

Semiring Provenance for Büchi Games: Strategy Analysis with Absorptive Polynomials

Erich Grädel

RWTH Aachen University
Aachen, Germany

graedel@logic.rwth-aachen.de

Niels Lücking

RWTH Aachen University
Aachen, Germany

niels.luecking@rwth-aachen.de

Matthias Naaf

RWTH Aachen University
Aachen, Germany

naaf@logic.rwth-aachen.de

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.

Due to space reasons, several proofs have been omitted and can be found in the full version [14].

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 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. [17] 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. [4, 5, 6]). 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 Sect. 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 [16] 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 [20], 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 [9]. 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 [1, 7, 8, 9, 15], following the successful development of semiring provenance in database theory and related fields (see e.g. [10, 12, 18, 19, 23, 24, 25]). 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. 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 on semirings $\mathbb{S}^\infty[X]$ of *generalized absorptive polynomials*, which we explain in Sect. 4.1, and which 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 $h: \mathbb{S}^\infty[X] \rightarrow \mathbb{K}$, see Theorem 17. 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. [2].

The theoretical underpinning of our method is 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 . The proof is somewhat involved and requires non-trivial machinery; details are given in the full version of this paper [14]. 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 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 $\mathbb{S}^\infty[X, \bar{X}]$ by a construction that has been used before in [15, 27, 9] 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.

2 Büchi Games and Strategies

A Büchi game is given by a tuple $G = (V, V_0, V_1, E, F)$ where V is a set of positions (here assumed to be finite), with a disjoint decomposition $V = V_0 \cup V_1$ 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 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 $G = (V, V_0, V_1, E, F)$ 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., [16]). For simplicity, we only consider strategies of Player 0, so unless mentioned otherwise, “strategy” always refers to a strategy for Player 0.

Definition 1. Given a Büchi game $G = (V, V_0, V_1, E, F)$, the *tree unraveling* from v_0 is the tree $T(G, v_0)$ whose nodes are all finite paths ρ from v_0 in G and whose edges are $\rho \rightarrow \rho v$ for $v \in V$. We often write a node of $T(G, v_0)$ as ρv to indicate a finite path ending in $v \in V$. The length of ρ is denoted by $|\rho|$ and we write $\rho \sqsubseteq \rho'$ 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 the strategy. Moreover, we only consider finite Büchi games and hence the tree unraveling and all strategies are finitely branching.

Definition 2. A *strategy* S (of Player 0) from v_0 in G is a subtree of $T(G, v_0)$ 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 S are winning.

We commonly write $\rho \in S$ instead of $\rho \in W$, and we often refer to paths of the form $\rho v \in S$ as *occurrences of v* in 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 Figure 1b for an example. For $v \in V_0$, we further write $S(\rho v) = w$ if $\rho v w$ is the (unique) successor of ρv in S . If S is positional, we may also write $S(v)$ to denote the unique successor of v chosen by S . We write $\text{Strat}_G(v)$ and $\text{WinStrat}_G(v)$ to denote the set of all (winning) strategies of Player 0 from position $v \in G$, and we drop G if the game is clear from the context.



(a) Rectangular nodes belong to Player 1, round nodes to Player 0, dashed nodes are in F .



(b) Depiction of an infinite strategy tree of a winning strategy for Player 0 from position v .

Figure 1: Running example of a Büchi game and a winning strategy.

Example 3. An example of a Büchi game is depicted in Figure 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 into account. \lrcorner


3 Strategies with Minimal Effort

Perfection is achieved, not when there is nothing more to add, but when there is nothing left to take away
— ANTOINE DE SAINT-EXUPÉRY

As a measure for the complexity or effort of a strategy, we consider the set of edges a strategy 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 4. Given an edge $e = vw \in E$ in a Büchi game G and a strategy S in G , we denote by $\#_e(S) = |\{\rho v \in S \mid \rho v \rightarrow \rho vw \text{ is an edge in } S\}| \in \mathbb{N} \cup \{\infty\}$ the number of times (possibly infinite) the edge e occurs in S . With each strategy S we associate its *edge profile*, the vector $\#_E(S) = (\#_e(S))_{e \in E}$.

Example 5. 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 consider the strategy with $n = 0$ as the simplest strategy (that wins with the least effort). \lrcorner

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 6. Let S_1, S_2 be two strategies in a Büchi game $G = (V, V_0, V_1, E, F)$. We say that S_1 *absorbs* S_2 , denoted $S_1 \succeq S_2$, if $\#_e(S_1) \leq \#_e(S_2)$ for all edges $e \in E$. If additionally $\#_e(S_1) < \#_e(S_2)$ for some $e \in E$, we say that S_1 *strictly absorbs* S_2 , denoted $S_1 \succ S_2$. They are *absorption-equivalent*, denoted $S_1 \equiv S_2$, if both $S_1 \succeq S_2$ and $S_2 \succeq S_1$. A strategy $S \in \text{Strat}(v)$ is *absorption-dominant from position* v , if there is no strategy $S' \in \text{Strat}(v)$ with $S' \succ S$. It is further *strictly absorption-dominant*, if there is no other strategy $S' \in \text{Strat}(v)$ with $S' \succeq S$, so no other strategy is absorption-equivalent to S .

Notice that absorption is simply the inverse pointwise order on the edge profiles. In particular, $S_1 \equiv S_2$ if, and only if, $\#_E(S_1) = \#_E(S_2)$. 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. [16] for a similar example).

Example 7. Consider the strategy S as depicted in Figure 1b. It is not positional, as the choice for position v' is not unique (both e and f occur in 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 S , and one can verify that 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 S , so that e is used after c , and f after b . \lrcorner




Figure 2: Venn diagram depicting the relation of various classes of winning strategies.

Strategies such as the one in Figure 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 [20] in the context of LTL on game graphs and has been further studied in [11]. Persistent strategies have also been called *weakly positional* in [16].

We say that a strategy *plays positionally* from a position $v \in V_0$ if the strategy makes a unique choice at position v . 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 Figure 2. We first observe that if a strategy S does not play positionally from v , we can always obtain a strategy S' with $S' \succeq S$ by swapping the choices at v . This leads to the first result (see [14] for details):

Proposition 8. *Strictly absorption-dominant strategies coincide with positional strategies.*

Towards persistent strategies, we observe, by applying a simple combinatorial fact known as Dickson's Lemma to edge profiles, that there are only finitely many absorption-dominant strategies from v up to absorption-equivalence. This leads to the following result:


Proposition 9. *Let $S \in \text{WinStrat}_G(v)$ be absorption-dominant from v , and let $w \in V_0$ be a position. If w occurs infinitely often in S , then S plays positionally from w .*

Proof sketch. Consider the infinitely many substrategies at occurrences of w in S . By Dickson's Lemma, there is one such substrategy S_w such that infinitely many of the substrategies are absorption-equivalent to S_w . In particular, S_w only uses edges that occur infinitely often in S . By positional determinacy, there is thus a positional winning strategy S_{pos} from w with the same property. If S would not be positional, then modifying S to always play S_{pos} from w would result in a strategy $S' \succ S$, a contradiction. \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 by Proposition 9; positions that occur finitely often cannot be visited twice in one play (this would be a redundant repetition).

Corollary 10. *Every absorption-dominant winning strategy in G is persistent.*

Example 11. For strictness, consider the following game (a modified part of Figure 1a):



Due to the self-loop b , only the positional strategies S_1 (always take a) and S_2 (always take b) are absorption-dominant from v . The strategy 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 S_2 . \lrcorner

As a consequence of Corollary 10, all moves after the first repeated position are determined by persistence and we can thus represent absorption-dominant strategies in a compact way.

Corollary 12. *Let G be a game with $n = |V|$ positions. Every winning strategy $S \in \text{WinStrat}_G(v)$ 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 [9].

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, following the definitions in [9].

Definition 13. 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 14. 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 \bigsqcup and infimum \bigsqcap) and for all non-empty chains $C \subseteq \mathbb{K}$, elements $a \in \mathbb{K}$ and $\circ \in \{+, \cdot\}$,

$$\bigsqcup(a \circ C) = a \circ \bigsqcup C, \quad \text{and} \quad \bigsqcap(a \circ C) = a \circ \bigsqcap C.$$

A semiring homomorphism $h: \mathbb{K}_1 \rightarrow \mathbb{K}_2$ on fully-continuous semirings is *fully continuous* if $h(\bigsqcup C) = \bigsqcup h(C)$ and $h(\bigsqcap C) = \bigsqcap h(C)$ 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 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 [9].

Definition 15. 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* [21], 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 $a^\infty = \bigsqcap_{n \in \mathbb{N}} a^n$ with natural properties such as $a \cdot a^\infty = a^\infty$ and $(a + b)^\infty = a^\infty + b^\infty$.

Example 16. Examples of semirings used in provenance analysis of databases and logics [18, 19, 16] are:

- 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* for logical statements. It is isomorphic to the *tropical semiring* $\mathbb{T} = (\mathbb{R}_+^\infty, \min, +, \infty, 0)$ which is 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. \lrcorner

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 [18, 15], are not absorptive and hence not well-suited for fixed-point logic. Instead, we rely on semirings of (generalized¹) *absorptive polynomials*. We summarize the definition and main properties given in [9]. 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 $\mathbb{N} \cup \{\infty\}$ (where $n < \infty$ for $n \in \mathbb{N}$). Monomial multiplication is defined as usual (with $n + \infty = \infty$).

The key ingredient is absorption among monomials. We say that a monomial m_1 *absorbs* 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^2b and ab^2 are incomparable. In an absorptive polynomial, we omit all monomials that would be absorbed, so absorptive polynomials are \succeq -*antichains* of monomials. Consequently, addition and multiplication are defined as usual, but afterwards we only keep the \succeq -maximal monomials. For example, $(ab^2 + a^2b) \cdot a^\infty = a^\infty b^2 + a^\infty b = a^\infty b$.

We write $\mathbb{S}^\infty[X]$ for the semiring of absorptive polynomials over the finite variable set X . The 0 and 1-elements are the empty polynomial and the single monomial 1 (with all zero exponents). This defines an absorptive, fully-continuous semiring [9]. In fact, $\mathbb{S}^\infty[X]$ is the most general such semiring:

Theorem 17 (Universal Property, [9]). *Every mapping $h: X \rightarrow \mathbb{K}$ into an absorptive, fully-continuous semiring \mathbb{K} uniquely extends to a fully-continuous semiring homomorphism $h: \mathbb{S}^\infty[X] \rightarrow \mathbb{K}$.*

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 $\psi(\mathbf{y}) = [\mathbf{lfp} R\mathbf{x}. \varphi(R, \mathbf{x})](\mathbf{y})$ and $\psi(\mathbf{y}) = [\mathbf{gfp} R\mathbf{x}. \varphi(R, \mathbf{x})](\mathbf{y})$. 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 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 [13].

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

¹The definition we use here generalizes the notion of absorptive polynomials in [10] by allowing ∞ as exponent.

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

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

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

Definition 18. 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 \mathbb{K} is *model-defining*, if for each literal L , exactly one of $\pi(L)$ and $\pi(\neg L)$ is 0.

We lift π to LFP-formulae in negation normal form, resulting in a semiring value $\pi[\psi]$, by interpreting \wedge and \vee as semiring operations \cdot and $+$, and similarly quantifiers as products or sums over the (finite) universe. For fixed-point formulae, we consider the induced operator F_φ analogous to the Boolean case (but acting on mappings $A^k \rightarrow \mathbb{K}$ instead of relations $R \subseteq A^k$) and compute its fixed point in the semiring \mathbb{K} . We refer to [9] for a formal definition; the details are not relevant for this paper, as we will only consider a fixed formula ψ (see Section 5). An important property of the resulting semantics is that they are preserved by fully-continuous semiring homomorphisms, in particular by polynomial evaluation in $\mathbb{S}^\infty[X]$ due to Theorem 17 (but not by polynomial evaluation of $\mathbb{N}[X]$ or formal power series!).

5 Case Study: 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. Figure 1a) and use an $\mathbb{S}^\infty[X]$ -interpretation π_{strat} to track moves (i.e., edge literals Euv) via their indeterminates. We assume the game graph to be fixed and do not wish to track information about the winning set F or the active player at a certain node, hence we simply map all other literals over $\tau = \{E, F, V_0, V_1\}$ (e.g., $Fv, V_0v, \neg Euv$) to 0 or 1, depending on whether they are true or false in the 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 19. Let $G = (V, V_0, V_1, E, F)$ 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 π_{strat} as follows (depending on G):

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

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 by the following formula (see, e.g., [3, 26]):

$$\begin{aligned} \text{win}_0(x) &:= [\mathbf{gfp} Yy. [\mathbf{lfp} Zz. \varphi(Y, Z, z)](y)](x), \\ \varphi(Y, Z, z) &:= \left(Fz \wedge ((V_0z \wedge \exists u(Ezu \wedge Yu)) \vee (V_1z \wedge \forall u(Ezu \rightarrow Yu))) \right) \\ &\quad \vee \left(\neg Fz \wedge ((V_0z \wedge \exists u(Ezu \wedge Zu)) \vee (V_1z \wedge \forall u(Ezu \rightarrow Zu))) \right). \end{aligned}$$

Given a \mathbb{K} -interpretation π for a Büchi game $G = (V, V_0, V_1, E, F)$, semiring semantics of the above formula induce² 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 semiring semantics 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 $v \in V$:

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

For most of this paper, we use the interpretation 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 :

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

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 of the semiring, then apply the above equations (i.e., the induced operator F_φ) 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, we then continue with the supremum/infimum).

Example 20. Recall the simple game from Example 5 ($a \xrightarrow{c} \bigcirc_v \xrightarrow{b} \bigcirc_w \xrightarrow{c} c$).

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 $(Z_v Z_w)^T$.

$$\begin{array}{l} \mathbf{Y}: \begin{pmatrix} 1 \\ 1 \end{pmatrix} \quad \begin{pmatrix} bc \\ c \end{pmatrix} \quad \begin{pmatrix} bc^2 \\ c^2 \end{pmatrix} \quad \rightsquigarrow \mathbf{Y}^* = \begin{pmatrix} bc^\infty \\ c^\infty \end{pmatrix} \\ \mathbf{Z}: \begin{pmatrix} 0 \\ 0 \end{pmatrix} \xrightarrow{F_\varphi} \begin{pmatrix} 0 \\ c \end{pmatrix} \xrightarrow{F_\varphi} \begin{pmatrix} bc \\ c \end{pmatrix} \quad \begin{pmatrix} 0 \\ 0 \end{pmatrix} \xrightarrow{F_\varphi} \begin{pmatrix} 0 \\ c^2 \end{pmatrix} \xrightarrow{F_\varphi} \begin{pmatrix} bc^2 \\ c^2 \end{pmatrix} \quad \dots \end{array}$$

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 5). \lrcorner

²Here we first translate $Ezu \rightarrow Yu$ to the formula $\neg Ezu \vee (Ezu \wedge Yu)$ in negation normal form.

5.3 Connection to Strategies

By mapping edges to semiring values, we can track edges through the fixed-point computation. In Example 20, 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 G . That is, they may assign arbitrary values to positive edge literals E_{vw} , but all other literals are mapped to 0 or 1 in accordance with G . To make the connection to strategies explicit, we first define semiring values for strategies based on their edges.

Definition 21. Let S be a strategy in a Büchi game $G = (V, V_0, V_1, E, F)$. Let \mathbb{K} be an absorptive, fully-continuous semiring and π an edge-tracking \mathbb{K} -interpretation on G . The \mathbb{K} -value of S is the product of the values for all edges appearing in S . Formally,

$$\pi[[S]] := \prod_{vw \in E} \pi(E_{vw})^{\#_{vw}(S)}.$$

The semiring value of win_0 can then be expressed as the sum over the values of all winning strategies. The proof of this result (see [14]) is based on a similar result for LFP model-checking games in [9] which is rather involved due to both infinite strategy trees and infinite fixed-point iterations.

Theorem 22 (Sum of Strategies). *Let G be a Büchi game and v a position in G . Let \mathbb{K} be an absorptive, fully-continuous semiring and π an edge-tracking \mathbb{K} -interpretation. Then,*

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

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 23. For the edge-tracking interpretation π_{strat} induced by the edge labels in Figure 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 Figure 1b (and its switched version, see Example 7). The values of all other strategies are strictly absorbed by one of these monomials. \lrcorner

5.4 Strategy Analysis

For the general case, we fix a Büchi game G and focus on the $\mathbb{S}^\infty[X]$ -interpretation π_{strat} with $X = \{X_{uv} \mid u, v \in G\}$. The values $\pi_{\text{strat}}[[S]]$ are monomials and we can read off the number of occurrences of each edge in S from the exponents, i.e., the monomial is a representation of the edge profile $\#_E(S)$. In particular, $\pi_{\text{strat}}[[S_1]] \succeq \pi_{\text{strat}}[[S_2]]$ if, and only if, $S_1 \succeq S_2$. The fact that absorptive polynomials are always finite [9] 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 22, 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 S in Figure 1b. The edge k occurs infinitely often in the strategy tree and we get $\pi_{\text{strat}}[S] = abcefhg^\infty k^\infty m^\infty$. However, k is clearly played only once in each play consistent with S , whereas edge m is played infinitely often. We cannot distinguish edges k and m just from $\pi_{\text{strat}}[S]$, 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.

Proposition 24. *Let $S \in \text{WinStrat}_G(v)$ be absorption-dominant from v , and let $e = uw \in E$ be an edge with $\#_e(S) = \infty$. Then there is a unique positional strategy $S_w \in \text{WinStrat}_G(w)$ such that $\pi_{\text{strat}}[S_w] \succeq \pi_{\text{strat}}[S]$. Moreover, S admits a play in which e occurs infinitely often if, and only if, e occurs in S_w .*

Example 25. Consider the strategy S in Figure 1b with $\pi_{\text{strat}}[S] = abcefhg^\infty k^\infty m^\infty$ and the edge k from u to w . Since edge n does not occur in $\pi_{\text{strat}}[S]$, the only winning strategy S_w from w we need to consider is the strategy that always stays at w , with $\pi_{\text{strat}}[S_w] = m^\infty \succeq \pi_{\text{strat}}[S]$. As k does not occur in S_w , we conclude that it occurs only finitely often (and hence at most once) in each play consistent with S .

If, on the other hand, we consider edge m (which also leads to position w), we see that m occurs in S_w and we can thus conclude that S contains a play visiting m infinitely often. \lrcorner

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 26. *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.

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

Regarding the complexity, we remark that it is not obvious how to compute $\pi_{\text{strat}}[\text{win}_0(v)]$, as depending on the semiring, the fixed-point iteration may be infinite (even if the game is finite). However, it has been shown in [22] that in absorptive, fully-continuous semirings, fixed points of polynomial systems, such as the ones induced by $\text{win}_0(v)$, can be computed using a number of semiring operations (including the infinitary power operation) that is polynomial in the size of the game.

Notice, however, that for $\pi_{\text{strat}}[\text{win}_0(v)]$, the number of monomials and hence the cost to compute semiring operations can become exponentially large. This is not avoidable, as the number of (positional) winning strategies can be exponential. More efficient algorithms can be obtained by modifying π_{strat} to track only some of the edges (mapping the rest to 1). The resulting interpretation remains edge-tracking, so the Sum-of-Strategies Theorem still applies and we can see which sets of tracked edges (if any) are required for a winning strategy. Even then, some of the questions in Corollary 26 can also be solved by direct methods which may be more efficient. Thus, the main benefit of semiring semantics in $\mathbb{S}^\infty[X]$ does not lie in a more efficient method to compute some specific winning strategy, but rather in providing a general and compact description of all important strategies at once, from which we can directly derive the answers to many different questions concerning the strategy analysis of a Büchi game.

5.5 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 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 $G = (V, V_0, V_1, E, F)$ be a Büchi game and $v \in G$ 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 [27], 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 [15, 27, 9] and extend our semiring by dual-indeterminates $\bar{X} = \{\bar{x} \mid x \in X\}$. The idea is to label a literal and its negation by corresponding indeterminates x and \bar{x} . We must then avoid monomials such as $x\bar{x}$, as they represent contradictory information. To this end, we consider the quotient of $\mathbb{S}^\infty[X \cup \bar{X}]$ with respect to the congruence generated by $x \cdot \bar{x} = 0$ for $x \in X$ and refer to the resulting quotient semiring as dual-indeterminate absorptive polynomials $\mathbb{S}^\infty[X, \bar{X}]$. This semiring inherits most of the properties of $\mathbb{S}^\infty[X]$. Most importantly, any assignment $h: X \cup \bar{X}$ that respects dual-indeterminates, i.e., $h(x) \cdot h(\bar{x}) = 0$, lifts to a fully-continuous homomorphism analogous to Theorem 17.

We then replace π_{strat} with an $\mathbb{S}^\infty[X, \bar{X}]$ -interpretation π_{rep}^\pm for $X = \{X_e \mid e \in E^\pm\}$: if $vw \in E^\pm$, we set $\pi_{\text{rep}}^\pm(Evw) = X_{vw}$ and $\pi_{\text{rep}}^\pm(\neg Evw) = \bar{X}_{vw}$, all other literals are mapped to 0 or 1 according to G . Notice that π_{rep}^\pm is not model-defining, but still satisfies $\pi_{\text{rep}}^\pm(L) \cdot \pi_{\text{rep}}^\pm(\neg L) = 0$ for all literals L .

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

Proposition 27. *In the above setting, the following holds:*

1. Let $m \in \pi_{\text{rep}}^\pm[\text{win}_0(v)]$ 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 $m \in \pi_{\text{rep}}^\pm[\text{win}_0(v)]$ 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 27, we illustrate the computation of minimal repairs in a small example.

Example 28. 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 $\mathbb{S}^\infty[X, \bar{X}]$ -interpretation π_{rep}^\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 \bar{X}_b for the repairs. \lrcorner

Proof of Proposition 27. We prove both statements by considering homomorphisms into the Boolean semiring \mathbb{B} . For the first statement, let $m \in \pi_{\text{rep}}^{\pm}[\text{win}_0(v)]$ be a monomial and let $h: X \cup \bar{X} \rightarrow \mathbb{B}$ 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 $X_e, \bar{X}_e \notin \text{var}(m)$ and $e \in E^+$ (do not add e without reason),
- $h(X_e) = \mathbf{1}$, if $X_e, \bar{X}_e \notin \text{var}(m)$ and $e \in E^-$ (do not remove e without reason).

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

For the second statement, let $R \subseteq E^{\pm}$ be a repair and consider the repaired game $G' \models \text{win}_0(v)$. As G' differs from G only by edges in E^{\pm} , there is a unique assignment $h: X \cup \bar{X} \rightarrow \mathbb{B}$ such that $h \circ \pi_{\text{rep}}^{\pm}$ corresponds to G' . Again, h lifts to a fully-continuous homomorphisms and we thus get $\mathbf{1} = h \circ \pi_{\text{rep}}^{\pm}[\text{win}_0(v)] = h(\pi_{\text{rep}}^{\pm}[\text{win}_0(v)])$. So there must be a monomial $m \in \pi_{\text{rep}}^{\pm}[\text{win}_0(v)]$ 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, $\bar{X}_e \in \text{var}(m) \cap X^{\pm}$ implies $h(\bar{X}_e) = \mathbf{1}$ 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 could drop exponents from $\mathbb{S}^{\infty}[X, \bar{X}]$ and work in the resulting, simpler semiring $\text{PosBool}[X, \bar{X}]$ (but again, the number of minimal repairs can be exponential in the size of the game). Further, the reverse analysis approach is not limited to questions about edges. We can also work with interpretations that track the target set F and thus answer questions such as *how to choose or modify the target set so that Player 0 wins?*

6 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 [9] 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. We remark that, of course, there are also limitations to this method and not all relevant questions about strategies can be answered directly by semiring valuations. As an example we show in the full version of this paper [14] that minimal cost computations provide serious obstacles to a semiring treatment.

The Sum-of-Strategies result motivates the general notion of absorption-dominant strategies which 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

- [1] C. Bourgaux, A. Ozaki, R. Peñaloza & L. Predoiu (2020): *Provenance for the Description Logic ELHr*. In: *Proceedings of IJCAI 2020*, pp. 1862–1869, doi:10.24963/ijcai.2020/258.
- [2] J. Bradfield & I. Walukiewicz (2018): *The mu-calculus and model checking*. In E. Clarke, T. Henzinger, H. Veith & R. Bloem, editors: *Handbook of Model Checking*, Springer, pp. 871–919, doi:10.1007/978-3-319-10575-8_26.
- [3] F. Canavoi, E. Grädel, S. Lessenich & W. Pakusa (2015): *Defining winning strategies in fixed-point logic*. In: *30th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pp. 366–377, doi:10.1109/LICS.2015.42.
- [4] K. Chatterjee, W. Dvorák, M. Henzinger & V. Loitzenbauer (2016): *Conditionally Optimal Algorithms for Generalized Büchi Games*. arXiv:1607.05850 [cs.DS]. Available at <http://arxiv.org/abs/1607.05850>.
- [5] K. Chatterjee & M. Henzinger (2014): *Efficient and Dynamic Algorithms for Alternating Büchi Games and Maximal End-Component Decomposition*. *J. ACM* 61(3), pp. 15:1–15:40, doi:10.1145/2597631.
- [6] K. Chatterjee, T. Henzinger & N. Piterman (2008): *Algorithms for Büchi Games*. arXiv:0805.2620 [cs.GT]. Available at <http://arxiv.org/abs/0805.2620>.
- [7] K. Dannert & E. Grädel (2019): *Provenance Analysis: A Perspective for Description Logics?* In C. Lutz et al., editor: *Description Logic, Theory Combination, and All That, Lecture Notes in Computer Science* 11560, Springer, pp. 266–285, doi:10.1007/978-3-030-22102-7_12.
- [8] K. Dannert & E. Grädel (2020): *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, Springer, pp. 53–79, doi:10.1007/978-3-030-64187-0_3.
- [9] K. Dannert, E. Grädel, M. Naaf & V. Tannen (2021): *Semiring Provenance for Fixed-Point Logic*. In C. Baier & J. Goubault-Larrecq, editors: *29th EACSL Annual Conference on Computer Science Logic*

- (CSL 2021), *Leibniz International Proceedings in Informatics (LIPIcs)* 183, Dagstuhl, pp. 17:1–17:22, doi:10.4230/LIPIcs.CSL.2021.17.
- [10] D. Deutch, T. Milo, S. Roy & V. Tannen (2014): *Circuits for Datalog Provenance*. In: *Proc. 17th International Conference on Database Theory ICDT*, OpenProceedings.org, pp. 201–212, doi:10.5441/002/icdt.2014.22.
- [11] J. Duparc (2003): *Positive games and persistent strategies*. In: *International Workshop on Computer Science Logic*, Springer, pp. 183–196, doi:10.1007/978-3-540-45220-1_17.
- [12] F. Geerts & A. Poggi (2010): *On database query languages for K-relations*. *Journal of Applied Logic* 8(2), pp. 173–185, doi:10.1016/j.jal.2009.09.001.
- [13] E. Grädel, P. G. Kolaitis, L. Libkin, M. Marx, J. Spencer, M. Y. Vardi, Y. Venema & S. Weinstein (2007): *Finite Model Theory and Its Applications*. Texts in Theoretical Computer Science. An EATCS Series, Springer, doi:10.1007/3-540-68804-8.
- [14] E. Grädel, N. Lücking & M. Naaf (2021): *Semiring Provenance for Büchi Games: Strategy Analysis with Absorptive Polynomials*. arXiv:2106.12892 [cs.LO]. Available at <https://arxiv.org/abs/2106.12892>.
- [15] E. Grädel & V. Tannen (2017): *Semiring Provenance for First-Order Model Checking*. arXiv:1712.01980 [cs.LO]. Available at <https://arxiv.org/abs/1712.01980>.
- [16] E. Grädel & V. Tannen (2020): *Provenance analysis for logic and games*. *Moscow Journal of Combinatorics and Number Theory* 9(3), pp. 203–228, doi:10.2140/moscow.2020.9.203. Preprint available at <https://arxiv.org/abs/1907.08470>.
- [17] E. Grädel, W. Thomas & T. Wilke, editors (2002): *Automata, Logics, and Infinite Games*. *Lecture Notes in Computer Science* 2500, Springer, doi:10.1007/3-540-36387-4.
- [18] T. Green, G. Karvounarakis & V. Tannen (2007): *Provenance semirings*. In: *Principles of Database Systems PODS*, pp. 31–40, doi:10.1145/1265530.1265535.
- [19] T. Green & V. Tannen (2017): *The Semiring Framework for Database Provenance*. In: *Proceedings of PODS*, ACM, pp. 93–99, doi:10.1145/3034786.3056125.
- [20] J. Marcinkowski & T. Truderung (2002): *Optimal complexity bounds for positive LTL games*. In: *International Workshop on Computer Science Logic*, Springer, pp. 262–275, doi:10.1007/3-540-45793-3_18.
- [21] M. Mohri (2002): *Semiring frameworks and algorithms for shortest-distance problems*. *Journal of Automata, Languages and Combinatorics* 7(3), pp. 321–350, doi:10.25596/jalc-2002-321.
- [22] M. Naaf (2021): *Computing Least and Greatest Fixed Points in Absorptive Semirings*. arXiv:2106.00399 [cs.LO]. Available at <https://arxiv.org/abs/2106.00399>.
- [23] A. Ozaki & R. Peñaloza (2018): *Provenance in Ontology-based Data Access*. In: *Description Logics 2018, CEUR Workshop Proceedings* 2211. Available at <http://ceur-ws.org/Vol-2211/paper-28.pdf>.
- [24] M. Raghothaman, J. Mendelson, D. Zhao, M. Naik & B. Scholz (2020): *Provenance-guided synthesis of Datalog programs*. *Proc. ACM Program. Lang.* 4, pp. 62:1–62:27, doi:10.1145/3371130.
- [25] P. Senellart (2017): *Provenance and Probabilities in Relational Databases: From Theory to Practice*. *SIGMOD Record* 46(4), pp. 5–15, doi:10.1145/3186549.3186551.
- [26] I. Walukiewicz (2002): *Monadic second-order logic on tree-like structures*. *Theoretical Computer Science* 275, pp. 311–346, doi:10.1016/S0304-3975(01)00185-2.
- [27] J. Xu, W. Zhang, A. Alawini & V. Tannen (2018): *Provenance Analysis for Missing Answers and Integrity Repairs*. *IEEE Data Eng. Bull.* 41(1), pp. 39–50. Available at <http://sites.computer.org/debull/A18mar/p39.pdf>.