchi game $\mathcal{G} = (V, V_0, V_1, E, F)$. We say that \mathcal{S}_1 *absorbs* \mathcal{S}_2 , denoted $\mathcal{S}_1 \succeq \mathcal{S}_2$, if $\#_e(\mathcal{S}_1) \leq \#_e(\mathcal{S}_2)$ for all edges $e \in E$. If additionally $\#_e(\mathcal{S}_1) < \#_e(\mathcal{S}_2)$ for some $e \in E$, we say that \mathcal{S}_1 *strictly absorbs* \mathcal{S}_2 , denoted $\mathcal{S}_1 \succ \mathcal{S}_2$. They are *absorption-equivalent*, denoted $\mathcal{S}_1 \equiv \mathcal{S}_2$, if both $\mathcal{S}_1 \succeq \mathcal{S}_2$ and $\mathcal{S}_2 \succeq \mathcal{S}_1$.

A strategy $\mathcal{S} \in \text{Strat}(v)$ is *absorption-dominant from position v* , if there is no strategy $\mathcal{S}' \in \text{Strat}(v)$ with $\mathcal{S}' \succ \mathcal{S}$. It is further *strictly absorption-dominant*, if there is no other strategy $\mathcal{S}' \in \text{Strat}(v)$ with $\mathcal{S}' \succeq \mathcal{S}$, so no other strategy is absorption-equivalent to \mathcal{S} .

Notice that absorption is simply the inverse pointwise order on the edge profiles. In particular, $\mathcal{S}_1 \equiv \mathcal{S}_2$ if, and only if, $\#_E(\mathcal{S}_1) = \#_E(\mathcal{S}_2)$. We next aim at understanding the relation between (strictly) absorption-dominant strategies and the standard notion of positional strategies. As a starter, we show that absorption-dominant strategies are not necessarily positional (cf. [GT20] for a similar example).

Example 3.4. Consider the strategy \mathcal{S} as depicted in Fig. 1b. It is not positional, as the choice for position v' is not unique (both e and f occur in \mathcal{S}). It is, however, absorption-dominant. As there are two paths to v' , every strategy must either use e or f twice, or use both edges. If e (or f) is used twice, then the strategy cannot absorb \mathcal{S} , and one can verify that \mathcal{S} absorbs all strategies using both e and f .

It is not strictly absorption-dominant, as we obtain an absorption-equivalent strategy by switching the two branches in the depiction of \mathcal{S} , so that e is used after c , and f after b . \square

¹Perfection is achieved, not when there is nothing more to add, but when there is nothing left to take away.

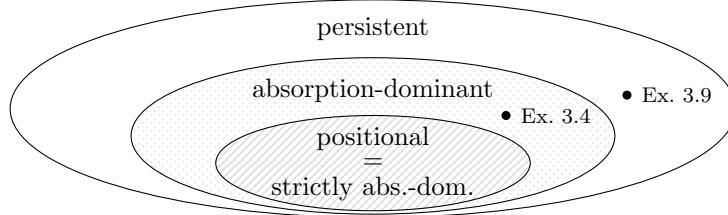


FIGURE 2. Venn diagram depicting various classes of winning strategies.

Strategies such as the one in Fig. 1b are not positional, but satisfy the weaker property that within each *play*, the strategy makes a unique decision for each position $v \in V_0$. This notion of strategies has been introduced as *persistent* strategies in [MT02] in the context of LTL on game graphs and has been further studied in [Dup03]. Persistent strategies have also been called *weakly positional* in [GT20].

We say that a strategy *plays positionally* from a position $v \in V_0$ if the strategy makes a unique choice at position v (not depending on the history of the play). A strategy that plays positionally from all positions in V_0 is positional. With this notation, we now clarify the relation between the different notions of strategies; a summary is shown in Fig. 2. We first observe that if a strategy \mathcal{S} does not play positionally from v , we can always obtain a strategy \mathcal{S}' with $\mathcal{S}' \succeq \mathcal{S}$ by swapping the choices at v , which leads to:

Proposition 3.5. *Strictly absorption-dominant strategies coincide with positional strategies.*

Proof. Let $\mathcal{S} \in \text{Strat}_{\mathcal{G}}(v)$ be a strategy from v . First assume towards a contradiction that \mathcal{S} is positional but not strictly absorption-dominant. That is, there is a different strategy $\mathcal{S}' \in \text{Strat}(v)$ with $\mathcal{S}' \succeq \mathcal{S}$. Since \mathcal{S}' is different from \mathcal{S} , there is a position w and a path ρw occurring in both strategies for which the strategies differ, i.e., we have $w_1 = \mathcal{S}(\rho w)$ and $w_2 = \mathcal{S}'(\rho w)$ with $w_1 \neq w_2$. Since \mathcal{S} is positional, the edge ww_2 does not occur in \mathcal{S} . Hence it occurs strictly more often in \mathcal{S}' , contradicting the assumption $\mathcal{S}' \succeq \mathcal{S}$.

We prove the other direction by contraposition. Let \mathcal{S} be non-positional, so there is a position w and two paths ρw and $\rho' w$ such that $\mathcal{S}(\rho w) \neq \mathcal{S}(\rho' w)$. Let $\mathcal{S}_{\rho w}$ and $\mathcal{S}_{\rho' w}$ be the substrategies of \mathcal{S} from ρw and $\rho' w$, respectively. First assume that $\rho \sqsubseteq \rho'$. We then consider the strategy \mathcal{S}' that behaves like \mathcal{S} , but switches to $\mathcal{S}_{\rho' w}$ at ρw . As every edge occurring in \mathcal{S}' also occurs in \mathcal{S} , we have $\mathcal{S}' \succeq \mathcal{S}$ and \mathcal{S} is not strictly absorption-dominant. The case $\rho' \sqsubseteq \rho$ is symmetric. If, on the other hand, ρ and ρ' are incomparable nodes in \mathcal{S} , we consider the strategy \mathcal{S}' that behaves like \mathcal{S} , but plays $\mathcal{S}_{\rho' w}$ from ρw and $\mathcal{S}_{\rho w}$ from $\rho' w$, swapping the two substrategies. Then $\mathcal{S}' \equiv \mathcal{S}$, so \mathcal{S} is not strictly absorption-dominant. \square

We next establish the relation to persistent strategies. To this end, we first show under which circumstances absorption-dominant strategies must make unique choices. Our proof needs the following combinatorial observation.

Lemma 3.6. *Let $v \in \mathcal{G}$ be a position. There are only finitely many absorption-dominant strategies from v up to absorption-equivalence.*

Proof. Consider the pointwise order on edge profiles induced by the standard order on $\mathbb{N} \cup \{\infty\}$. By definition, a strategy \mathcal{S} is absorption-dominant from v if, and only if, its edge profile $\#_E(\mathcal{S})$ is minimal among all strategies from v (and absorption-equivalent strategies have the same edge profile). By a simple combinatorial fact known as Dickson's lemma, every set of edge profiles contains only finitely many minimal elements. \square

Proposition 3.7. *Let $\mathcal{S} \in \text{WinStrat}_{\mathcal{G}}(v)$ be absorption-dominant from v , and let $w \in V_0$ be a position. If w occurs infinitely often in \mathcal{S} , then \mathcal{S} plays positionally from w .*

Proof. Consider the infinitely many substrategies at occurrences of w in \mathcal{S} . By Lemma 3.6, there is one such substrategy \mathcal{S}_w such that infinitely many of the substrategies are absorption-equivalent to \mathcal{S}_w . This means that every edge occurring in \mathcal{S}_w also occurs in infinitely many substrategies and hence infinitely often in the full strategy \mathcal{S} . Notice that \mathcal{S}_w is winning from w , as it is a substrategy of the winning strategy \mathcal{S} . Consider the subgame of \mathcal{G} containing only edges occurring in \mathcal{S}_w . Clearly, Player 0 wins from w (using \mathcal{S}_w) and by positional determinacy, there is thus a positional winning strategy \mathcal{S}_{pos} from w using only edges that occur in \mathcal{S}_w and hence infinitely often in \mathcal{S} .

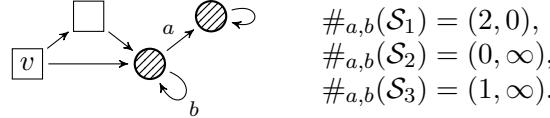
Now consider the strategy $\mathcal{S}' \in \text{WinStrat}_{\mathcal{G}}(v)$ that behaves like \mathcal{S} , but always uses \mathcal{S}_{pos} from w . Then $\mathcal{S}' \succeq \mathcal{S}$ by construction of \mathcal{S}_{pos} . Further, \mathcal{S}_{pos} is positional and makes a unique choice $\mathcal{S}_{\text{pos}}(w)$. If \mathcal{S} would not play positionally from w , then there would be some path ρw such that $\mathcal{S}(\rho w) = w' \neq \mathcal{S}_{\text{pos}}(w)$. But then the edge ww' never occurs in \mathcal{S}' , so $\mathcal{S}' \succ \mathcal{S}$ and \mathcal{S} would not be absorption-dominant. \square

With this important insight, we can deduce that the absorption-dominant winning strategies (from some position v) are a (strict) subset of the persistent strategies: An absorption-dominant strategy must play positionally from positions that occur infinitely often; repetitions of positions that occur finitely often are always redundant.

Corollary 3.8. *Every absorption-dominant winning strategy in \mathcal{G} is persistent.*

Proof. Let $\mathcal{S} \in \text{WinStrat}_{\mathcal{G}}(v)$ be absorption-dominant from v . Assume towards a contradiction that \mathcal{S} is not persistent, so there is a position w and a play of the form $\rho_1 w \rho_2 w \rho_3$ such that \mathcal{S} makes different decisions at w , say $\mathcal{S}(\rho_1 w) = w_1$ and $\mathcal{S}(\rho_1 w \rho_2 w) = w_2$ with $w_1 \neq w_2$. By Proposition 3.7, w can only occur finitely often in \mathcal{S} . Hence the edge ww_1 also occurs finitely often, say n times. Let \mathcal{S}'_w be the substrategy of \mathcal{S} from $\rho_1 w \rho_2 w$. Now consider the strategy $\mathcal{S}' \in \text{WinStrat}_{\mathcal{G}}(v)$ that behaves like \mathcal{S} , but switches to \mathcal{S}'_w at $\rho_1 w$. By construction, \mathcal{S}' uses each edge at most as often as \mathcal{S} . Moreover, one occurrence of the edge ww_1 is removed, so this edge occurs at most $n - 1$ times in \mathcal{S}' . Hence $\mathcal{S}' \succ \mathcal{S}$, contradicting the absorption-dominance of \mathcal{S} . \square

Example 3.9. For strictness, consider the following game (a modified part of Fig. 1a):



Due to the self-loop b , only the positional strategies \mathcal{S}_1 (always take a) and \mathcal{S}_2 (always take b) are absorption-dominant from v . The strategy \mathcal{S}_3 that, depending on Player 1's choice, either takes edge a or loops indefinitely using edge b is persistent, but not absorption-dominant: it is strictly absorbed by \mathcal{S}_2 . \square

As a consequence of Corollary 3.8, all moves after the first repeated position are determined by persistence. We can thus represent absorption-dominant strategies in a compact way and strengthen Lemma 3.6 as follows.

Corollary 3.10. *Let \mathcal{G} be a game with $n = |V|$ positions. Every winning strategy $\mathcal{S} \in \text{WinStrat}_{\mathcal{G}}(v)$ that is absorption-dominant from v can be uniquely represented by a subtree of the tree unraveling of height at most n . In particular, the number of absorption-dominant winning strategies is finite.*

4. A WHIRLWIND TOUR OF SEMIRING SEMANTICS

This section gives an overview on semiring semantics for fixed-point logics, with a focus on the semirings relevant for the case study. For a complete account, we refer to [DGNT21].

4.1. Semirings. Semirings are algebraic structures with two binary operations, usually denoted $+$ and \cdot , which we use to interpret the logical connectives \vee and \wedge . While semirings are very general structures, we make additional assumptions to ensure well-defined and meaningful semiring semantics for logics with fixed-point operators.

Definition 4.1. A *commutative semiring* is an algebraic structure $(\mathbb{K}, +, \cdot, 0, 1)$, with $0 \neq 1$, such that $(\mathbb{K}, +, 0)$ and $(\mathbb{K}, \cdot, 1)$ are commutative monoids, \cdot distributes over $+$, and $0 \cdot a = a \cdot 0 = 0$. It is *idempotent* if $a + a = a$ for all $a \in \mathbb{K}$.

All semirings we consider are commutative, so we omit *commutative* in the following. Towards fixed-point logic, we compute least and greatest fixed points with respect to the natural order $\leq_{\mathbb{K}}$ (see below) and to ensure that they exist, we require $\leq_{\mathbb{K}}$ to be a complete lattice (in fact, suprema and infima of chains would suffice, but in idempotent semirings this is equivalent). We additionally impose a natural *continuity* requirement which is crucial to our proofs, but does not seem to be a strong restriction in practice (we are not aware of any natural complete-lattice semirings that are not continuous). Regarding notation, a *chain* is a totally ordered set $C \subseteq \mathbb{K}$ and we write $a \circ C = \{a \circ c \mid c \in C\}$ for $a \in \mathbb{K}$.

Definition 4.2. In an idempotent semiring $(\mathbb{K}, +, \cdot, 0, 1)$, the *natural order* $\leq_{\mathbb{K}}$ is the partial order defined by $a \leq_{\mathbb{K}} b \Leftrightarrow a + b = b$. We say that \mathbb{K} is *fully continuous* if $\leq_{\mathbb{K}}$ is a complete lattice (with supremum \bigcup and infimum \bigcap) and for all non-empty chains $C \subseteq \mathbb{K}$, elements $a \in \mathbb{K}$ and $\circ \in \{+, \cdot\}$,

$$\bigcup(a \circ C) = a \circ \bigcup C, \quad \text{and} \quad \bigcap(a \circ C) = a \circ \bigcap C.$$

A semiring homomorphism $h: \mathbb{K}_1 \rightarrow \mathbb{K}_2$ on fully-continuous semirings is *fully continuous* if $h(\bigcup C) = \bigcup h(C)$ and $h(\bigcap C) = \bigcap h(C)$ for all non-empty chains $C \subseteq \mathbb{K}_1$.

By the Knaster-Tarski theorem, every $\leq_{\mathbb{K}}$ -monotone function $f: \mathbb{K} \rightarrow \mathbb{K}$ on a fully-continuous semiring has a least fixed point $\mathbf{lfp}(f)$ and a greatest fixed point $\mathbf{gfp}(f)$ in \mathbb{K} , and this suffices to guarantee well-defined semiring semantics of fixed-point logics. However, from a provenance perspective we further want this semantics to be meaningful in the sense that the value of a formula provides insights into why the formula holds. It turns out that this is the case if we additionally require the semiring to be absorptive (see [DGNT21]).

Definition 4.3. A semiring \mathbb{K} is *absorptive* if $a + ab = a$ for all $a, b \in \mathbb{K}$.

We remark that absorption is equivalent to \mathbb{K} being *0-closed* or *bounded* [Moh02], that is, $1+a=1$. If \mathbb{K} is idempotent, then absorption is further equivalent to multiplication being decreasing, that is, $a \cdot b \leq_{\mathbb{K}} a, b$. Clearly, every absorptive semiring is idempotent and thus partially ordered by $\leq_{\mathbb{K}}$, with 1 as top element. If we additionally assume full continuity, we can extend any absorptive semiring by an infinitary power operation $a^\infty = \bigcap_{n \in \mathbb{N}} a^n$ with natural properties such as $a \cdot a^\infty = a^\infty$, $(ab)^\infty = a^\infty b^\infty$ and $(a+b)^\infty = a^\infty + b^\infty$.

Example 4.4. Here is a short, non-exhaustive list of semirings used in provenance analysis of databases and logics [GKT07, GT17b, GT20].

- The *Boolean semiring* $\mathbb{B} = (\{\mathbf{0}, \mathbf{1}\}, \vee, \wedge, \mathbf{0}, \mathbf{1})$ is the standard habitat of logical truth. It is absorptive and (trivially) fully continuous.
- $\mathbb{N} = (\mathbb{N}, +, \cdot, 0, 1)$ is used for counting evaluation strategies for a logical statement. It is not absorptive and hence not well suited for fixed-point logics.
- The *Viterbi semiring* $\mathbb{V} = ([0, 1], \max, \cdot, 0, 1)$ is used to compute *confidence scores* of logical statements. It is isomorphic to the *tropical semiring* $\mathbb{T} = (\mathbb{R}_+^\infty, \min, +, \infty, 0)$ used for measuring the cost of evaluation strategies. Both are absorptive and fully continuous.
- The *min-max semiring* (A, \max, \min, a, b) on a totally ordered set (A, \leq) with least and greatest elements a and b can be used to model access privileges. It is absorptive and fully continuous. \square

From now on, all semirings we consider are commutative, absorptive and fully continuous. Besides the application semirings listed above, we are particularly interested in universal semirings of polynomials to represent abstract information. We can then use fully-continuous homomorphisms to specialize the computed information to application semirings as needed, as these homomorphisms preserve fixed points. The common examples of semirings of polynomials $\mathbb{N}[X]$ and formal power series $\mathbb{N}^\infty[[X]]$, as used for provenance analysis of FO and Datalog in [GKT07, GT17a], are not absorptive and hence not well-suited for fixed-point logic.

Instead, we rely on semirings of (generalized²) *absorptive polynomials* (cf. [DGNT21]). Essentially, an absorptive polynomial such as $ab^3 + c^\infty$ is a sum of monomials over a finite set of variables X , but without coefficients and with exponents from $\mathbb{N} \cup \{\infty\}$. Monomial multiplication is defined as usual by adding exponents (with $n + \infty = \infty$). The key ingredient is absorption among monomials: we say that a monomial m_1 *absorbs* m_2 ($m_1 \succeq m_2$), if all its exponents are smaller (or equal). Formally, $m_1 \succeq m_2$ if $m_1(x) \leq m_2(x)$ for all $x \in X$, where $m_1(x)$ denotes the exponent of x in m_1 . For example, $ab^2 \succeq a^\infty b^2$ and $a \succeq ab$, but $a^2 b$ and ab^2 are incomparable. In an absorptive polynomial, we omit all monomials that would be absorbed, so absorptive polynomials are precisely the \succeq -*antichains* of monomials. Consequently, addition and multiplication are defined as usual, but afterwards we only keep the \succeq -maximal monomials. For example, $(ab^2 + a^2 b) \cdot a^\infty = a^\infty b^2 + a^\infty b = a^\infty b$.

Definition 4.5. We write $\mathbb{S}^\infty[X]$ for the semiring of absorptive polynomials over the finite variable set X . Its elements are the \succeq -antichains of monomials (written as sums) with the operations described above. The neutral elements are the empty polynomial (denoted by 0) and the single monomial 1 (with all zero exponents).

This defines an absorptive, fully-continuous semiring [DGNT21]. In fact, $\mathbb{S}^\infty[X]$ is the most general such semiring:

Theorem 4.6 (Universal property, [DGNT21]). *Every mapping $h: X \rightarrow \mathbb{K}$ into an absorptive, fully-continuous semiring \mathbb{K} uniquely extends to a fully-continuous semiring homomorphism $h: \mathbb{S}^\infty[X] \rightarrow \mathbb{K}$ (by means of polynomial evaluation).*

4.2. Logic. We consider here the fixed-point logic LFP that extends first-order logic FO by least and greatest fixed-point formulae of the form $\psi(\mathbf{y}) = [\mathbf{lfp} \, Rx. \, \varphi(R, \mathbf{x})](\mathbf{y})$ and $\psi(\mathbf{y}) = [\mathbf{gfp} \, Rx. \, \varphi(R, \mathbf{x})](\mathbf{y})$. Here, R is a relation symbol occurring only positively in φ and \mathbf{x}, \mathbf{y} are variable tuples of matching arity. Given a (Boolean) model \mathfrak{A} and a tuple \mathbf{a} of

²The definition we use here generalizes the one in [DMRT14] by allowing ∞ as exponent.

elements of \mathfrak{A} , the formula $\psi(\mathbf{a})$ holds in \mathfrak{A} , denoted $\mathfrak{A} \models \psi(\mathbf{a})$, if \mathbf{a} is contained in the least (or greatest) fixed point of the operator $F_\varphi: R \mapsto \{\mathbf{a} \mid \mathfrak{A} \models \varphi(R, \mathbf{a})\}$ that maps a relation R to the relation consisting of those tuples for which φ holds. For more background and a precise definition, we refer to [GKL⁺07].

In order to generalize Boolean semantics to semiring semantics, we first adapt the notion of a model \mathfrak{A} . Instead of determining for each literal whether it is true or false in \mathfrak{A} , we assign to each literal a semiring value, interpreting 0 as *false* and all other values as *nuances of true*. Special care is required to ensure that the assignment is consistent with respect to opposing literals (this is not always necessary, but often desirable). In the following, let \mathbb{K} be a semiring, A a finite universe and τ a relational signature (we drop A and τ if clear from the context). We denote the set of (instantiated) literals as

$$\text{Lit}_{A,\tau} = \{R\mathbf{a}, \neg R\mathbf{a} \mid R \in \tau \text{ of arity } k, \mathbf{a} \in A^k\} \cup \{\mathbf{a} = \mathbf{b}, \mathbf{a} \neq \mathbf{b} \mid \mathbf{a}, \mathbf{b} \in A^k\}.$$

Given a literal α , we write $\neg\alpha$ for the opposing literal (identifying $\neg\neg\alpha$ and α). The role of the Boolean model \mathfrak{A} is then replaced by a semiring interpretation π that assigns semiring values to all literals.

Definition 4.7. Let \mathbb{K} be a semiring. A \mathbb{K} -interpretation (over finite A and τ) is an assignment $\pi: \text{Lit}_{A,\tau} \rightarrow \mathbb{K}$ that maps true (in)equalities to 1 and false (in)equalities to 0. We say that π is *model-defining*, if for each literal α , exactly one of $\pi(\alpha)$ and $\pi(\neg\alpha)$ is 0.

For Büchi games, we will always use the signature $\tau = \{E, F, V_0, V_1\}$, where E is a binary and F, V_0, V_1 are unary relation symbols. We can then view a game $\mathcal{G} = (V, V_0, V_1, E, F)$ as a τ -structure. Notice that we do not distinguish between the edge relation E of \mathcal{G} and the relation symbol E ; it will always be clear from the context what we refer to. The set of instantiated literals $\text{Lit}_{V,\tau}$ then contains, e.g., $E v_1 v_2$ and $\neg F v_1$, where $v_1, v_2 \in \mathcal{G}$.

We lift K -interpretations π from literals to LFP-formulae in negation normal form (nnf), resulting in a semiring value $\pi[\psi]$, by interpreting the logical connectives as semiring operations. For fixed-point formulae, we consider the induced operator F_φ analogous to the Boolean case, but acting on functions $g: A^k \rightarrow \mathbb{K}$ instead of relations $R \subseteq A^k$ (which can be seen as functions $R: A^k \rightarrow \mathbb{B}$, justifying our generalisation). We extend the natural order to such functions by pointwise comparison. More formally, given a \mathbb{K} -interpretation π over signature τ , we denote by $\pi[R \mapsto g]$ the \mathbb{K} -interpretation over $\tau \cup \{R\}$ obtained from π by adding values $g(\mathbf{a})$ for the instantiated atoms $R\mathbf{a}$. The analogue of the Boolean operator F_φ is then the operator F_π^φ that maps a function $g: A^k \rightarrow \mathbb{K}$ to the function

$$F_\pi^\varphi(g): \mathbf{a} \mapsto \pi[R \mapsto g][\varphi(R, \mathbf{a})].$$

With this in mind, we define the following natural generalization of Boolean semantics.

Definition 4.8. A \mathbb{K} -interpretation $\pi: \text{Lit}_A(\tau) \rightarrow \mathbb{K}$ (for finite A and τ) in a fully-continuous semiring \mathbb{K} extends to a \mathbb{K} -valuation $\pi: \text{LFP}(\tau) \rightarrow \mathbb{K}$ by mapping an LFP-sentence $\psi(\mathbf{a})$ in negation normal form to a value $\pi[\psi]$ using the rules

$$\begin{aligned} \pi[\psi \vee \varphi] &:= \pi[\psi] + \pi[\varphi], & \pi[\exists x \psi(x)] &:= \sum_{a \in A} \pi[\varphi(a)], & \pi[\neg\psi] &:= \pi[\text{nnf}(\neg\psi)], \\ \pi[\psi \wedge \varphi] &:= \pi[\psi] \cdot \pi[\varphi], & \pi[\forall x \psi(x)] &:= \prod_{a \in A} \pi[\varphi(a)], \end{aligned}$$

and, for fixed-point formulae,

$$\pi[[\text{lfp } Rx.\varphi(R, \mathbf{x})](\mathbf{a})] := \text{lfp}(F_\pi^\varphi)(\mathbf{a}), \quad \pi[[\text{gfp } Rx.\varphi(R, \mathbf{x})](\mathbf{a})] := \text{gfp}(F_\pi^\varphi)(\mathbf{a}).$$

Interpretation & application	π_{strat} <i>strategy tracking</i>	π_{repair} <i>reverse analysis of moves</i>	π_{target} <i>target synthesis</i>
Semiring	$\mathbb{S}^\infty[X]$	$\mathbb{S}^\infty[X, \bar{X}]$ or $\text{PosBool}[X, \bar{X}]$	$\text{PosBool}[X, \bar{X}]$
$\pi(Evw)$	$X_{vw}/0$	\bar{X}_{vw} (or 1/0)	1/0
$\pi(\neg Evw)$	1/0	\bar{X}_{vw} (or 1/0)	1/0
$\pi(Fv)$	1/0	1/0	X_v
$\pi(\neg Fv)$	1/0	1/0	\bar{X}_v
other literals	1/0	1/0	1/0

FIGURE 3. Semiring interpretations used in this paper (the notation $a/0$ indicates the value a if the literal is true, and 0 if it is false in \mathcal{G}).

An important property of the resulting semantics is that it is preserved by fully-continuous semiring homomorphisms, in particular by polynomial evaluation in $\mathbb{S}^\infty[X]$ due to Theorem 4.6 (but not by polynomial evaluation of $\mathbb{N}[X]$ or formal power series!).

Theorem 4.9 (Fundamental property, [DGNT21]). *Let $h: \mathbb{K}_1 \rightarrow \mathbb{K}_2$ be a fully-continuous semiring homomorphism. Then for every \mathbb{K}_1 -interpretation π , the mapping $h \circ \pi$ is a \mathbb{K}_2 -interpretation and $h(\pi[\varphi]) = (h \circ \pi)[\varphi]$, for every $\varphi \in \text{LFP}$.*

$$\begin{array}{ccc}
 \begin{array}{c} \text{Lit}_A(\tau) \\ \pi \swarrow \quad \searrow h \circ \pi \\ K_1 \xrightarrow{h} K_2 \end{array} & \Rightarrow & \begin{array}{c} \text{LFP} \\ \pi \swarrow \quad \searrow h \circ \pi \\ K_1 \xrightarrow{h} K_2 \end{array}
 \end{array}$$

5. COMPUTING STRATEGIES WITH SEMIRING SEMANTICS

This section connects the previous sections on semiring semantics and absorption-dominant strategies. We focus on the formula for the winning region in a Büchi game and show that its value under semiring semantics can be understood in terms of (absorption-dominant) winning strategies.

5.1. The Semiring Interpretation. We want to use semiring semantics to analyze moves in winning strategies. For this reason, we label edges with indeterminates X (cf. Fig. 1a) and use an $\mathbb{S}^\infty[X]$ -interpretation π_{strat} to track moves (i.e., edge literals Evw for positions v, w) via their indeterminates. We assume the game graph to be fixed and do not wish to track information about the target set F or the active player at a certain node, hence we simply map all other literals over $\tau = \{E, F, V_0, V_1\}$, such as Fv , V_0v and $\neg Evw$, to 0 or 1, depending on whether they are true or false in the fixed game. The resulting interpretation is almost Boolean and hence behaves very similar to the original game, except that we remember which edges are used in the evaluation of a formula.

Definition 5.1. Let $\mathcal{G} = (V, V_0, V_1, E, F)$ be a Büchi game and let $X = \{X_{vw} \mid vw \in E\}$ be a set of indeterminates for all edges. We define the $\mathbb{S}^\infty[X]$ -interpretation $\pi_{\text{strat}}: \text{Lit}_{V, \tau} \rightarrow$

$\mathbb{S}^\infty[X]$ as follows (depending on \mathcal{G}):

$$\begin{aligned}\pi_{\text{strat}}(Evw) &= X_{vw} \text{ for all edges } vw \in E, \\ \pi_{\text{strat}}(\alpha) &= \begin{cases} 1, & \text{if } \mathcal{G} \models \alpha, \\ 0, & \text{if } \mathcal{G} \not\models \alpha, \end{cases} \text{ for all other literals } \alpha \in \text{Lit}_{V,\tau}.\end{aligned}$$

For the applications in Section 6, we may consider other interpretations which are defined in a similar way, but also track negative edge literals or the target set F . An overview is given in Fig. 3. In this section, we always work with π_{strat} .

5.2. The Formula. It is well known that the winning region (of Player 0) in a Büchi game is definable in fixed-point logic. Intuitively, the winning region is the largest set Y such that from each position in Y , Player 0 can enforce a visit to $Y \cap F$ (after at least one move). In LFP, we can express the winning region as follows (see, e.g., [CGLP15, Wal02]):

$$\text{win}_0(x) := [\mathbf{gfp} Yy. [\mathbf{lfp} Zz. \varphi(Y, Z, z)](y)](x),$$

where

$$\begin{aligned}\varphi(Y, Z, z) := & \left(Fz \wedge ((V_0z \wedge \exists u(Ezu \wedge Yu)) \vee (V_1z \wedge \forall u(Ezu \rightarrow Yu))) \right) \\ & \vee \left(\neg Fz \wedge ((V_0z \wedge \exists u(Ezu \wedge Zu)) \vee (V_1z \wedge \forall u(Ezu \rightarrow Zu))) \right).\end{aligned}$$

Given a \mathbb{K} -interpretation π for a Büchi game $\mathcal{G} = (V, V_0, V_1, E, F)$, semiring semantics of the above formula induces³ the following fixed-point computation. To simplify the presentation, we introduce two families of variables, $\mathbf{Y} = (Y_v)_{v \in V}$ and $\mathbf{Z} = (Z_v)_{v \in V}$ that take values in \mathbb{K} . We can then express the resulting semiring valuation as $\pi[\text{win}_0(v)] = Y_v^*$ where $\mathbf{Y}^* = (Y_v^*)_{v \in V}$ is the *greatest* solution to the equation system

$$\mathbf{Y} = \mathbf{Z}^*(\mathbf{Y})$$

where, in turn, $\mathbf{Z}^*(\mathbf{Y})$ is the *least* solution, given values $\mathbf{Y} = (Y_v)_{v \in V}$, to the equation system consisting of the following equation for all $v \in V$:

$$\begin{aligned}Z_v &= \pi(Fv) \cdot \left((\pi(V_0v) \cdot \sum_{w \in V} (\pi(Evw) \cdot Y_w)) + (\pi(V_1v) \cdot \prod_{w \in V} (\pi(\neg Evw) + \pi(Evw) \cdot Y_w)) \right) \\ &+ \pi(\neg Fv) \cdot \left((\pi(V_0v) \cdot \sum_{w \in V} (\pi(Evw) \cdot Z_w)) + (\pi(V_1v) \cdot \prod_{w \in V} (\pi(\neg Evw) + \pi(Evw) \cdot Z_w)) \right).\end{aligned}$$

For most of this paper, we use π_{strat} to track only *moves* of winning strategies. As π_{strat} maps most of the literals to 0 or 1, we can simplify the equations depending on v :

	$v \in F$	$v \notin F$
$v \in V_0$	$Z_v = \sum_{w \in vE} \pi(Evw) \cdot Y_w$	$Z_v = \sum_{w \in vE} \pi(Evw) \cdot Z_w$
$v \in V_1$	$Z_v = \prod_{w \in vE} \pi(Evw) \cdot Y_w$	$Z_v = \prod_{w \in vE} \pi(Evw) \cdot Z_w$

³Here we first translate $Ezu \rightarrow Yu$ to the formula $\neg Ezu \vee (Ezu \wedge Yu)$ in negation normal form.

A good way to think about (and compute) the least and greatest solutions is the fixed-point iteration. The idea is to start with each Z_v set to the least element 0 of the semiring, then apply the above equations (i.e., the induced operator $F_{\pi_{\text{strat}}}^\varphi$) to compute a next, larger semiring value and repeat this process until a fixed-point is reached (notice that the iteration can also be infinite, the fixed-point is then the supremum or infimum).

Example 5.2. Recall the simple game from Example 3.2 $(\xrightarrow[a]{v} \bigcirc \xrightarrow[b]{w} \bigcirc \xrightarrow[c]{})$.

Using the interpretation π_{strat} corresponding to the edge labels, we obtain the following fixed-point iteration. We write the tuples \mathbf{Y} and \mathbf{Z} as vectors $(Y_v \ Y_w)^T$ and $(Z_v \ Z_w)^T$.

$$\begin{array}{ccccccc} \mathbf{Y} : & \begin{pmatrix} 1 \\ 1 \end{pmatrix} & & \begin{pmatrix} bc \\ c \end{pmatrix} & & \begin{pmatrix} bc^2 \\ c^2 \end{pmatrix} & \begin{pmatrix} bc^n \\ c^n \end{pmatrix} \\ & \searrow & & \downarrow & & \downarrow & \downarrow \\ \mathbf{Z} : & \begin{pmatrix} 0 \\ 0 \end{pmatrix} & \xrightarrow{F^\varphi} & \begin{pmatrix} 0 \\ c \end{pmatrix} & \xrightarrow{F^\varphi} & \begin{pmatrix} 0 \\ c^2 \end{pmatrix} & \xrightarrow{F^\varphi} \dots & \dots \end{array}$$

We obtain the overall result $\pi_{\text{strat}}[\text{win}_0(v)] = Y_v^* = \prod_n bc^n = bc^\infty$ corresponding to the unique absorption-dominant strategy using edge b once and c infinitely often (cf. Example 3.2). \lhd

5.3. Connection to Strategies. By mapping edges to semiring values, we can track edges through the fixed-point computation. In Example 5.2, the resulting semiring value revealed how often each edge is used in the unique absorption-dominant winning strategy. We now generalize this observation. For simplicity, we only consider \mathbb{K} -interpretations π that are *edge tracking* for a given game \mathcal{G} . That is, they may assign arbitrary values to positive edge literals Evw , but all other literals are mapped to 0 or 1 in accordance with \mathcal{G} . To make the connection to strategies explicit, we first define semiring values for strategies based on the appearance of edges.

Definition 5.3. Let \mathcal{S} be a strategy in a Büchi game $\mathcal{G} = (V, V_0, V_1, E, F)$. Let \mathbb{K} be an absorptive, fully-continuous semiring and π an edge-tracking \mathbb{K} -interpretation on \mathcal{G} . The \mathbb{K} -value of \mathcal{S} is the product of the values for all edges appearing in \mathcal{S} . Formally,

$$\pi[\mathcal{S}] := \prod_{vw \in E} \pi(Evw)^{\#_{vw}(\mathcal{S})},$$

where infinite exponents are interpreted by the infinitary power operation of the semiring.

The semiring value of win_0 can then be expressed as the sum over the values of all winning strategies. A direct proof is not completely straightforward, as fixed-point iterations and strategy trees can both be infinite (even if \mathcal{G} is finite). Instead, we make use of a similar sum-of-strategies result for model-checking games for LFP (see [DGNT21]).

Theorem 5.4 (Sum of Strategies). *Let \mathcal{G} be a Büchi game and v a position in \mathcal{G} . Let \mathbb{K} be an absorptive, fully-continuous semiring and π an edge-tracking \mathbb{K} -interpretation. Then,*

$$\pi[\text{win}_0(v)] = \sum \{ \pi[\mathcal{S}] \mid \mathcal{S} \in \text{WinStrat}_{\mathcal{G}}(v) \text{ is absorption-dominant from } v \}.$$

It is in fact this central result that motivated the notion of *absorption-dominant* strategies. However, as we have already discussed, these may also be interesting in their own right if one is interested in minimal winning strategies.

Example 5.5. For the edge-tracking interpretation π_{strat} induced by the edge labels in Fig. 1a, we obtain

$$\pi_{\text{strat}}[\![\text{win}_0(v)]\!] = (abcd)^\infty + abc e^2 h^2(gkm)^\infty + abc f^2(gkm)^\infty + abc e f h(gkm)^\infty.$$

There are four monomials, corresponding to four equivalence classes of absorption-dominant strategies. Each monomial reveals the edges that appear in the corresponding strategies, so we see that the first three monomials belong to positional (and hence uniquely defined) strategies. The last monomial belongs to the non-positional strategy shown in Fig. 1b (and its switched version, see Example 3.4). The values of all other strategies are strictly absorbed by one of these monomials. \square

5.4. Proof of the Sum-of-Strategies Theorem. To prove Theorem 5.4, we show that it follows from a more general sum-of-strategies theorem [DGNT21, Thm. 23]⁴. The general theorem expresses the semiring semantics of arbitrary LFP-formula by means of winning strategies in the corresponding model-checking game; this does not immediately apply to Büchi games, but it is not difficult to see (and we explain below) that the model-checking game for $\text{win}_0(v)$ on a Büchi game \mathcal{G} has the same structure as the original game graph \mathcal{G} .

Towards the proof, we first sketch the model-checking game $\text{MC}(\mathcal{G}, v)$ of $\text{win}_0(v)$ on a Büchi game \mathcal{G} . We then briefly recapitulate how semiring values of winning strategies are defined in [DGNT21], and we show that we can simplify $\text{MC}(\mathcal{G}, v)$ without changing these values so that it becomes almost identical to \mathcal{G} . This allows us to derive Theorem 5.4 from the general sum-of-strategies result on $\text{MC}(\mathcal{G}, v)$.

Model-Checking Game. In general, the model-checking game of an LFP-formula ψ in a structure \mathfrak{A} is a parity game whose positions are pairs of a subformula $\varphi(\mathbf{x})$ of ψ and an instantiation of the free variables, conveniently written as $\varphi(\mathbf{a})$ for some tuple $\mathbf{a} \subseteq \mathfrak{A}$. The positions belong either to Verifier (who wants to prove $\mathfrak{A} \models \varphi(\mathbf{a})$) or to Falsifier (who wants to prove $\mathfrak{A} \not\models \varphi(\mathbf{a})$). Edges allow the current player to move from a position $\varphi(\mathbf{a})$ to a direct subformula of $\varphi(\mathbf{a})$, or from fixed-point literals (such as Yu in our case) back to the entire fixed-point formula. Terminal positions arise from literals $\varphi(\mathbf{a})$ and are won by Verifier precisely if $\mathfrak{A} \models \varphi(\mathbf{a})$. The model-checking game we are interested in is shown in Fig. 4, so we refrain from a complete definition, see e.g. [AG11, Chap. 4] for more background.

Here we are only concerned with the model-checking game of $\text{win}_0(v)$ in a Büchi game \mathcal{G} , and we denote this game as $\text{MC}(\mathcal{G}, v)$. Since $\text{win}_0(v)$ only has a single alternation of fixed-point operators, $\text{MC}(\mathcal{G}, v)$ has only two priorities and can thus be equivalently represented as a Büchi game (with Verifier as Player 0). Recall the formula from Section 5.2 (split into subformulae to ease referencing and with implication already rewritten):

$$\text{win}_0(x) := [\mathbf{gfp} \, Yy. [\mathbf{lfp} \, Zz. \varphi(Y, Z, z)](y)](x),$$

where

$$\varphi(z) := (Fz \wedge \vartheta_1(z)) \vee (\neg Fz \wedge \vartheta_2(z)),$$

$$\vartheta_1(z) := ((V_0 z \wedge \exists u(Ezu \wedge Yu)) \vee (V_1 z \wedge \forall u(\neg Ezu \vee (Ezu \wedge Yu)))),$$

$$\vartheta_2(z) := ((V_0 z \wedge \exists u(Ezu \wedge Zu)) \vee (V_1 z \wedge \forall u(\neg Ezu \vee (Ezu \wedge Zu)))),$$

⁴A complete proof can be found in the full version [DGNT19, Section 6].

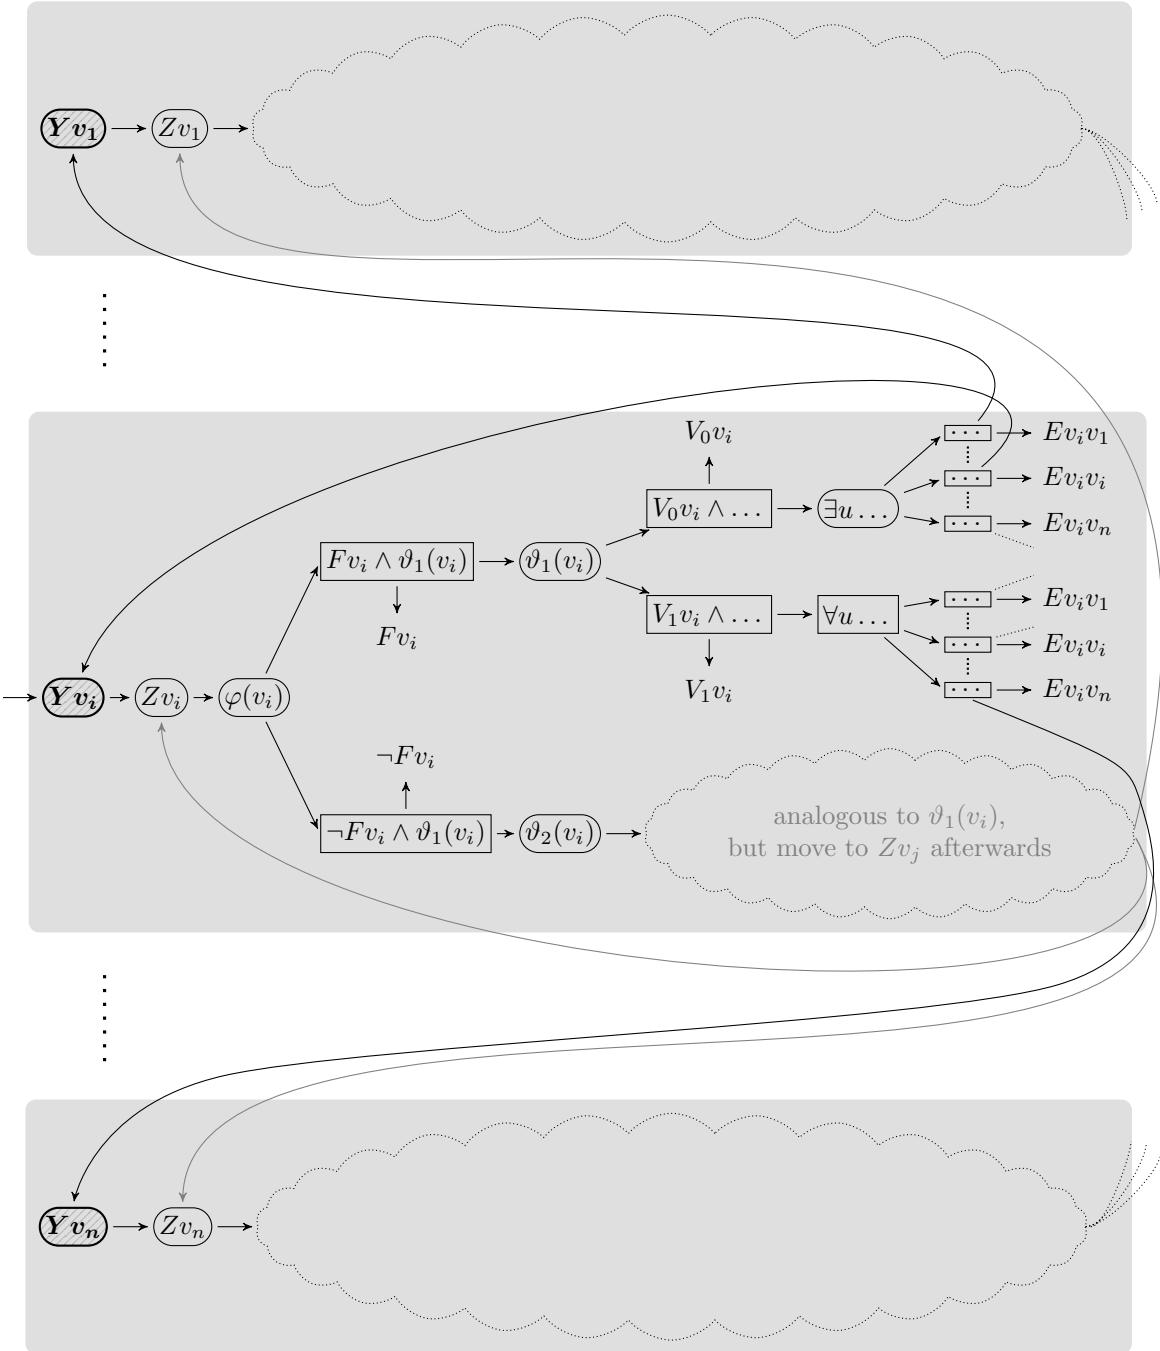


FIGURE 4. Illustration of the model-checking game $MC(\mathcal{G}, v_i)$ for a Büchi game with positions $V = \{v_1, \dots, v_n\}$. Rounded nodes belong to Verifier, rectangular nodes to Falsifier. Nodes without border are terminal positions representing literals, dashed nodes are target positions.

A depiction of the complete model-checking game $\text{MC}(\mathcal{G}, v)$ is shown in Fig. 4 (with some unavoidable omissions due to space reasons).

Semiring Values for the Model-Checking Game. We define strategies for $\text{MC}(\mathcal{G}, v)$ as we did for Büchi games, always taking the perspective of Verifier. To avoid confusion, we use the letter \mathcal{M} to denote strategies in the model-checking game, and \mathcal{S} for strategies in \mathcal{G} . We say that a strategy \mathcal{M} is *winning* if on each infinite play, target nodes (i.e., nodes of the form Yv) occur infinitely often; we impose no restrictions on finite plays (which end in literals). In the following, we only consider strategies for $\text{MC}(\mathcal{G}, v)$ that are winning.

Now consider literals, the terminal positions of the model-checking game. As in Definition 4.7, we write Lit (omitting V and τ) for the set of instantiated literals over the signature of win_0 and the set of positions of \mathcal{G} . Given a \mathbb{K} -interpretation π , it assigns to each literal $\alpha \in \text{Lit}$ the value $\pi(\alpha)$. Following [DGNT21], we define the value of a winning strategy \mathcal{M} by counting literals.

Definition 5.6. Let π be a \mathbb{K} -interpretation in an absorptive, fully-continuous semiring. Let \mathcal{G} be a Büchi game and \mathcal{M} a winning strategy in $\text{MC}(\mathcal{G}, v)$. For a literal $\alpha \in \text{Lit}$, we write $\#\alpha(\mathcal{M}) \in \mathbb{N} \cup \{\infty\}$ for the number of occurrences of α in \mathcal{M} (represented as a strategy tree). The \mathbb{K} -value of \mathcal{M} is

$$\pi[\mathcal{M}] := \prod_{\alpha \in \text{Lit}} \pi(\alpha)^{\#\alpha(\mathcal{M})}.$$

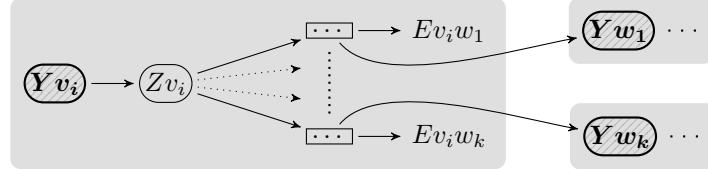
Notice that for the Boolean interpretation $\pi: \text{Lit} \rightarrow \mathbb{B}$ that maps all literals to truth values according to \mathcal{G} , we have $\pi[\mathcal{M}] \neq 0$ if, and only if, \mathcal{M} is a winning strategy (in the classical sense) in the standard Boolean model-checking game for $\text{win}_0(v)$ and \mathcal{G} . By assigning non-Boolean values for the literals, in particular indeterminates in $\mathbb{S}^\infty[X]$, we can track how the literals affect the truth of $\text{win}_0(v)$. With this terminology, applying the sum-of-strategies theorem in [DGNT21] to $\text{win}_0(v)$ gives:

Theorem 5.7. *Let π be a \mathbb{K} -interpretation into an absorptive, fully-continuous semiring. Then $\pi[\text{win}_0(v)] = \sum \{\pi[\mathcal{M}] \mid \mathcal{M} \text{ is a winning strategy in } \text{MC}(\mathcal{G}, v)\}$.*

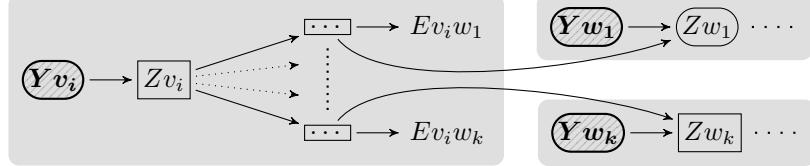
In order to use this result in our context, we will show that the winning strategies \mathcal{M} in $\text{MC}(\mathcal{G}, v)$ correspond to winning strategies \mathcal{S} in \mathcal{G} . For edge-tracking interpretations π , only the edge literals Evw are relevant and thus the corresponding strategies \mathcal{M} and \mathcal{S} have the same \mathbb{K} -value, so Theorem 5.7 will then imply Theorem 5.4.

From the Model-Checking Game to the Büchi Game. Let π be an edge-tracking \mathbb{K} -interpretation for \mathcal{G} , so most of the literals are mapped to 0 or 1. We can then remove the corresponding terminal positions in $\text{MC}(\mathcal{G}, v)$ and in some cases also their predecessors. For instance, consider a position $\varphi(\mathbf{a})$ in $\text{MC}(\mathcal{G}, v)$ from which Falsifier can move to a literal α with $\pi(\alpha) = 0$. Then every strategy \mathcal{M} that visits $\varphi(\mathbf{a})$ must also visit α , thus having value $\pi[\mathcal{M}] = 0$, so we can ignore this strategy for the sum in Theorem 5.7. Hence, replacing the position $\varphi(\mathbf{a})$ by its successor α does not change the sum. On the other hand, if $\pi(\alpha) = 1$, then visiting α does not affect the \mathbb{K} -value of \mathcal{M} and hence we can remove α . Similar reasoning applies to positions of Verifier. Moreover, we can always skip over non-target positions with a unique successor, as they neither affect gameplay nor the \mathbb{K} -values of strategies.

With these insights, we can simplify the model-checking game in Fig. 4 quite a bit. If, say, $v_i \in F$ and $v_i \in V_0$, then the central part of the picture simplifies to:



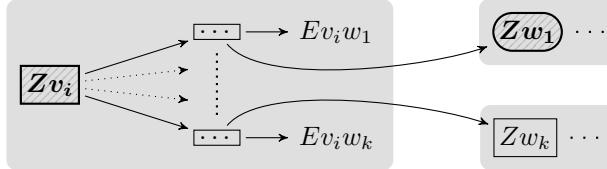
Here, $v_i E = \{w_1, \dots, w_k\}$ are the successors of v_i in \mathcal{G} , so Verifier's moves from Zv_i are only those corresponding to actual edges of \mathcal{G} . The other situations are similar, here is the case $v_i \notin F$, $v_i \in V_1$:



where again $\{w_1, \dots, w_k\}$ are the successors of v_i in \mathcal{G} . Notice that Zv_i belongs to Falsifier (as a result of skipping several positions). In general, Zv_i belongs to Verifier precisely if v_i belongs to Player 0 in \mathcal{G} .

We can now identify the entire subgraph from Yv_i up to the edge literals Ev_iw_j (that is, the gray rectangle) with the position v_i in \mathcal{G} , and we call this the *gadget for v_i* . Indeed, if $v_i \in V_0$ then Verifier chooses a successor $w \in v_i E$ and moves to the corresponding gadget, in analogy to Player 0 choosing a successor in \mathcal{G} . Similarly, Falsifier chooses a successor if it is Player 1's turn in \mathcal{G} .

What remains to discuss are the target positions. Each gadget has two entry points Yv and Zv , and only Yv is a target position. Notice that when we move from a gadget for v to a gadget for w , we use the entry point Yw if, and only if, $v \in F$ (where F is the target set of the original game \mathcal{G}). Hence any play that visits infinitely many target positions Yw also visits infinitely many gadgets for positions $v \in F$ (the predecessors of w). We can thus change the target set without affecting the winning strategies: Instead of the positions Yw , we set the target set to $\{Zv \mid v \in F\}$. The positions Yw are then regular positions with unique successors and can thus be removed. As an example, say we have $v_i \in F$, $w_1 \in F$ and $w_k \notin F$. The previous picture then becomes:



Proof of Theorem 5.4. Let $\widetilde{\text{MC}}(\mathcal{G}, v)$ be the game that results from $\text{MC}(\mathcal{G}, v)$ by applying all of the above-mentioned simplifications. Hence $\widetilde{\text{MC}}(\mathcal{G}, v)$ contains for each position v a gadget with unique entry point Zv that belongs to Verifier exactly if $v \in V_0$, and Zv is a target position exactly if $v \in F$. Moreover, the gadget for v is directly connected to the gadget for w if, and only if, the edge vw exists in \mathcal{G} . It is now easy to see that every winning strategy \mathcal{S} in the original Büchi game \mathcal{G} induces a unique winning strategy \mathcal{M} in $\widetilde{\text{MC}}(\mathcal{G}, v)$: whenever \mathcal{S} visits a position v , then \mathcal{M} visits the gadget for v (via the unique entry point Zv). Conversely, every winning strategy \mathcal{M} uniquely induces a winning strategy \mathcal{S} (which moves to v when \mathcal{M} enters the gadget for v), and we thus say that \mathcal{M} and \mathcal{S} are *corresponding* winning strategies.

Lemma 5.8. *Let \mathcal{M} be a winning strategy in $\widetilde{\text{MC}}(\mathcal{G}, v)$. Let \mathcal{S} be a winning strategy in \mathcal{G} so that \mathcal{M} and \mathcal{S} are corresponding strategies. Then, $\pi[\mathcal{M}] = \pi[\mathcal{S}]$.*

Proof. As we removed all other literals from $\text{MC}(\mathcal{G}, v)$, the only literals occurring in \mathcal{M} are edge literals of the form Evw , so $\pi[\mathcal{M}] = \prod_{vw \in E} \pi(Evw)^{\#_{Evw}(\mathcal{M})}$. An edge literal Evw occurs in \mathcal{M} whenever \mathcal{M} transitions from the gadget for v to the gadget for w , and this happens whenever the edge vw occurs in \mathcal{S} . So $\pi[\mathcal{M}] = \prod_{vw \in E} \pi(Evw)^{\#_{vw}(\mathcal{S})} = \pi[\mathcal{S}]$. \square

This closes the proof of the Sum-of-Strategies Theorem: since our modifications did not affect the sum over all winning strategies, Theorem 5.7 implies

$$\begin{aligned} & \pi[\text{win}_0(v)] \\ &= \sum \{ \pi[\mathcal{M}] \mid \mathcal{M} \text{ is a winning strategy in } \text{MC}(\mathcal{G}, v) \} \\ &= \sum \{ \pi[\mathcal{M}] \mid \mathcal{M} \text{ is a winning strategy in } \widetilde{\text{MC}}(\mathcal{G}, v) \} \\ &= \sum \{ \pi[\mathcal{S}] \mid \mathcal{S} \text{ corresponds to a winning strategy in } \widetilde{\text{MC}}(\mathcal{G}, v) \} \\ &= \sum \{ \pi[\mathcal{S}] \mid \mathcal{S} \in \text{WinStrat}_{\mathcal{G}}(v) \}, \end{aligned}$$

and restricting the sum to winning strategies \mathcal{S} that are absorption-dominant from v does not change the overall value. \square

6. CASE STUDY: APPLICATIONS OF SEMIRING SEMANTICS

We now have all of the necessary groundwork to consider applications of semiring semantics for Büchi games. This section discusses what information the Sum-of-Strategies Theorem provides about winning strategies, how semiring semantics helps to find minimal repairs and why it is not well suited for cost computations.

6.1. Strategy Analysis. We begin with the question what information we can derive from the Sum-of-Strategies Theorem. To this end, we fix a Büchi game \mathcal{G} and focus on the $\mathbb{S}^\infty[X]$ -interpretation π_{strat} with $X = \{X_{uv} \mid u, v \in \mathcal{G}\}$. The values $\pi_{\text{strat}}[\mathcal{S}]$ are monomials and we can read off the number of occurrences of each edge in \mathcal{S} from the exponents, i.e., the monomial is a representation of the edge profile $\#_E(\mathcal{S})$. In particular, $\pi_{\text{strat}}[\mathcal{S}_1] \succeq \pi_{\text{strat}}[\mathcal{S}_2]$ if, and only if, $\mathcal{S}_1 \succeq \mathcal{S}_2$. The fact that absorptive polynomials are always finite [DGNT21] is thus another way to see that the number of absorption-dominant strategies is finite.

What can we learn from the polynomial $\pi_{\text{strat}}[\text{win}_0(v)]$? First, $\pi_{\text{strat}}[\text{win}_0(v)] \neq 0$ holds if, and only if, Player 0 has a winning strategy from v . By Theorem 5.4, we can further derive information about all absorption-dominant strategies. More precisely, we learn which edges each absorption-dominant strategy uses and how often they appear in the strategy tree. Knowing the edge profile immediately reveals whether the strategy is positional and what the positional choices are. By counting monomials, we can thus count the positional strategies, as well as the absorption-dominant strategies up to absorption-equivalence.

We can further answer questions such as: can Player 0 still win if we remove edge e ? This is the case if, and only if, the polynomial $\pi_{\text{strat}}[\text{win}_0(v)]$ contains a monomial without the variable X_e (if there is a winning strategy without e , then there is also an absorption-dominant strategy and hence a monomial without X_e). Going further, a more interesting question is: can Player 0 still win if edge e may only be used finitely often in each play? The

answer is not immediately obvious. Consider for example the strategy \mathcal{S} in Fig. 1b. The edge k occurs infinitely often in the strategy tree and we get $\pi_{\text{strat}}[\mathcal{S}] = abcefhg^\infty k^\infty m^\infty$. However, k is clearly played only once in each play consistent with \mathcal{S} , whereas edge m is played infinitely often. We cannot distinguish edges k and m just from $\pi_{\text{strat}}[\mathcal{S}]$, but we can do so if we compute $\pi_{\text{strat}}[\text{win}_0(w)]$ for all positions $w \in V$, by the following criterion (notice that all of these values are computed anyway for the fixed-point iteration).

Proposition 6.1. *Let $\mathcal{S} \in \text{WinStrat}_{\mathcal{G}}(v)$ be absorption-dominant from v , and let $e = uw \in E$ be an edge with $\#_e(\mathcal{S}) = \infty$. Then there is a unique (positional) strategy $\mathcal{S}_w \in \text{WinStrat}_{\mathcal{G}}(w)$ such that $\pi_{\text{strat}}[\mathcal{S}_w] \succeq \pi_{\text{strat}}[\mathcal{S}]$. Moreover, \mathcal{S} admits a play in which e occurs infinitely often if, and only if, e occurs in \mathcal{S}_w .*

Proof. Consider the strategy tree \mathcal{S} and let ρw be an occurrence of w in \mathcal{S} . By assumption, w occurs infinitely often in \mathcal{S} . But then, for all successors $\rho w w'$ of ρw , also w' occurs infinitely often in \mathcal{S} . Indeed, either $w \in V_1$ and every occurrence of w must be followed by an occurrence of w' ; or $w \in V_0$ and \mathcal{S} plays positionally from w by Proposition 3.7, so again every occurrence of w is followed by an occurrence of w' in \mathcal{S} . By induction, it follows that \mathcal{S} plays positionally from w and from all positions occurring below ρw . In particular, the substrategy \mathcal{S}_w that \mathcal{S} plays from ρw (and any other occurrence of w) is positional. As a substrategy, we trivially have $\pi_{\text{strat}}[\mathcal{S}_w] \succeq \pi_{\text{strat}}[\mathcal{S}]$. To see that \mathcal{S}_w is unique with this property, notice that a strategy $\mathcal{S}' \in \text{Strat}_{\mathcal{G}}(w)$ deviating from \mathcal{S}_w must play an edge which does not occur in \mathcal{S}_w and hence also not in \mathcal{S} , so $\pi_{\text{strat}}[\mathcal{S}'] \not\succeq \pi_{\text{strat}}[\mathcal{S}]$.

For the second statement, first assume that \mathcal{S} admits a play $vv_1v_2v_3\dots$ in which the edge $e = uw$ occurs infinitely often. Let $v_i = w$ be the first occurrence of w . Then the remaining play $v_iv_{i+1}v_{i+2}\dots$ is consistent with \mathcal{S}_w and hence e occurs (infinitely often) in \mathcal{S}_w . Conversely, assume that e occurs in \mathcal{S}_w , so there is some $\rho \in V^*$ such that $w\rho uw \in \mathcal{S}_w$. As \mathcal{S}_w is positional, we can repeat ρuw to obtain the infinite play $w(\rho uw)^\omega$ consistent with \mathcal{S}_w . And since \mathcal{S}_w is a substrategy of \mathcal{S} , this induces an infinite play of the form $\rho'w(\rho uw)^\omega$ in \mathcal{S} which indeed uses the edge uw infinitely often. \square

Example 6.2. Consider the strategy \mathcal{S} in Fig. 1b with $\pi_{\text{strat}}[\mathcal{S}] = abcefhg^\infty k^\infty m^\infty$ and the edge k from u to w . Since edge n does not occur in $\pi_{\text{strat}}[\mathcal{S}]$, the only winning strategy \mathcal{S}_w from w we need to consider is the strategy that always stays at w , with $\pi_{\text{strat}}[\mathcal{S}_w] = m^\infty \succeq \pi_{\text{strat}}[\mathcal{S}]$. As k does not occur in \mathcal{S}_w , we conclude that it occurs only finitely often (and hence at most once) in each play consistent with \mathcal{S} .

If, on the other hand, we consider edge m (which also leads to position w), we see that m occurs in \mathcal{S}_w and we can thus infer that \mathcal{S} contains a play visiting m infinitely often. \square

Summarizing the results of this section, we see that semiring semantics in $\mathbb{S}^\infty[X]$ is very informative and allows us to derive important information about the winning strategies.

Corollary 6.3. *From the polynomial $\pi_{\text{strat}}[\text{win}_0(v)]$, we can efficiently (in the size of the polynomial) derive the following information:*

- *whether Player 0 wins from v ,*
- *the edge profiles of all absorption-dominant winning strategies from v ,*
- *the number and precise shape of all positional winning strategies from v ,*
- *whether Player 0 can still win from v if only a subset of the edges is allowed.*

Given the polynomials $\pi_{\text{strat}}[\text{win}_0(v)]$, for all positions v , we can further derive for each (equivalence class of an) absorption-dominant strategy and each edge, how often the edge can occur in a play consistent with the strategy.

6.2. Reverse Analysis. Instead of tracking strategies in a fixed game, we may also ask questions such as: assuming Player 1 wins from v , what are minimal modification to \mathcal{G} such that instead Player 0 wins? The generality of semiring semantics enables us to answer such questions by choosing appropriate semirings and interpretations.

More precisely, let $\mathcal{G} = (V, V_0, V_1, E, F)$ be a Büchi game and $v \in \mathcal{G}$ a position from which Player 1 wins. Let $E^- \subseteq E$ and $E^+ \subseteq V^2 \setminus E$ be sets of edges we are allowed to delete or add, respectively. We call a set of edges in $E^\pm := E^- \cup E^+$ a *repair* if Player 0 wins when these edges are deleted or added. Our goal is to determine all (preferably minimal) repairs. We achieve this by evaluating $\text{win}_0(v)$ in a modified polynomial semiring, similar to the computation of repairs for database queries in [XZAT18], except that here we need absorptive polynomials to deal with fixed points.

Dual-Indeterminates. To track negative information, such as the absence of an edge, we follow the approach in [GT17a, XZAT18, DGNT21] and extend our semiring by dual-indeterminates $\overline{X} = \{\overline{x} \mid x \in X\}$. The idea is to label a literal and its negation by corresponding indeterminates x and \overline{x} . We must then avoid monomials such as $x\overline{x}$, as they represent contradictory information. To this end, we consider the quotient of $\mathbb{S}^\infty[X \cup \overline{X}]$ with respect to the congruence generated by $x \cdot \overline{x} = 0$ for $x \in X$ and refer to the resulting quotient semiring as *dual-indeterminate absorptive polynomials* $\mathbb{S}^\infty[X, \overline{X}]$. This semiring inherits most of the properties of $\mathbb{S}^\infty[X]$. Most importantly, any assignment $h: X \cup \overline{X} \rightarrow \mathbb{K}$ into an absorptive, fully-continuous semiring \mathbb{K} that respects dual-indeterminates, i.e., $h(x) \cdot h(\overline{x}) = 0$, lifts to a fully-continuous homomorphism analogous to Theorem 4.6.

We then replace π_{strat} by an $\mathbb{S}^\infty[X, \overline{X}]$ -interpretation π_{repair}^\pm with $X = \{X_e \mid e \in E^\pm\}$: if $vw \in E^\pm$, we set $\pi_{\text{repair}}^\pm(Evw) = X_{vw}$ and $\pi_{\text{repair}}^\pm(\neg Evw) = \overline{X_{vw}}$, all other literals are mapped to 0 or 1 according to \mathcal{G} (cf. Fig. 3). Notice that π_{repair}^\pm is neither model-defining nor edge-tracking, but still satisfies $\pi_{\text{repair}}^\pm(\alpha) \cdot \pi_{\text{repair}}^\pm(\neg\alpha) = 0$ for all literals α .

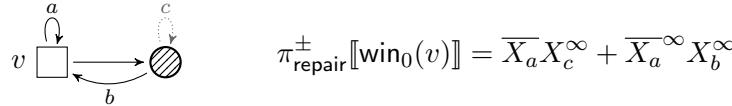
Back and Forth between Monomials and Models. Let $X^\pm = \{X_e \mid e \in E^+\} \cup \{\overline{X}_e \mid e \in E^-\}$. Given $Y \subseteq X^\pm$, we further write $E(Y) = \{e \mid X_e \in Y \text{ or } \overline{X}_e \in Y\}$ for the set of edges mentioned in Y . We denote the set of all (dual-)indeterminates occurring in a monomial m by $\text{var}(m) = \{x \in X \cup \overline{X} \mid m(x) > 0\}$. By examining what combinations of indeterminates from X^\pm occur in the monomials of $\pi_{\text{repair}}^\pm[\text{win}_0(v)]$, we can read off all minimal repairs as follows.

Proposition 6.4. *In the above setting, the following holds:*

- (1) *Let $m \in \pi_{\text{repair}}^\pm[\text{win}_0(v)]$ be a monomial. Then the set $E(\text{var}(m) \cap X^\pm)$ is a repair.*
- (2) *Let $R \subseteq E^\pm$ be a repair. Then there is a monomial $m \in \pi_{\text{repair}}^\pm[\text{win}_0(v)]$ such that $E(\text{var}(m) \cap X^\pm) \subseteq R$. If R is minimal, then $E(\text{var}(m) \cap X^\pm) = R$.*

Before proving Proposition 6.4, we illustrate the computation of minimal repairs in a small example.

Example 6.5. In the following game, Player 1 wins from v . We are interested in the minimal repairs with $E^+ = \{c\}$ and $E^- = \{a, b\}$.



Evaluating $\text{win}_0(v)$ in the $\mathbb{S}^\infty[X, \overline{X}]$ -interpretation π_{repair}^\pm described above results in two monomials. The first yields the repair $\{a, c\}$, the second yields the minimal repair $\{a\}$ (notice that $X_b \notin X^\pm$, as edge b is already present). The reason why we get two monomials is that we track also positive usage of edge b by X_b , but are only interested in the negative indeterminate $\overline{X_b}$ for the repairs. \square

Proof of Proposition 6.4. We prove both statements by considering homomorphisms into the Boolean semiring \mathbb{B} . For the first statement, let $m \in \pi_{\text{repair}}^\pm[\text{win}_0(v)]$ be a monomial and let $h: X \cup \overline{X} \rightarrow \mathbb{B}$ be the unique function that respects dual-indeterminates and satisfies

- $h(x) = \mathbf{1}$, for all $x \in \text{var}(m)$
- $h(X_e) = \mathbf{0}$, if $X_e, \overline{X_e} \notin \text{var}(m)$ and $e \in E^+$ (do not add e without reason),
- $h(X_e) = \mathbf{1}$, if $X_e, \overline{X_e} \notin \text{var}(m)$ and $e \in E^-$ (do not remove e without reason).

Then, h lifts to a fully-continuous semiring homomorphism $h: \mathbb{S}^\infty[X, \overline{X}] \rightarrow \mathbb{B}$ with $h(m) = \mathbf{1}$. Moreover, $h \circ \pi_{\text{repair}}^\pm$ is a Boolean interpretation which corresponds to a Boolean model \mathcal{G}' . Since semiring semantics are preserved by fully-continuous homomorphisms, we have $h \circ \pi_{\text{repair}}^\pm[\text{win}_0(v)] = h(\pi_{\text{repair}}^\pm[\text{win}_0(v)]) \geq h(m) = \mathbf{1}$ and hence $\mathcal{G}' \models \text{win}_0(v)$. By the choice of h , the model \mathcal{G}' is equal to \mathcal{G} except that we add all edges $e \in X^+$ with $X_e \in \text{var}(m)$, and remove all $e \in X^-$ with $\overline{X_e} \in \text{var}(m)$. Hence \mathcal{G}' results from \mathcal{G} by adding or deleting the edges $E(\text{var}(m) \cap X^\pm)$, and since $\mathcal{G}' \models \text{win}_0(v)$, this set is a repair as claimed.

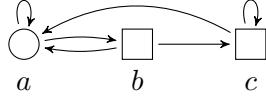
For the second statement, let $R \subseteq E^\pm$ be a repair and consider the repaired game $\mathcal{G}' \models \text{win}_0(v)$. As \mathcal{G}' differs from \mathcal{G} only by edges in E^\pm , there is a unique assignment $h: X \cup \overline{X} \rightarrow \mathbb{B}$ such that $h \circ \pi_{\text{repair}}^\pm$ corresponds to \mathcal{G}' . Again, h lifts to a fully-continuous homomorphisms and we thus get $\mathbf{1} = h \circ \pi_{\text{repair}}^\pm[\text{win}_0(v)] = h(\pi_{\text{repair}}^\pm[\text{win}_0(v)])$. So there must be a monomial $m \in \pi_{\text{repair}}^\pm[\text{win}_0(v)]$ with $h(m) = \mathbf{1}$. Consider the set $\text{var}(m) \cap X^\pm$. If $X_e \in \text{var}(m) \cap X^\pm$, then $h(X_e) = \mathbf{1}$ and hence $e \in R$ by construction of h . Further, $\overline{X_e} \in \text{var}(m) \cap X^\pm$ implies $h(\overline{X_e}) = \mathbf{1}$ and thus again $e \in R$ by construction of h . This proves $E(\text{var}(m) \cap X^\pm) \subseteq R$. If R is minimal, we have equality: otherwise $E(\text{var}(m) \cap X^\pm)$ would be a smaller repair by the first statement, contradicting minimality. \square

We remark that these results ignore the exponents of the monomials, so we can drop exponents from $\mathbb{S}^\infty[X, \overline{X}]$ (in other words, we use exponents from \mathbb{B} instead of $\mathbb{N} \cup \{\infty\}$) and work in the resulting, simpler semiring $\text{PosBool}[X, \overline{X}]$ (the dual-indeterminate quotient of the semiring $\text{PosBool}[X]$, see e.g. [GT17b]).

6.3. Target synthesis. The reverse analysis approach and the proof technique based on homomorphisms are not limited to questions about edges, but are general concepts of semiring provenance analysis. As an example of a different application, we consider the synthesis of the target set F . More precisely, we consider a game arena with positions $V = V_0 \dot{\cup} V_1$ and edges E and want to compute all minimal choices for the set F so that Player 0 wins the resulting Büchi game from some fixed starting position $u \in V$.

Similar to the computation minimal repairs, we can solve this task with an interpretation π_{target} over the dual-indeterminate semiring $\text{PosBool}[X, \overline{X}]$ which interprets most literals, including edge literals, by Boolean values, but tracks the target set F using corresponding pairs of dual-indeterminates $\pi_{\text{target}}(Fv) = X_v$ and $\pi_{\text{target}}(\neg Fv) = \overline{X_v}$ for each position $v \in V$ (cf. Fig. 3). We can then derive all possible minimal choices for F from the polynomial $\pi_{\text{target}}[\text{win}_0(u)]$, as illustrated in the following example.

Example 6.6. Consider the following game arena. What are the minimal choices of the target set F so that Player 0 wins from position a , or b ?



$$\begin{aligned}\pi_{\text{target}}[\text{win}_0(a)] &= X_a + \overline{X_a}X_bX_c \\ \pi_{\text{target}}[\text{win}_0(b)] &= X_aX_bX_c + X_a\overline{X_b}X_c + \overline{X_a}X_bX_c\end{aligned}$$

Using the $\text{PosBool}[X, \overline{X}]$ -interpretation π_{target} , we can derive that to win from a , the target set must contain at least $\{a\}$ or $\{b, c\}$. To win from b , it must contain at least $\{a, b, c\}$, $\{a, c\}$ or $\{b, c\}$ (notice that $\{a, b, c\}$ is a valid choice for F , but not minimal due to the presence of negative indeterminates). This covers all minimal possible choices of the target set. \square

In general, the positive indeterminates X_v in each monomial of $\pi_{\text{target}}[\text{win}_0(u)]$ induce one possible choice of the target set F , and conversely every minimal choice of F occurs as a monomial; this follows by the same arguments as in Proposition 6.4.

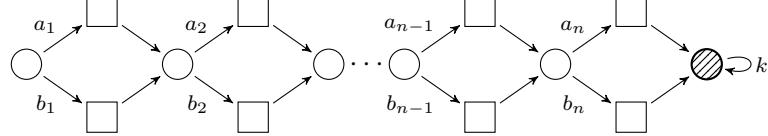
We remark that we can do slightly better, as we do not need to track negative information for the synthesis problem and hence do not need the negative indeterminates $\overline{X_v}$. In the above example, we would then obtain the polynomials $X_a + X_bX_c$ and $X_aX_c + X_bX_c$ for positions a and b , respectively, which correspond exactly to the minimal choices for F . This can be achieved by setting $\pi_{\text{target}}(\neg Fv) = 1$ for all $v \in V$ and observing that this has the same effect as omitting the subformula $\neg Fv$ from win_0 . Since the resulting formula is equivalent in Boolean semantics, the reasoning in the proof of Proposition 6.4 can be adapted to this setting. However, here we presented the general dual-indeterminate approach which does not depend on the actual formula we consider and hence also works in many other scenarios of provenance analysis, beyond the analysis of Büchi games.

6.4. Complexity. The previous applications show that once we have computed the polynomial $\pi_{\text{strat}}[\text{win}_0(v)]$ or $\pi_{\text{repair}}^{\pm}[\text{win}_0(v)]$, it is easy to derive information about strategies or minimal repairs – but how efficient is the computation of the polynomial in the first place? Recall that for a fixed formula, the LFP model-checking problem for a classical structure can be solved in polynomial time in the size of the structure (here the number of positions). The analogous problem in semiring semantics, to compute the value $\pi_{\text{strat}}[\text{win}_0(v)]$, is more involved since depending on the semiring, fixed-point iterations can be infinite, and contrary to the Boolean case, it is not trivial that fixed points can be computed efficiently. However, it has been shown in [Naa21] that in absorptive, fully-continuous semirings, a least or greatest fixed point of a polynomial system, such as the ones induced by LFP-formulae, can be computed using only a polynomial number of semiring operations (including the infinitary power operation). By evaluating nested fixed points recursively, we can thus compute the semiring value of a fixed⁵ formula in a number of semiring operations that is polynomial in the size of the structure (here in the number of positions).

Whether this computation is efficient depends on the complexity of the semiring operations and hence on the semiring. Indeed, although the *number* of semiring operations is polynomial, the resulting polynomial $\pi_{\text{strat}}[\text{win}_0(v)]$ can nevertheless have an exponential number of monomials. This is, in general, unavoidable as both the number of (positional) winning strategies as well as the number of minimal repairs can be exponential in the size

⁵In general, this applies to all LFP formulae in which the alternation depth, the arity of fixed-point relations and the number of free variables in any subformula is bounded by a constant.

of the game. For instance, games of the following form have an exponential number of positional (and hence absorption-dominant) winning strategies:



Algorithmically, the computation of the polynomials $\pi_{\text{strat}}[\text{win}_0(v)]$ or $\pi_{\text{repair}}^\pm[\text{win}_0(v)]$ is thus infeasible in many cases. It should be noted that the questions addressed in Corollary 6.3 can be solved by direct methods, and just for finding *some* winning strategy or repair, these will in many cases be more efficient than computing the polynomials. Thus, the main benefit of semiring semantics in $\mathbb{S}^\infty[X]$ does not lie in a more efficient algorithmic method to compute some specific winning strategy or repair, but rather in providing a general and compact description of many important strategies at once, from which we can directly derive the answers to a number of different questions concerning the strategy analysis of a Büchi game. This is attractive in cases where the relevant polynomials are reasonably small.

One scenario where we can achieve this is, when only some of the edges are tracked: instead of the interpretation π_{strat} that assigns indeterminates to all edges, we map edges we are not interested in to 1. The resulting $\mathbb{S}^\infty[X]$ -interpretation remains edge-tracking, so the Sum-of-Strategies Theorem still applies and we can see which sets of the tracked edges are at least required for a winning strategy. In particular, $\text{win}_0(v)$ evaluates to 1 if there is a winning strategy that avoids all tracked edges. If this information is sufficient, for instance if we know that a certain part of the game must be visited and only care about edges within this part, we can make provenance analysis more efficient by only tracking small parts of a potentially large game.

6.5. Limitations. We have seen that semiring semantics in polynomial semirings such as $\mathbb{S}^\infty[X]$ reveals useful information about strategies. However, there are also limitations to this framework when it comes to certain typical applications of provenance analysis such as cost computation. Given cost annotations of the edges, say in the tropical semiring $\mathbb{T} = (\mathbb{R}_+^\infty, \min, +, \infty, 0)$ (which is absorptive and fully continuous), do semiring semantics provide a sensible cost measure for the evaluation of $\text{win}_0(v)$? Or, more intuitively, for the cost that Player 0 has to pay to win? For first-order logic and acyclic games, this is indeed the case [GT17a, GT20], but here fixed-point computations and infinite plays complicate the situation. To see this, let π be an edge-tracking \mathbb{T} -interpretation that assigns to each edge vw in \mathcal{G} a cost $\pi(Evw) \neq \infty$ (cost 0 is allowed). By the Sum-of-Strategies Theorem, we can view $\pi[\text{win}_0(v)]$ as the minimum over the cost of each strategy, and in the tropical semiring, this cost can be expressed as follows (computed over real numbers):

$$\pi[\mathcal{S}] = \sum_{vw \in E} \#_{vw}(\mathcal{S}) \cdot \pi(Evw).$$

That is, Player 0 has to pay for each *occurrence* of an edge in the strategy tree. While this is certainly a possible cost measure for a strategy, it is debatable whether it is an intuitive one. Many edges will occur infinitely often in a strategy, and then $\infty \cdot \pi(Evw)$ is either 0 or ∞ , and the latter leads to the overall value $\pi[\mathcal{S}] = \infty$. Even if an edge occurs only once per play, but infinitely often in \mathcal{S} , Player 0 has to pay the cost $\infty \cdot \pi(Evw)$. Instead, it might be more intuitive to define the cost of a strategy so that

- (1) Player 0 only has to pay once for an edge, no matter how often it occurs (think of an “unlocking fee”),
- (2) Player 0 only has to pay for the maximal cost of any play consistent with \mathcal{S} , but not for all plays simultaneously.

We claim that neither of these options is possible without adapting our notion of semiring semantics. Notice that we can actually solve (1) by first computing the polynomial $\pi_{\text{strat}}[\text{win}_0]$ in $\mathbb{S}^\infty[X]$, then dropping exponents (or directly working in $\text{PosBool}[X]$), and then instantiating each variable X_{vw} by its cost to obtain a value in \mathbb{T} . But this is only an indirect solution and computing $\pi_{\text{strat}}[\text{win}_0]$ may incur an exponential blowup even though we only compute a single cost value.

A direct computation, in the sense that $\pi[\mathcal{S}]$ yields the desired cost value for some suitable π , is not possible. To see this for (1), note that we multiply the cost for each occurrence of an edge (in \mathbb{T} , semiring multiplication is defined as summation on real numbers, but we stick to the general vocabulary). Say we have two edges vw and $v'w'$ with the same cost $\pi(Evw) = \pi(Ev'w') = c$. To pay only once for vw , we would need $\pi(Evw) \cdot \pi(Evw) = c$, but at the same time $\pi(Evw) \cdot \pi(Ev'w') = c^2$ for different edges, a contradiction. The issue is that the information on whether two edges are equal is abstracted away by π .

A different argument explains why (2) is not possible. At a position in V_0 , we want to *minimize* the cost over all possible choices (corresponding to the existential quantification in win_0). For consecutive edges, we have to *add* costs up, requiring a second operation. But for positions in V_1 , to fulfil (2) we must *maximize* the cost over all possible choices, thus requiring a third operation. Hence this cost measure is not expressible in semirings with only two operations and we would need different algebraic structures.

7. CONCLUSION

Based on a recent line of research on semiring provenance analysis that lead from database theory to semiring semantics for LFP, we reported here on a case study that puts semiring semantics to use for a strategy analysis in Büchi games. The choice of Büchi games has been motivated on one side by their relevance for applications in the synthesis and verification of reactive systems, on the other side because they provide one of the simplest non-trivial cases of infinite games for which the definability of winning positions requires an alternation between least and greatest fixed points – and can thus not be treated by simpler classes of semirings such as the ω -continuous ones used for Datalog and reachability games.

The aim of the case study was to illustrate how semiring semantics can be applied to more complex games, featuring infinite plays and complicated winning conditions, and what kind of insights it provides (or fails to provide) about the winning strategies in the game. This is captured in the central Sum-of-Strategies Theorem and its applications. This non-trivial result can be seen as a simpler version of the general sum-of-strategies characterization in terms of model-checking games in [DGNT21] and it essentially identifies the value of the statement that Player 0 wins with the sum of the valuations of all (absorption-dominant) winning strategies. While this applies to the class of all absorptive, fully-continuous semirings, the most important semirings for our analysis are generalized absorptive polynomials $\mathbb{S}^\infty[X]$. Due to their universal property, these provide the most general information, allowing us to read off the edge profiles of all absorption-dominant strategies.

With this information, we can count positional strategies, we can determine whether a particular move is needed (once or even infinitely often) for winning the game, and we

can compute minimal “repairs” for a game. The method of semiring valuations is rather flexible; we can use different semirings than $\mathbb{S}^\infty[X]$, and we can tailor the set of moves that we track to make the resulting polynomial smaller and its computation more efficient. Of course, not all relevant questions about strategies can be answered directly by semiring valuations, and as an example for a limitation of this method, we have shown that minimal cost computations provide serious obstacles to a semiring treatment.

The Sum-of-Strategies Theorem motivates the notion of absorption-dominant strategies which is of interest in its own right as it captures strategies that are minimal with respect to the multiplicities of the edges they use. To understand these strategies, we discussed how they relate to other classes of simple strategies, namely to positional and persistent strategies, and we have shown that these form a strict hierarchy.

Finally, we remark that although Büchi games have been chosen as the topic of our case study, the method of semiring valuations in absorptive semirings is not confined to this case. In principle, it can be applied to different formulae and generalizes in particular to other games such as parity games, as long as the winning positions are definable in fixed-point logic. The win-formula for parity games is more complicated, and is parametrised by the number of priorities, and so the Sum-of-Strategies Theorem requires different technical details, but can be established along the same lines.

REFERENCES

- [AG11] K. Apt and E. Grädel, editors. *Lectures in Game Theory for Computer Scientists*. Cambridge University Press, 2011. doi:10.1017/CBO9780511973468.
- [BOPP20] C. Bourgault, A. Ozaki, R. Peñaloza, and L. Predoiu. Provenance for the description logic ELHr. In *Proceedings of IJCAI 2020*, pages 1862–1869, 2020. doi:10.24963/ijcai.2020/258.
- [BW18] J. Bradfield and I. Walukiewicz. The mu-calculus and model checking. In E. Clarke, T. Henzinger, H. Veith, and R. Bloem, editors, *Handbook of Model Checking*, pages 871–919. Springer, 2018. doi:10.1007/978-3-319-10575-8_26.
- [CDHL16] K. Chatterjee, W. Dvorák, M. Henzinger, and V. Loitzenbauer. Conditionally optimal algorithms for generalized Büchi games. arXiv:1607.05850 [cs.DS], 2016. URL: <http://arxiv.org/abs/1607.05850>.
- [CGLP15] F. Canavoi, E Grädel, S. Lessenich, and W. Pakusa. Defining winning strategies in fixed-point logic. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 366–377, 2015. doi:10.1109/LICS.2015.42.
- [CH14] K. Chatterjee and M. Henzinger. Efficient and dynamic algorithms for alternating Büchi games and maximal end-component decomposition. *J. ACM*, 61(3):15:1–15:40, 2014. doi:10.1145/2597631.
- [CHP08] K. Chatterjee, T. Henzinger, and N. Piterman. Algorithms for Büchi games. arXiv:0805.2620 [cs.GT], 2008. URL: <http://arxiv.org/abs/0805.2620>.
- [DG19] K. Dannert and E. Grädel. Provenance analysis: A perspective for description logics? In C. Lutz et al., editor, *Description Logic, Theory Combination, and All That*, volume 11560 of *Lecture Notes in Computer Science*, pages 266–285. Springer, 2019. doi:10.1007/978-3-030-22102-7_12.
- [DG20] K. Dannert and E. Grädel. Semiring provenance for guarded logics. In *Hajnal Andréka and István Németi on Unity of Science: From Computing to Relativity Theory through Algebraic Logic*, Outstanding Contributions to Logic, pages 53–79. Springer, 2020. doi:10.1007/978-3-030-64187-0_3.
- [DGNT19] K. Dannert, E. Grädel, M. Naaf, and V. Tannen. Generalized absorptive polynomials and provenance semantics for fixed-point logic. arXiv:1910.07910 [cs.LO], 2019. Full version of [DGNT21]. URL: <https://arxiv.org/abs/1910.07910>.
- [DGNT21] K. Dannert, E. Grädel, M. Naaf, and V. Tannen. Semiring Provenance for Fixed-Point Logic. In C. Baier and J. Goubault-Larrecq, editors, *29th EACSL Annual Conference on Computer Science Logic (CSL 2021)*, volume 183 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 17:1–17:22, Dagstuhl, 2021. doi:10.4230/LIPIcs.CSL.2021.17.

- [DMRT14] D. Deutch, T. Milo, S. Roy, and V. Tannen. Circuits for datalog provenance. In *Proc. 17th International Conference on Database Theory ICDT*, pages 201–212. OpenProceedings.org, 2014. doi:10.5441/002/icdt.2014.22.
- [Dup03] J. Duparc. Positive games and persistent strategies. In *International Workshop on Computer Science Logic*, pages 183–196. Springer, 2003. doi:10.1007/978-3-540-45220-1_17.
- [GKL⁺07] E. Grädel, P. G. Kolaitis, L. Libkin, M. Marx, J. Spencer, M. Y. Vardi, Y. Venema, and S. Weinstein. *Finite Model Theory and Its Applications*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2007. doi:10.1007/3-540-68804-8.
- [GKT07] T. Green, G. Karvounarakis, and V. Tannen. Provenance semirings. In *Principles of Database Systems PODS*, pages 31–40, 2007. doi:10.1145/1265530.1265535.
- [GLN21] E. Grädel, N. Lücking, and M. Naaf. Semiring provenance for Büchi games: Strategy analysis with absorptive polynomials. In Pierre Ganty and Davide Bresolin, editors, *Proceedings 12th International Symposium on Games, Automata, Logics, and Formal Verification (GandALF 2021)*, volume 346 of *EPTCS*, pages 67–82, 2021. doi:10.4204/EPTCS.346.5.
- [GP10] F. Geerts and A. Poggi. On database query languages for K-relations. *Journal of Applied Logic*, 8(2):173–185, 2010. doi:10.1016/j.jal.2009.09.001.
- [GT17a] E. Grädel and V. Tannen. Semiring provenance for first-order model checking. arXiv:1712.01980 [cs.LO], 2017. URL: <https://arxiv.org/abs/1712.01980>.
- [GT17b] T. Green and V. Tannen. The semiring framework for database provenance. In *Proceedings of PODS*, pages 93–99. ACM, 2017. doi:10.1145/3034786.3056125.
- [GT20] E. Grädel and V. Tannen. Provenance analysis for logic and games. *Moscow Journal of Combinatorics and Number Theory*, 9(3):203–228, 2020. Preprint available at <https://arxiv.org/abs/1907.08470>. doi:10.2140/moscow.2020.9.203.
- [GTW02] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002. doi:10.1007/3-540-36387-4.
- [Moh02] M. Mohri. Semiring frameworks and algorithms for shortest-distance problems. *Journal of Automata, Languages and Combinatorics*, 7(3):321–350, 2002. doi:10.25596/jalc-2002-321.
- [MT02] J. Marcinkowski and T. Trudlerung. Optimal complexity bounds for positive LTL games. In *International Workshop on Computer Science Logic*, pages 262–275. Springer, 2002. doi:10.1007/3-540-45793-3_18.
- [Naa21] M. Naaf. Computing least and greatest fixed points in absorptive semirings. In U. Fahrenberg, M. Gehrke, L. Santocanale, and M. Winter, editors, *Relational and Algebraic Methods in Computer Science (RAMiCS 2021)*, volume 13027 of *Lecture Notes in Computer Science*, pages 344–361. Springer, 2021. doi:10.1007/978-3-030-88701-8_21.
- [OP18] A. Ozaki and R. Peñaloza. Provenance in ontology-based data access. In *Description Logics 2018*, volume 2211 of *CEUR Workshop Proceedings*, 2018. URL: <http://ceur-ws.org/Vol-2211/paper-28.pdf>.
- [RMZ⁺20] M. Raghothaman, J. Mendelson, D. Zhao, M. Naik, and B. Scholz. Provenance-guided synthesis of datalog programs. *Proc. ACM Program. Lang.*, 4:62:1–62:27, 2020. doi:10.1145/3371130.
- [Sen17] P. Senellart. Provenance and probabilities in relational databases: From theory to practice. *SIGMOD Record*, 46(4):5–15, 2017. doi:10.1145/3186549.3186551.
- [Wal02] I. Walukiewicz. Monadic second-order logic on tree-like structures. *Theoretical Computer Science*, 275:311–346, 2002. doi:10.1016/S0304-3975(01)00185-2.
- [XZAT18] J. Xu, W. Zhang, A. Alawini, and V. Tannen. Provenance analysis for missing answers and integrity repairs. *IEEE Data Eng. Bull.*, 41(1):39–50, 2018. URL: <http://sites.computer.org/debull/A18mar/p39.pdf>.