Rheinisch-Westfälische Technische Hochschule Aachen
Lehr- und Forschungsgebiet
Mathematische Grundlagen der Informatik

Semirings for Provenance Analysis of Fixed-point Logics and Games

Matthias Ferdinand Naaf

Masterarbeit
im Studiengang Informatik

vorgelegt der
Fakultät für Mathematik, Informatik und Naturwissenschaften der
Rheinisch-Westfälischen Technischen Hochschule Aachen

im Juni 2019

Gutachter: Prof. Dr. Erich Grädel
Prof. Dr. Martin Grohe

Abstract

Semiring provenance originated in database theory and was recently applied to logic with the goal of answering questions such as "why does a formula hold in a particular model?", "how many proofs does it have?" or, if we only have low confidence in certain facts, "what is the overall confidence in the formula?". This is achieved by evaluating formulae in certain semirings instead of using standard truth values.

This thesis defines semiring semantics for least fixed-point logic (LFP) including both least and greatest fixed points. We show that for certain semirings, these semantics can be seen as an extension of standard semantics by multiple truth values which we interpret as provenance information. A particular focus is on the algebraic and order-theoretic properties of semirings that lead to reasonable semantics, including the interplay with negation, duality of least and greatest fixed points and homomorphisms.

We survey several candidates of provenance semirings for the analysis of LFP. Our results show that (generalized) absorptive polynomials $\mathbb{S}^\infty[X]$ provide reasonable information and have an important universal property in terms of provenance-preserving homomorphisms, rendering them the most interesting semiring for LFP. We characterize $\mathbb{S}^\infty[X]$-semantics by means of winning strategies in the associated model checking game. Intuitively, this result shows that provenance analysis in $\mathbb{S}^\infty[X]$ (and, due to universality, in many absorptive semirings) provides useful information about shortest proofs.

Contents

1 Introduction **1**

2 Semiring Semantics for First-Order Logic **5**
2.1 Semirings . . . 5
2.2 First-Order Logic . . . 7

3 Towards Fixed-Point Logic **11**
3.1 Least Fixed-Point Logic . . . 11
3.2 Lifting Semiring Semantics . . . 12

4 Semirings and Fixed Points **17**
4.1 Order Theory . . . 17
4.2 Suitable Semirings . . . 20
4.3 Examples . . . 22
4.4 Semiring Properties . . . 24
4.5 Related Concepts . . . 29
4.6 Homomorphisms and Function Semirings . . . 32

5 Semiring Semantics for Fixed-Point Logic **35**
5.1 Definition . . . 35
5.2 Semiring vs. Standard Semantics . . . 37
5.3 Duality . . . 40
5.4 Fundamental Property . . . 45
5.5 Logical Fragments and Continuity . . . 47

6 Provenance Semirings **55**
6.1 Counting . . . 56
6.2 Formal Power Series . . . 57
6.3 Why-Semiring . . . 60
6.4 Absorptive Polynomials . . . 63
6.5 PosBool-Semiring . . . 82
6.6 Polynomials in Dual Indeterminates . . . 83

7 Understanding Provenance via Games **93**
7.1 Preliminaries: Parity Games . . . 93
7.2 Excursion: Strategy Computation with Semirings . . . 95
7.3 Model Checking Games . . . 99
7.4 Characterization for Absorptive Polynomials . . . 103
7.5 Thoughts Beyond Absorption . . . 121

8 Conclusion **125**

1 Introduction

Provenance analysis is concerned with understanding the information flow through computations such as the evaluation of logical formulae. The idea to use semirings for this purpose originated in database theory where, given a relational database and some query, one wants to understand *why* a particular tuple satisfies the query, *how* an answer was established or from *which* facts a certain tuple was derived. Put differently, the goal is to obtain information about the origin of a particular answer – its provenance. Several approaches for data provenance have been developed, for instance in [BKWC01]. The seminal work of Green, Karvounarakis and Tannen [GKT07] unified many of these approaches by introducing the framework of **semiring provenance for databases**.

The central idea is to annotate facts with values from a *commutative semiring* in order to track their provenance. Semirings are algebraic structures with two binary operations, addition and multiplication, that satisfy basic laws of algebra. Compared to the more widely used algebraic structures of rings and fields, they require neither inverse elements for multiplication nor for addition. Semirings thus provide a high degree of abstraction which has already proven useful in formal language theory [DK09] and to generalize constraint satisfaction problems [BMR97]. Green et al. use them to generalize the evaluation of queries: Starting with an *annotation* of database entries with values from a commutative semiring, these values are then updated using the semiring operations alongside the evaluation of the query. In the end, each answer tuple is annotated with a semiring value that provides information about the provenance of this particular answer. By using different semirings, many provenance applications can be covered. For instance, computations in the Boolean semiring $\mathbb{B}$ are equivalent to *set semantics* while natural numbers $\mathbb{N}$ count multiplicities and lead to *bag semantics*. The most interesting semirings are polynomials, as they allow to *track* the influence of certain facts by annotating them with variables. For an overview of semiring provenance for databases and an extensive list of references, we refer to [GT17b]. A key insight of Green et al. is that *semiring homomorphisms* preserve provenance computations for positive relational algebra. Polynomials $\mathbb{N}[X]$ thus provide the *most general* provenance information due to their *universal property*. For datalog queries, Green et al. propose formal power series $\mathbb{N}^\infty[\![X]\!]$ as completion of $\mathbb{N}[X]$.

Recently, Grädel and Tannen applied semiring provenance to logic [GT17a]. Instead of a database query, a finite model $\mathfrak{A}$ and a formula φ are given, say in first order-logic. **Semiring provenance for logic** is then concerned with the question *how* the truth of φ in $\mathfrak{A}$ can be established. An important detail is that database queries are usually positive formulae whereas first-order logic admits negation. In semiring provenance, we always view the value 0 as *false* and all other values as *nuances of true*, which has the consequence that it is not possible to interpret negation directly. Instead, Grädel and Tannen define semiring semantics via the *negation normal form*. Given a formula φ in negation normal form, the idea is, in analogy to database provenance, to replace the evaluation of φ in terms of truth values with computations in a commutative semiring. Starting from an annotation of literals $R\mathbf{a}$ and $\neg R\mathbf{a}$ with semiring values, addition is used to interpret *alternative use of information* such as disjunction and existential quantification, while multiplication is used for *joint use of information* in conjunctions and universal quantification. A common application is to count proofs using the semiring of natural numbers. As an example, consider the following formula on the given model. We annotate all literals Exy by 0 or 1 depending on whether the corresponding edge is present in the graph.

$$\varphi(u, v) = \exists y(Euy \wedge Eyv)$$

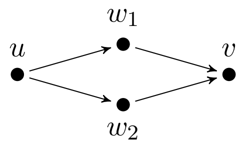

The formula $Euy \wedge Eyv$ then evaluates to $1 \cdot 1 = 1$ if we set y to w_1 or w_2, and it evaluates to 0 otherwise. Using the sum over all choices to interpret the existential quantification results in the value $1 + 1 + 0 + 0 = 2$ for $\varphi(u, v)$, which indeed is the number of proofs for $\varphi(u, v)$ in the given graph. Different semirings allow for further applications:

- Computations in the Boolean semiring $\mathbb{B}$ are equivalent to standard truth semantics.
- When we annotate literals with *confidence scores* or *access levels* (both semirings are introduced later on), we can compute the overall confidence in the formula φ or the minimal access level required to establish the truth of φ in $\mathfrak{A}$.
- Polynomial semirings such as $\mathbb{N}[X]$ allow us to *track* certain literals by mapping them to variables. The resulting polynomial provides information about the different proofs of φ and tells us in which proofs the tracked literals occur.

For reasonable literal annotations, the evaluation in a semiring can be regarded as a generalization of standard semantics by several truth values similar to the idea of *many-valued logics*. For example, the semiring $(\{0, \frac{1}{2}, 1\}, \max, \min, 0, 1)$ can be interpreted as providing a ternary truth value $\frac{1}{2}$. However, our aim is not merely to *extend* standard semantics, but rather to interpret the additional semiring values as provenance information in order to better *understand* standard semantics.

In [GT17a], Grädel and Tannen define semiring provenance for first-order logic and show that it has reasonable properties. They introduce *dual-indeterminate polynomials* $\mathbb{N}[X, \bar{X}]$ to represent negative information and observe that provenance computations are closely related to *proof trees*. In the more recent work [GT19], they extend semiring provenance to reachability and safety games as well as positive least fixed-point logic (posLFP). In particular, they show that one can equivalently define semiring provenance for posLFP via the evaluation of formulae and via the corresponding *model checking games*. Grädel and Tannen further introduce the semiring $\mathbb{S}^{\infty}[X]$ of generalized absorptive polynomials which plays a central role in this thesis. This semiring is also used in [Mrk18] which discusses provenance analysis of temporal logics such as CTL, using ω-continuous semirings for the positive fragment and absorptive polynomials for full CTL including greatest fixed points. Related work considers semiring provenance for guarded logics [DG19] and logics with team semantics [Huw18].

In this thesis, we are concerned with semiring provenance for least fixed-point logic including both least and greatest fixed points. Our overall focus is on the algebraic and order theoretic properties of semirings that lead to useful provenance information. As an example of the challenges arising from fixed points, consider the following two informally stated formulae on the given graph.

$$\begin{aligned} \varphi_1(u,v) &\coloneqq \text{“there is a path from } u \text{ to } v\text{”} \\ \varphi_2(u) &\coloneqq \text{“there is an infinite path from } u\text{”} \end{aligned}$$

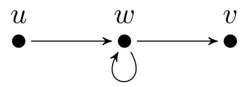

Formula φ_1 can easily be described by a least fixed point. Interpreting $\varphi_1(u,v)$ in the semiring $\mathbb{N}$ by annotating all edges with 1 yields the number of paths from u to v. However, the graph above contains a cycle, so there are infinitely many such paths (allowing node repetitions). Indeed, the fixed-point iteration induced by $\varphi_1(u,v)$ in $\mathbb{N}$ does not terminate and no fixed-point is reached. We thus have to complete $\mathbb{N}$ to $\mathbb{N}^\infty$ by adding a greatest element ∞ which is then the fixed-point corresponding to $\varphi_1(u,v)$.

In general, we define semiring semantics for LFP similar to [GT19] and show that a certain form of *chain-completeness* is sufficient to obtain well-defined semantics. This results in a more general definition than the one for CTL in [Mrk18] which introduced the rather restrictive notion of *absorptive lattice semirings*.

In the above example, the value ∞ is reasonable for $\varphi_1(u,v)$ due to the infinite number of paths. If we also allow greatest fixed points, we often obtain unexpected results. For example, $\varphi_2(u)$ can be expressed by a greatest fixed point and it is clear that there is only one infinite path from u. However, evaluating $\varphi_2(u)$ in $\mathbb{N}^\infty$ yields the value ∞ which does not correspond to the number of infinite paths. Even worse, using formal power series $\mathbb{N}^\infty[\![X]\!]$ (which work well for datalog [GKT07] and posLFP [GT19]) and annotating all edges by variables results in the overall value of 0 (which we usually interpret as false), although $\varphi_2(u)$ clearly holds in the given graph.

To prevent such inconsistencies with standard semantics, we impose additional requirements on the semirings and then discuss how the results of [GT17a] for FO can be lifted to LFP, including the compatibility with models and the interplay with homomorphisms and negation. The latter is especially interesting for LFP, as it comprises the *duality* of least and greatest fixed points. As seen for $\varphi_2(u)$ above, greatest fixed points lead to a number of questions. Which polynomial semirings can we use in place of $\mathbb{N}^\infty[\![X]\!]$ to obtain useful provenance information? Going further, can we characterize provenance in these semirings in the same way that $\mathbb{N}[X]$-provenance for FO corresponds to proof trees? While we can easily provide finite witnesses for the truth of $\varphi_1(u)$, how can we describe proofs of formulae like $\varphi_2(u)$ involving greatest fixed points? We show that these questions have positive answers for *absorptive* semirings when we consider strategies in *model checking games* as proofs or witnesses of formulae.

The **organization of this thesis** is as follows. In chapter 2, we recapitulate semiring provenance for FO as a reference point. Chapter 3 then introduces LFP and poses challenges and questions arising from the interpretation of fixed-point formulae. The following chapter 4 addresses these questions by introducing the required concepts for semirings, most importantly chain-completeness and fixed-point theorems, including many examples. The semantics in these semirings are then studied in chapter 5 in terms of their relation to standard semantics, homomorphisms and duality. Chapter 6 discusses provenance semirings based on polynomials with a focus on absorptive polynomials $\mathbb{S}^\infty[X]$ and their universal property. Computations in $\mathbb{S}^\infty[X]$ are then characterized in chapter 7 by means of strategies in model checking games, before we conclude in chapter 8 and outline topics for future work.

2 Semiring Semantics for First-Order Logic

This chapter introduces semiring interpretations for formulae which lay the foundation for provenance analysis. The idea is to map the logical connectives $\vee$ and $\wedge$ to the semiring operations $+$ and $\cdot$, thereby replacing evaluation in terms of truth values with computations in semirings. The choice of the semiring then determines which information the evaluation yields – from truth values to confidence scores or counting proofs.

2.1 Semirings

Semirings are algebraic structures with two operations, usually denoted by $+$ and $\cdot$, that satisfy common laws of algebra. We follow the terminology of [GT17a], more background on semirings can be found in [DK09, Gol99]. For the definition, we use the notion of *monoids* which are algebraic structures with an associative binary operation and a neutral element. A monoid is commutative if the operation is commutative.

Definition 2.1. A *commutative semiring* is an algebraic structure $(S, +, \cdot, 0, 1)$ with $0 \neq 1$ such that $(S, +, 0)$ and $(S, \cdot, 1)$ are commutative monoids and additionally,

(1) distrubutivity: $a \cdot (b + c) = ab + ac$, for all $a, b, c \in S$,

(2) 0 is annihilating: $0 \cdot a = 0$, for all $a \in S$.

In the following, semiring always refers to a commutative semiring. The reason to require commutativity is that we want to interpret the formulae $\varphi \wedge \vartheta$ and $\vartheta \wedge \varphi$ in the same way. In addition, commutativity simplifies provenance computation and allows us to use polynomial semirings. We write S instead of $(S, +, \cdot, 0, 1)$ if the context is clear. In many cases, we consider semirings with additional properties:

Definition 2.2. A semiring S is said to be:

- *idempotent*, if $a + a = a$ for all $a \in S$,
- *absorptive*, if $a + ab = a$ for all $a, b \in S$,
- *multiplicative idempotent*, if $a \cdot a = a$ for all $a \in S$.

These subclasses of semirings are especially important later on for the analysis of fixed-point logics. For now, we only note that every absorptive semiring is also idempotent and that both properties correspond to logical equalities, i.e., $\varphi \vee \varphi \equiv \varphi$ and $\varphi \vee (\varphi \wedge \vartheta) \equiv \varphi$. More interesting for first-order logic is the notion of *positive* semirings which ensures that computations with positive values remain positive.

Definition 2.3. Let S be a semiring. Then

- S is *+-positive* if $a + b = 0$ implies $a = 0$ and $b = 0$, for all $a, b \in S$;
- S has *divisors of* 0 if there are $a, b \in S$ with $a, b \neq 0$ and $ab \neq 0$;
- S is *positive* if it is +-positive and has no divisors of 0.

Two examples of semirings are the Boolean semiring $(\mathbb{B}, \vee, \wedge, \bot, \top)$ with $\mathbb{B} = \{\bot, \top\}$ which represents truth values and the semiring of natural numbers $(\mathbb{N}, +, \cdot, 0, 1)$ for counting. Both are positive, but only $\mathbb{B}$ is absorptive and has idempotent operations. An extensive list of examples is presented in section 4.3.

An important aspect of using semirings is that we can compute provenance information in a general way with polynomial semirings, thereby covering different applications with a single computation. To achieve this, we use homomorphisms to switch from universal semirings to application-specific ones.

Definition 2.4. A *semiring homomorphism* is a function $h : S \to T$ on semirings S, T such that:

(1) h preserves neutral elements: $h(0) = 0$ and $h(1) = 1$,

(2) h is additive: $h(a + b) = h(a) + h(b)$ for all $a, b \in S$,

(3) h is multiplicative: $h(ab) = h(a) \cdot h(b)$ for all $a, b \in S$.

To close the introduction of semirings, let us give an example of an important canonical homomorphism which connects semiring elements and truth values.

Example 2.5. Let S be a semiring and consider the function $\dagger_S : S \to \mathbb{B}$ defined by

$$\dagger_S(a) = \begin{cases} \top, & \text{if } a \neq 0 \\ \bot, & \text{if } a = 0 \end{cases}$$

which we call *truth projection*. Then $\dagger_S$ is a semiring homomorphism if, and only if, S is positive [GT17a].

2.2 First-Order Logic

Following the notion of K-interpretations[1] from [GT17a], we show how one can use semirings to interpret first-order logic. We give an overview of the main results in [GT17a] concerning the properties of the resulting semantics, which serves as a reference point when discussing semiring interpretations for fixed-point logic.

Notation

For the remainder of this work, fix a finite relational signature $\tau = \{R, P, \dots\}$, a set of variables $\mathcal{V}$ and a *finite* non-empty universe A. We usually denote elements of A by $a, b, \dots$ and tuples from A^k (for some k) by bold symbols $\mathbf{a}, \mathbf{b}$. Similarly, we use x, y, z for variables and $\mathbf{x}, \mathbf{y}, \mathbf{z}$ for variable tuples. Semiring semantics are based on the interpretation of literals, so given A and τ, we define

$$\mathsf{Lit}_{A,\tau} = \{R\mathbf{a}, \neg R\mathbf{a} \mid R \text{ a } k\text{-ary relation symbol from } \tau,\ \mathbf{a} \in A^k\}$$

Note that $R\mathbf{a} \in \mathsf{Lit}_A$ is an *evaluated* literal, as $\mathbf{a}$ is an element of A^k and not a variable tuple. To ease notation, we always identify $R\mathbf{a}$ and $\neg\neg R\mathbf{a}$ and often write L to refer to either $R\mathbf{a}$ or $\neg R\mathbf{a}$. We usually assume τ to be fixed and thus write Lit_A for $\mathsf{Lit}_{A,\tau}$.

A (variable) valuation is a mapping $\alpha : \mathcal{V} \to A$. The valuation $\alpha[x/a]$ is like α except that it maps x to a (for some $x \in \mathcal{V}$, $a \in A$). We use the notation $\alpha[\mathbf{x}/\mathbf{a}]$ to replace each variable in the tuple $\mathbf{x}$ with the corresponding value in $\mathbf{a}$. We use standard syntax (and semantics) for FO and denote the negation normal form of a formula φ (where negations only appear in literals) by $\mathsf{nnf}(\varphi)$.

S-Interpretations

We are now ready to define semiring interpretations. In standard semantics, a k-ary relation symbol is interpreted by a subset of A^k which one can also view as a mapping $A^k \to \mathbb{B}$ (where elements of the subset are mapped to $\top$). Semiring interpretations provide a mapping $A^k \to S$ for an arbitrary semiring S and extend the classic view in two ways. We treat 0 as *false* and all other semiring values as *true*, so we have different *nuances of truth* available which can carry additional information. The second generalization is that we allow inconsistent interpretations of literals. That is, opposing literals are interpreted independent of each other, so we may assign both to positive values (or both to 0). However, we often want to exclude such interpretations in the context of provenance analysis, as they do not correspond to actual models.

[1] We call them S-interpretations (for *semiring*) and usually denote them by ℓ (for *literal mapping*).

Definition 2.6. Given a commutative semiring S, an S-*interpretation* is a literal mapping

$$\ell : \mathsf{Lit}_A \to S$$

We lift ℓ to FO formulae together with a variable valuation $\alpha : \mathcal{V} \to A$:

$$\llbracket R\mathbf{x} \rrbracket_\ell^\alpha = \ell(R\,\alpha(\mathbf{x})) \qquad \llbracket \neg R\mathbf{x} \rrbracket_\ell^\alpha = \ell(\neg R\,\alpha(\mathbf{x}))$$

$$\llbracket \varphi \vee \vartheta \rrbracket_\ell^\alpha = \llbracket \varphi \rrbracket_\ell^\alpha + \llbracket \vartheta \rrbracket_\ell^\alpha \qquad \llbracket \varphi \wedge \vartheta \rrbracket_\ell^\alpha = \llbracket \varphi \rrbracket_\ell^\alpha \cdot \llbracket \vartheta \rrbracket_\ell^\alpha$$

$$\llbracket \exists x\, \varphi \rrbracket_\ell^\alpha = \sum_{a \in A} \llbracket \varphi \rrbracket_\ell^{\alpha[x/a]} \qquad \llbracket \forall x\, \varphi \rrbracket_\ell^\alpha = \prod_{a \in A} \llbracket \varphi \rrbracket_\ell^{\alpha[x/a]}$$

where we handle negation via negation normal form and interpret equality atoms by 0 or 1 according to their truth value in standard semantics:

$$\llbracket x \sim y \rrbracket_\ell^\alpha = \begin{cases} 1, & \text{if } \alpha(x) \sim \alpha(y) \\ 0, & \text{otherwise} \end{cases} \qquad \text{where } \sim \in \{=, \neq\}$$

$$\llbracket \neg\varphi \rrbracket_\ell^\alpha = \llbracket \mathsf{nnf}(\neg\varphi) \rrbracket_\ell^\alpha$$

We omit α and write $\llbracket \varphi \rrbracket_\ell$ if α is irrelevant, for instance if φ is a sentence. For a formula $\varphi(x)$ with a free variable x, we further write $\llbracket \varphi(a) \rrbracket_\ell^\alpha$ for $\llbracket \varphi \rrbracket_\ell^{\alpha[x/a]}$ to simplify notation.

Remark: From the point of view of semiring semantics, there is no need to treat equality atoms in a special way – we could as well include them in Lit_A (and assign semiring values via ℓ). We therefore do not explicitly consider equality atoms from now on to simplify proofs. The reason to define them as above is that we usually do not want to track equalities in provenance analysis.

To establish a correspondence between S-interpretations and standard semantics, Grädel and Tannen [GT17a] consider the **canonical truth interpretation** $\ell_{\mathfrak{A}}$ and the **canonical counting interpretation** $\ell_{\#\mathfrak{A}}$ for a given τ-structure $\mathfrak{A}$, defined by

$$\ell_{\mathfrak{A}} : \mathsf{Lit}_A \to \mathbb{B},\ \ell_{\mathfrak{A}}(L) = \begin{cases} \top, & \text{if } \mathfrak{A} \models L \\ \bot, & \text{otherwise} \end{cases}$$

$$\ell_{\#\mathfrak{A}} : \mathsf{Lit}_A \to \mathbb{N},\ \ell_{\#\mathfrak{A}}(L) = \begin{cases} 1, & \text{if } \mathfrak{A} \models L \\ 0, & \text{otherwise} \end{cases}$$

and they indeed show that $\ell_{\mathfrak{A}}$ corresponds to standard semantics while $\ell_{\#\mathfrak{A}}$ counts the number of proofs using the additional values in $\mathbb{N}$.

Proposition 2.7. *Let $\mathfrak{A}$ be a τ-structure with universe A and φ an FO sentence. Then*

(1) $[\![\varphi]\!]_{\ell_{\mathfrak{A}}} = \top \iff \mathfrak{A} \models \varphi$,

(2) $[\![\varphi]\!]_{\ell_{\#\mathfrak{A}}}$ *equals the number of different proof trees witnessing* $\mathfrak{A} \models \varphi$.

The other direction, from S-interpretations to models, does not hold in general, as S-interpretations may assign literals in an inconsistent way.

Definition 2.8. Let ℓ be an S-interpretation. Then S is said to be

- *consistent*, if $\ell(L) = 0$ or $\ell(\neg L) = 0$ for every literal $L \in \mathsf{Lit}_A$,
- *complete*, if $\ell(L) \neq 0$ or $\ell(\neg L) \neq 0$ for every literal $L \in \mathsf{Lit}_A$,
- *model-defining*, if ℓ is both consistent and complete.

In a positive semiring, we can equivalently say that ℓ is model-defining if $\ell(L) \cdot \ell(\neg L) = 0$ and $\ell(L) + \ell(\neg L) \neq 0$ for all literals. A model-defining S-interpretation uniquely defines the **induced model** $\mathfrak{A}_\ell$ by means of

$$\mathfrak{A}_\ell \models L \quad \iff \quad \ell(L) \neq 0$$

for all literals $L \in \mathsf{Lit}_A$. If S is positive, this property lifts to FO sentences.

Proposition 2.9. *Let ℓ be a model-defining S-interpretation for a positive semiring S. Then $\mathfrak{A}_\ell \models \varphi$ if, and only if, $[\![\varphi]\!]_\ell \neq 0$ for any FO sentence φ.*

Another way to think about this result is by means of homomorphisms. If S is positive, the truth projection $\dagger_S : S \to \mathbb{B}$ is a homomorphism. Following the above definitions, one can see that

$$\dagger_S \circ \ell = \ell_{\mathfrak{A}_\ell} \qquad \text{and hence} \qquad [\![\varphi]\!]_{\dagger_S \circ \ell} = \top \iff \mathfrak{A}_\ell \models \varphi$$

where $\ell_{\mathfrak{A}_\ell}$ is the canonical truth interpretation for the model induced by ℓ. The composition into $\dagger_S \circ \ell$ has the convenient property that it lifts to formulae:

$$[\![\varphi]\!]_{\dagger_S \circ \ell} = \dagger_S\big([\![\varphi]\!]_\ell\big)$$

In other words, we can first compute the interpretation of φ in the semiring S and later switch to $\mathbb{B}$ via the homomorphism $\dagger_S$. As $\dagger_S$ is the *truth* projection, this means that computations in S *preserve truth*. In [GT17a], Grädel and Tannen state the following general formulation of this composition that allows us to compute provenance in a general way and later switch to specific semirings via homomorphisms. We prove an adaption for fixed-point logic in chapter 5.

Theorem 2.10 (fundamental property)**.** *Let $h : S \to T$ be a semiring homomorphism and $\ell : \mathsf{Lit}_A \to S$ be an S-interpretation. Then $h \circ \ell : \mathsf{Lit}_A \to T$ is a T-interpretation and we have $h(\llbracket \varphi \rrbracket_\ell) = \llbracket \varphi \rrbracket_{h\circ\ell}$ for all FO sentences φ.*

$$
\begin{tikzcd}
 & \mathsf{Lit}_A \arrow[dl, "\ell"'] \arrow[dr, "h\circ\ell"] & \\
S \arrow[rr, "h"] & & T
\end{tikzcd}
\quad\Longrightarrow\quad
\begin{tikzcd}
 & \mathrm{FO} \arrow[dl, "\ell"'] \arrow[dr, "h\circ\ell"] & \\
S \arrow[rr, "h"] & & T
\end{tikzcd}
$$

To close the overview, let us formulate the following consistency and completeness results from [GT17a]. These simple observations show that (reasonable) semiring interpretations continue to make sense in the presence of negation which justifies the definition based on the negation normal form.

Proposition 2.11. *Let $\ell : \mathsf{Lit}_A \to S$ be an S-interpretation. Then the following holds:*

(1) *If ℓ is consistent, then $\llbracket \varphi \rrbracket_\ell = 0$ or $\llbracket \neg\varphi \rrbracket_\ell = 0$ for all FO sentences φ.*

(2) *If $\ell(L) \cdot \ell(\neg L) = 0$ for all $L \in \mathsf{Lit}_A$, then $\llbracket \varphi \rrbracket_\ell \cdot \llbracket \neg\varphi \rrbracket_\ell = 0$ for all FO sentences φ.*

If S is positive, the above statements are equivalent and we additionally have the following completeness property:

(3) *If S is positive and ℓ is complete, then $\llbracket \varphi \rrbracket_\ell + \llbracket \neg\varphi \rrbracket_\ell \neq 0$ for all FO sentences φ.*

As a conclusion, we can say that semiring semantics in terms of model-defining interpretations have similar properties as standard semantics, but generalize the latter by providing several truth values. We thus obtain additional information which can give insights into why or how a formula is satisfied. The fundamental property is a very important result as it allows to unify provenance computations by working with universal semirings such as polynomials $\mathbb{N}[X]$. From a theoretical point of view, it is also a strong indication that semiring semantics yield a generalization with desirable properties.

3 Towards Fixed-Point Logic

In this chapter, we extend FO by a fixed-point operator which leads to least fixed-point logic (LFP). After providing a definition of semiring interpretations in analogy to standard semantics, we pose several questions and requirements on the semirings that need to be fulfilled to obtain well-defined semantics with reasonable properties. These questions are addressed in the subsequent chapters.

3.1 Least Fixed-Point Logic

Least fixed-point logic extends FO by a least fixed-point operator lfp. As FO is closed under negation, LFP can also express greatest fixed points due to the well-known duality between the two notions of fixed points. As for FO, we work with formulae in negation normal form and therefore express greatest fixed-points explicitly using the gfp operator to avoid negations. We define the syntax and standard semantics below, more background on LFP can, for instance, be found in [GKL+07, chapter 3].

Recall that we always assume a relational signature τ and a finite universe A. Formulae with fixed-point operators are written as

$$\varphi(\mathbf{y}) = [\mathsf{lfp}\, R\,\mathbf{x}.\ \vartheta](\mathbf{y}) \qquad \text{or} \qquad \varphi(\mathbf{y}) = [\mathsf{gfp}\, R\,\mathbf{x}.\ \vartheta](\mathbf{y})$$

where $\mathbf{x}$ is a k-tuple of pairwise different variables, $\mathbf{y}$ is a k-tuple of variables[2], R is a k-ary relation symbol and ϑ is a formula over the signature $\tau \cup \{R\}$. Both $\mathbf{x}$ and R may occur in ϑ (we sometimes write $\vartheta(R, \mathbf{x})$ to emphasize this), but R must only occur positively (i.e., not behind negations). For ease of presentation, we assume that all lfp- and gfp-subformulae of an LFP formula use different names for the relation symbol R. We write $[\mathsf{fp}\, R\, x.\ \vartheta](\mathbf{y})$ when we want to refer to both kinds of fixed-point formulae.

Given a structure $\mathfrak{A}$, Each formula $\vartheta(R, \mathbf{x})$ induces an *update operator*

$$F^\vartheta : \mathcal{P}(A^k) \to \mathcal{P}(A^k), \quad R \mapsto \{\mathbf{a} \mid \mathfrak{A} \models \vartheta(R, \mathbf{a})\}$$

where k is the arity of the relation symbol R. If R only occurs positively in ϑ, then F^ϑ is monotone and thus has a least (and greatest) fixed point on the complete lattice $\mathcal{P}(A^k)$, which we denote by $\mathsf{lfp}\, F^\vartheta$ (and $\mathsf{gfp}\, F^\vartheta$). We then define

$$\mathfrak{A} \models [\mathsf{lfp}\, R\,\mathbf{x}.\ \vartheta](\mathbf{a}) \quad \Longleftrightarrow \quad \mathbf{a} \in \mathsf{lfp}\, F^\vartheta$$

The fixed-point can also be described iteratively. Starting with $R_0 = \emptyset$, we set $R_{\beta+1} = F^\vartheta(R_\beta)$ and $R_\lambda = \bigcup_{\beta<\lambda} R_\beta$ for ordinals $\beta \in \mathsf{On}$ and limit ordinals $\lambda \in \mathsf{On}$ (where On

[2] When allowing non-relational signatures, $\mathbf{y}$ may be a tuple of arbitrary terms.

denotes the class of ordinals). Due to monotonicity of F^ϑ, this defines an ascending sequence. As A^k is a set, this *fixed-point iteration* always terminates for some ordinal α which is called the *closure ordinal* of F^ϑ or ϑ (if A^k is finite as in our case, then α is finite as well). It follows by monotonicity of F^ϑ that $\mathsf{lfp}\, F^\vartheta = R_\alpha$. An analogous iteration is possible for $\mathsf{gfp}\, F^\vartheta$, starting with $R_0 = A^k$ and obtaining a decreasing sequence.

Example 3.1. A simple property that is expressible in LFP but not in FO is graph reachability in a directed graph with edge relation E. The following formula asserts that there is a (possibly empty) path from u to v:

$$\varphi_{\text{path}}(u, v) = [\mathsf{lfp}\, R\, x.\ x = v \lor \exists y(Exy \land R\, y)](u)$$

The easiest way to reason about LFP formulae is by the fixed-point iteration:

- $R_0 = \emptyset$,
- $R_1 = \{v\}$,
- $R_2 = \{v\} \cup \{u \mid u \to v\} = \{u \mid u \xrightarrow{<2} v\}$
- $R_n = \{u \mid u \xrightarrow{<n} v\}$
- $R_\omega = \{u \mid \text{there is a path from } u \text{ to } v\}$

where $u \xrightarrow{<n} v$ means that there is a (possibly empty) path from u to v with less than n edges. We reconsider this example with semiring semantics below.

3.2 Lifting Semiring Semantics

In order to lift S-interpretations from FO to LFP, let us first take a different view on the semantics defined above. F^ϑ operates on subsets of A^k which we can alternatively describe by mappings $A^k \to \mathbb{B}$. We now generalize this to mappings $A \to S$ into an arbitrary semiring S.

Definition 3.2. A k-ary S-*valuation* is a function $\pi : A^k \to S$ mapping k-tuples from the universe to values in a semiring S. We denote their set as $\mathsf{Val}_{S,k}$ or just Val_k (if S is clear from the context).

In analogy to standard semantics, we define the update operator F^ϑ_ℓ of a formula $\vartheta(R, \mathbf{x})$ under an S-interpretation ℓ. Instead of sets, F^ϑ_ℓ now operates on valuations, but is otherwise define in the same way.

Definition 3.3. Let $\vartheta(R, \mathbf{x})$ be a formula over the signature $\tau \cup \{R\}$ where R and $\mathbf{x}$ have arity k and R only occurs positively. Given a variable valuation α and an S-interpretation ℓ, this induces the following *update operator* $F_\ell^{\vartheta,R,\mathbf{x},\alpha}$ or simply F_ℓ^ϑ (as $R, \mathbf{x}, \alpha$ are usually clear from the context):

$$F_\ell^\vartheta : \mathsf{Val}_k \to \mathsf{Val}_k, \quad \pi \mapsto \pi'$$

where π' is the S-valuation defined by

$$\pi' : A^k \to S, \quad \mathbf{a} \mapsto [\![\vartheta]\!]_{\ell[R/\pi]}^{\alpha[\mathbf{x}/\mathbf{a}]}$$

Here, $\ell[R/\pi]$ denotes the S-interpretation that maps literals $R\mathbf{a}$ to $\pi(\mathbf{a})$ (for all $\mathbf{a} \in A^k$) and otherwise behaves like ℓ. Note that the value of $\neg R\mathbf{a}$ is not affected; it is irrelevant, as we require that R only occurs positively in ϑ.

To simplify notation in proofs when referring to π', we write $[\![\vartheta]\!]_\ell^\bullet$ for the mapping $\mathbf{a} \mapsto [\![\vartheta]\!]_\ell^{\alpha[\mathbf{x}/\mathbf{a}]}$ if α and $\mathbf{x}$ are clear from the context.

With this notation, we can lift S-interpretations to LFP formulae as follows. For the definition to be well-defined, we have to ensure that the required fixed points always exist. This imposes several requirements on the semiring which we make precise in the following. For now, assume that we have a suitable semiring S together with some order (which induces a pointwise order on valuations) in which the required fixed points exist.

Definition 3.4. Let ℓ be an S-interpretation for a *suitable* semiring S and α a valuation. We define $[\![\varphi]\!]_\ell^\alpha$ inductively as in definition 2.6 and additionally set

$$\Big[\!\!\Big[[\mathsf{lfp}\, R\mathbf{x}.\, \vartheta](\mathbf{y})\Big]\!\!\Big]_\ell^\alpha = \big(\mathsf{lfp}\, F_\ell^\vartheta\big)(\alpha(\mathbf{y}))$$

$$\Big[\!\!\Big[[\mathsf{gfp}\, R\mathbf{x}.\, \vartheta](\mathbf{y})\Big]\!\!\Big]_\ell^\alpha = \big(\mathsf{gfp}\, F_\ell^\vartheta\big)(\alpha(\mathbf{y}))$$

where $\mathsf{lfp}\, F_\ell^\vartheta$, $\mathsf{gfp}\, F_\ell^\vartheta$ denote the least and greatest fixed points of the update operator.

Recall that definition 2.6 defined the interpretation of negation in terms of negation normal form, i.e., $[\![\neg\varphi]\!]_\ell = [\![\mathsf{nnf}(\varphi)]\!]_\ell$. In standard semantics, we use the following observation to obtain $\mathsf{nnf}(\varphi)$ which is based on the duality of least and greatest fixed points:

$$[\mathsf{lfp}\, R\mathbf{x}.\, \vartheta](\mathbf{y}) \equiv \neg[\mathsf{gfp}\, R\mathbf{x}.\, \overline{\vartheta}](\mathbf{y}) \quad \text{where } \overline{\vartheta} \text{ is the formula}^{3}\ \neg\vartheta[R/\neg R]$$

In order to justify the interpretation of negation, we have to consider how this duality translates to (suitable) semirings, similar to our justification for FO in proposition 2.11.

[3] We write $\vartheta[R/\neg R]$ to replace every positive literal of the form $R\mathbf{x}$ in ϑ by its negation $\neg R\mathbf{x}$.

For a first idea of the requirements on a suitable semiring, let us reconsider the reachability example, this time over the semiring $\mathbb{N}$ and two different graphs. As in the previous example, we use the fixed-point iteration, now defined in terms of valuations: $\bot$, $F^\vartheta(\bot)$, $F^\vartheta(F^\vartheta(\bot))$, ... where $\bot$ is the valuation $a \mapsto \bot$ that maps every element to the least semiring element (which is 0 in $\mathbb{N}$). We again rely on the semiring for this iteration to be well-defined and to yield the least fixed point.

Example 3.5. Recall example 3.1 where we defined

$$\varphi_{\text{path}}(u, v) = [\mathsf{lfp}\, R\, x.\ x = v \vee \exists y(Exy \wedge R\, y)](u)$$

Let us first consider an $\mathbb{N}$-interpretation corresponding to the graph shown on the right, so we set $\ell(Exy) = 1$ if $(x, y) \in E$ and $\ell(Exy) = 0$ otherwise. We want to compute $[\![\varphi_{\text{path}}(u, v)]\!]_\ell$. This leads to the following fixed-point iteration, where we write a valuation π as tuple $(\pi(u), \pi(w_1), \pi(w_2), \pi(v), \pi(z))$.

- $\pi_0 = (0, 0, 0, 0, 0)$
- $\pi_1 = (0, 0, 0, 1, 0)$
- $\pi_2 = (0, 1, 1, 1, 0)$
- $\pi_3 = (2, 1, 2, 1, 0)$
- $\pi_4 = (3, 1, 2, 1, 0)$
- $\pi_5 = \pi_4$

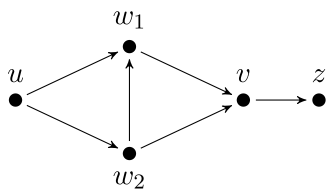

As with standard semantics, the iteration forms an ascending chain (we compare valuations pointwise) and we reach the least fixed point after four steps. The overall result is $[\![\varphi_{\text{path}}(u, v)]\!]_\ell = \pi_4(u) = 3$. We observe that π_4 maps a node to a positive value if, and only if, there is a path from the node to v. Moreover, the value equals the number of different paths to v. From u, we indeed have three paths uw_1v, uw_2v and uw_2w_1v. So in this case, the evaluation in $\mathbb{N}$ is possible and provides useful information.

However, $\mathbb{N}$ is not a suitable semiring. To see this, consider a graph with cycles as shown on the right. Writing valuations π as $(\pi(u), \pi(w), \pi(v))$, we get the following iteration:

- $\pi_0 = (0, 0, 0)$
- $\pi_1 = (0, 0, 1)$
- $\pi_2 = (0, 1, 1)$
- $\pi_3 = (1, 2, 1)$
- $\pi_4 = (2, 3, 1)$
- $\pi_n = (n - 2, n - 1, 1)$

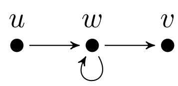

The iteration is again an ascending chain but, unlike the first example, it does not terminate after finitely many steps. Note that this is in line with our intuition, as the cycle admits an infinite number of paths from u to v.

As in standard semantics, we continue the iteration with π_ω, which is the supremum of the previous steps. In our case, $\pi_\omega(u)$ would be the supremum $\bigsqcup\{n - 2 \mid 2 \leq n < \omega\}$ which does not exist in $\mathbb{N}$, so we cannot interpret φ_{path} in this case.

So what does it mean for a semiring to be suitable? First, it must be equipped with an order such that we can talk about *least* and *greatest* fixed points. We further want the update operator F_ℓ^ϑ to be monotone as in standard semantics. For an FO formula ϑ, F_ℓ^ϑ is defined by means of the semiring operations $+$ and $\cdot$, so we want these operations to be monotone. Most importantly, we require the existence of (least and greatest) fixed points of monotone operators. As we see in the example, this is related to the existence of suprema of ascending chains and, for greatest fixed points, infima of descending chains.

The reachability example only uses least fixed points and is therefore contained in the positive fragment of LFP (which we consider in section 5.5). There have already been discussions of semiring provenance for similar logics such as datalog [GKT07] and positive fragments of CTL [Mrk18] or LFP [GT17a] which propose ω-continuous semirings. We introduce this notion in the following chapter; for now, think of an ordered semiring in which suprema of ω-chains $x_0 \leq x_1 \leq x_2 \leq \dots$ always exist and satisfy some additional properties. The ω-completion of $\mathbb{N}$ is the semiring $\mathbb{N}^\infty$ which adds an element ∞ with the usual semantics. The example then yields $\pi_\omega(u) = \infty$ as intended.

In this work, the focus is on full LFP including greatest fixed points. While we will see that $\mathbb{N}^\infty$ is also suitable for this task, we cannot work with ω-continuous semirings in general. Let us formulate the following central questions to better understand the requirements of and desired properties induced by suitable semirings.

(Q1) How to define an order on semirings such that F_ℓ^ϑ is always monotone?

(Q2) Which semirings guarantee the existence of least and greatest fixed points?

(Q3) Are semiring semantics for LFP compatible with standard semantics, as for FO?

(Q4) Does the fundamental property apply to LFP, as for FO?

(Q5) Are certain logical fragments easier to interpret (e.g., no alternation of fixed points)?

(Q6) Do semirings preserve the duality of least and greatest fixed points?

Question (Q5) is already partially answered above, as ω-continuous semirings suffice if we omit gfp-formulae. The common solution to (Q1), which we adopt, is to consider naturally ordered semirings in which the order is induced by addition (e.g. [GKT07]). For (Q2), note that the domain of standard semantics is the powerset lattice $\mathcal{P}(A^k)$. This is a complete lattice, which guarantees the existence of both kinds of fixed points. We may therefore consider semirings with an order that is a complete lattice. Indeed, [Mrk18] defines a notion of *absorptive lattice semirings* for provenance analysis of CTL. In these semirings, the order must be a complete lattice and several additional assumptions must hold. We present a more general definition of suitable semirings based on chain-completeness and discuss the properties of the resulting semantics for LFP in the following chapters to answer the questions stated above.

4 Semirings and Fixed Points

This chapter augments the algebraic structure of semirings with order-theoretic concepts to find semirings suitable for the interpretation of fixed-point logic. We focus on naturally ordered semirings whose order is chain-complete or a complete lattice in order to guarantee the existence of fixed points. Following a comprehensive list of examples, we discuss how these concepts and the algebraic properties of semirings relate to each other and to similar notions in the literature.

4.1 Order Theory

Let us start with some background on order theory and fixed points. For more details, we refer to [DP02]. Given a partially ordered set (poset) $(P, \leq)$, a *chain* is a (possibly empty) totally ordered subset of P. An ascending (or descending) *ω-chain* is a sequence $(x_n)_{n<\omega}$ with $x_n \leq x_{n+1}$ (or $x_n \geq x_{n+1}$) for all $n < \omega$, where ω is the smallest infinite ordinal. Unless stated otherwise, ω-chain usually refers to ascending ω-chains. Regarding further notation, we use $\bigsqcup S$ and $\bigsqcap S$ to denote the *supremum* and *infimum* of a set S and $\bot$, $\top$ always refer to the *least* and *greatest* elements. Let us start by introducing the notion of chain-completeness which is crucial for the existence of fixed points.

Definition 4.1. A poset $(P, \leq)$ is *chain-sup-complete* (*chain-inf-complete*) if every chain $C \subseteq P$ has a supremum $\bigsqcup C \in P$ (an infimum $\bigsqcap C \in P$). In particular, P must have a least element $\bot$ (greatest element $\top$) as supremum (infimum) of the empty chain.

If it is both chain-sup- and chain-inf-complete, we say that $(P, \leq)$ is *fully chain-complete*. We abbreviate fully chain-complete partial orders by *cpo* (and sup-cpo, inf-cpo).

Remark: In the literature, chain-completeness usually only requires suprema of chains (e.g., [DP02]) and thus coincides with what we call chain-sup-completeness. For the interpretation of logic, we need the existence of both suprema and infima to guarantee the existence of both least and greatest fixed points. Contrary to the literature, we thus use the abbreviation cpo to mean fully chain-complete partial orders (instead of requiring only suprema of chains). This deviation continues for continuous functions which, in our definition, must preserve both suprema and infima.

In some cases, considering ω-chains is sufficient for our applications and we use the following notion. Note that ω-completeness is weaker in two aspects: It limits the cardinality of chains and it does not consider infima.

Definition 4.2. A poset $(P, \leq)$ is *ω-complete* if it has a least element $\bot$ and each ascending ω-chain has a supremum in P.

If P additionally has a greatest element $\top$ and each descending ω-chain has an infimum in P, we say that P is *fully ω-complete* (we rarely consider such orders).

An alternative to complete partial orders are complete lattices which contain suprema and infima not only for chains, but for arbitrary sets.

Definition 4.3. A poset $(P, \leq)$ is a *lattice* if every two elements $a, b \in P$ have a supremum $a \sqcup b$ and an infimum $a \sqcap b$ in P. A lattice $(P, \leq)$ is *complete* if every subset $A \subseteq P$ has a supremum $\bigsqcup A$ and an infimum $\bigsqcap A$ in P.

We now turn our attention to fixed points and consider functions $f : P \to P$ on a poset P. A *fixed point* of f is an element $x \in P$ with $f(x) = x$ and we denote the *least* and *greatest* fixed point of f (if they exist) by $\mathsf{lfp}(f)$ and $\mathsf{gfp}(f)$, respectively. The central properties related to fixed points are monotonicity and continuity.

Definition 4.4. Let $f : P \to Q$ be a function on posets $(P, \leq)$ and $(Q, \preceq)$. Then f is *monotone* if $a \leq b$ implies $f(a) \preceq f(b)$ for all $a, b \in P$.

Definition 4.5. If P is a cpo, then f is *continuous* if it respects suprema and infima of nonempty chains in the following way (note that this implies the existence of the suprema/infima in Q):

(1) $f(\bigsqcup C) = \bigsqcup f(C)$ for every nonempty chain $C \subseteq P$

(2) $f(\bigsqcap C) = \bigsqcap f(C)$ for every nonempty chain $C \subseteq P$

If f only respects suprema (infima), we say that f is *sup-continuous* (*inf-continuous*).

If P is ω-complete, then f is *ω-continuous* if it respects suprema of ω-chains, i.e., $f(\bigsqcup_{n<\omega} x_n) = \bigsqcup_{n<\omega} f(x_n)$ for ascending ω-chains $(x_n)_{n<\omega}$. If P is fully ω-complete, then f is *fully ω-continuous* if it additionally respects infima of descending ω-chains.

In the above definition, we use the notation $f(C) = \{f(c) \mid c \in C\}$ to express continuity in a convenient way. As a first observation, note that every continuous or ω-continuous function f is also monotone (by considering the chain $\{a, b\}$ of two elements $a \leq b$). If P is a cpo and f is continuous, the following observation attributed to Kleene implies the existence of $\mathsf{lfp}(f)$ and $\mathsf{gfp}(f)$ (cf. [DP02, CPO Fixpoint Theorem I]).

Theorem 4.6 (Kleene). *If $(P, \leq)$ is ω-complete and $f : P \to P$ is ω-continuous, then*

$$\mathsf{lfp}(f) = \bigsqcup\{f^n(\bot) \mid n < \omega\}$$

Proof sketch. Consider the *Kleene iteration*

$$\bot, \quad f(\bot), \quad f(f(\bot)), \quad \ldots \quad f^n(\bot), \quad \ldots$$

As f is monotone, this iteration defines an ascending ω-chain which has a supremum in P. This supremum is a fixed point by ω-continuity of f:

$$f\Big(\bigsqcup\{f^n(\bot) \mid n < \omega\}\Big) = \bigsqcup f(\{f^n(\bot) \mid n < \omega\}) = \bigsqcup\{f^n(\bot) \mid n < \omega\}$$

To see that this is the *least* fixed point, one can show inductively using the monotonicity of f that $f^n(\bot) \leq x$ for every fixed point x and thus $\bigsqcup\{f^n(\bot) \mid n < \omega\} \leq x$. □

Corollary 4.7. *If $(P, \leq)$ is a cpo and $f : P \to P$ is continuous, then*

$$\mathsf{lfp}(f) = \bigsqcup\{f^n(\bot) \mid n < \omega\} \quad \textit{and} \quad \mathsf{gfp}(f) = \bigsqcap\{f^n(\top) \mid n < \omega\}$$

Proof. Every cpo is also ω-complete and every continuous function is also ω-continuous, so this follows from theorem 4.6 (for $\mathsf{gfp}(f)$, we apply the theorem to the dual order). □

Theorem 4.6 also shows that suprema and infima of ω-chains suffice to express least and greatest fixed points of continuous functions (despite the cardinality of P). If f is not continuous but only monotone, this no longer holds in general, but the existence of fixed points is still guaranteed. A proof for least fixed points in sup-cpos (which can be dualized for greatest fixed points) can be found in [DP02, CPO Fixpoint Theorem II]. We present an alternative proof that extends Kleene iteration using ordinals.

Theorem 4.8. *Let $(P, \leq)$ be a cpo and $f : P \to P$ be monotone. Then $\mathsf{lfp}(f)$ and $\mathsf{gfp}(f)$ exist.*

Proof. We show the existence of $\mathsf{lfp}(f)$, the proof for $\mathsf{gfp}(f)$ is analogous (using $\top$ and infima). We extend Kleene iteration by transfinite recursion as follows:

- $x_0 := \bot$
- $x_{\beta+1} := f(x_\beta)$ for all ordinals $\beta \in \mathsf{On}$
- $x_\gamma := \bigsqcup\{f(x_\beta) \mid \beta < \gamma\}$ for limit ordinals $\gamma \in \mathsf{On}$

Using monotonicity, one can show by induction that this defines an ascending chain and hence the suprema exist due to chain-completeness of P. As P is a set, the chain must stagnate, so there must be an ordinal $\alpha \in \mathsf{On}$ with $x_{\alpha+1} = x_\alpha$ and hence $f(x_\alpha) = x_\alpha$.

To prove that x_α is indeed the least fixed point, one shows by induction on $\beta \in \mathsf{On}$ that $x_\beta \leq y$ for any fixed point y, again by using that f is monotone. □

This theorem also applies to complete lattices. In addition, complete lattices admit a characterization of these fixed points as a consequence of the well-known Knaster-Tarski theorem (cf. [DP02, chapter 2]).

Theorem 4.9 (Knaster-Tarski)**.** *Let $(P, \leq)$ be a complete lattice and $f : P \to P$ be monotone. Then $\mathsf{lfp}(f)$ and $\mathsf{gfp}(f)$ exist and can be described as follows:*

$$\mathsf{lfp}(f) = \bigsqcap \{x \in P \mid x \geq f(x)\}$$
$$\mathsf{gfp}(f) = \bigsqcup \{x \in P \mid x \leq f(x)\}$$

These results motivate a notion of suitable semirings whose order forms a cpo or a complete lattice. To justify this requirement, we state the following converse from [DP02] which is based on results of [Mar76] (which we adapt from sup-cpos to cpos).

Theorem 4.10. *If P is a poset (lattice) in which every monotone function $f : P \to P$ has a least and greatest fixed point, then P is a cpo (complete lattice).*

4.2 Suitable Semirings

To define an order on a given semiring S, we use the *natural order* induced by addition. That is, $a \leq a+b$ for all elements $a, b \in S$. This relation is always reflexive and transitive, but not always antisymmetric. One example are rings in which $a + (b - a) = b$ for any elements a, b. We want the semiring to form a poset, so we cannot use rings such as $\mathbb{Z}$ or $\mathbb{R}$. Instead, we use the semirings $\mathbb{N}$ and $\mathbb{R}_+$ (the nonnegative real numbers) without negative values. To justify our choice, we note that the natural order coincides with the standard order on $\mathbb{N}$ and $\mathbb{R}_+$.

Definition 4.11. A semiring $(S, +, \cdot, 0, 1)$ is *naturally ordered* if the relation

$$a \leq b \quad \iff \quad \text{there is } c \in S \text{ with } a + c = b$$

is a partial order on S. We call $\leq$ the *natural order* on S.

The reason to define the order in terms of addition is that the semiring operations are then always monotone. This monotonicity lifts to update operators F^ϑ, addressing one of our requirements on suitable semirings.

Proposition 4.12. *Let S be a naturally ordered semiring. Then addition and multiplication are monotone in each argument.*

Proof. Let $a_1, a_2, b \in S$ with $a_1 \leq a_2$. By definition of the natural order, we have $a_1 + a' = a_2$ for some $a' \in S$. Using the semiring axioms, we obtain

$$\begin{aligned} a_1 + b &\leq (a_1 + b) + a' = (a_1 + a') + b = a_2 + b \\ a_1 \cdot b &\leq a_1 b + a' b = (a_1 + a') \cdot b = a_2 \cdot b \end{aligned}$$

□

We have seen above that we need chain-completeness or a complete lattice to guarantee the existence of fixed points. The following definition is therefore essential for a suitable semiring. The most important notion for the interpretation of full LFP is the one of cpo semirings which, as we show in the next chapter, is already sufficient to obtain well-defined semiring semantics.

Definition 4.13. Let S be a naturally ordered commutative semiring. We say that S is a *cpo semiring*, *lattice semiring* or *(fully) ω-complete semiring* if the natural order is a cpo, complete lattice or (fully) ω-complete, respectively.

Many natural examples of semirings have the additional property that both operations preserve suprema and infima of chains. Following the common notion of ω-continuous semirings (e.g., [DK09, GKT07]), we propose an analogous definition for cpo semirings.

Definition 4.14. A *(fully) ω-continuous semiring* is a (fully) ω-complete semiring in which addition and multiplication are (fully) ω-continuous in each argument. A *continous semiring* is a cpo semiring in which both operations are continuous in each argument.

That is, a cpo semiring S is continuous if for all $a \in S$, chains $\emptyset \neq C \subseteq S$ and $\circ \in \{+, \cdot\}$,

$$a \circ \bigsqcup C = \bigsqcup (a \circ C), \quad a \circ \bigsqcap C = \bigsqcap (a \circ C)$$

where we write $a \circ C$ for the set $\{a \circ c \mid c \in C\}$ to simplify notation.

The last step towards suitable semirings is to transfer the idea of *positive* semirings into the context of cpo semirings. We again want to ensure that computations with positive values remain positive, which leads to the following definition. As with positive semirings, this is not a hard requirement for a suitable semiring, but is needed for compatibility with standard semantics (if we interpret positive values as truth).

Definition 4.15. Let S be a cpo semiring. Then S is *chain-positive* if for all chains $C \subseteq S$ with $0 \notin C$, we have $\bigsqcap C \neq 0$.

We discuss how continuous semirings, cpo semirings and lattice semirings relate to each other in the remainder of this chapter. It is also interesting how the properties of the order relate to properties of the semiring such as idempotence or absorption. Before going into details, we consider some examples of suitable semirings.

4.3 Examples

Let us give a few examples which are useful for provenance analysis. Most of the semirings and their applications are taken from [GKT07, GT17a] and not all are suitable for fixed-point logics, but will serve as a starting point. All of the semirings presented below are commutative, positive and naturally ordered. We start with a list of semirings for different applications in provenance analysis.

- The Boolean semiring $\mathbb{B} = (\mathbb{B}, \vee, \wedge, \bot, \top)$ where $\mathbb{B} = \{\bot, \top\}$ is the simplest semiring we consider. It represents truth values and interpretations using this semiring correspond to standard semantics of logics. $\mathbb{B}$ is absorptive and both operations are idempotent, so it is a distributive lattice with join $\vee$ and meet $\wedge$. Due to being finite, it is also continuous and chain-positive.
- The natural numbers form the semiring $\mathbb{N} = (\mathbb{N}, +, \cdot, 0, 1)$ which can be used for counting proofs or for bag semantics in databases. $\mathbb{N}$ is not idempotent and lacks a greatest element. If we consider $\mathbb{N}^\infty = \mathbb{N} \cup \{\infty\}$ with the usual semantics[4] of ∞ instead, we have a greatest element and thus obtain a continuous lattice semiring. It is also chain-positive, as the infimum of a set is simply its minimum.
- $\mathbb{T} = (\mathbb{R}_+^\infty, \min, +, \infty, 0)$ and $(\mathbb{N}^\infty, \min, +, \infty, 0)$ are the *tropical* semirings which appear in min-cost computations and tropical geometry. Both are absorptive and the natural order is the inverse of the standard order on $\mathbb{R}$ or $\mathbb{N}$, so ∞ is the least element. Due to the underlying domains, both are continuous lattice semirings, but not chain-positive as we have descending chains in $\mathbb{R}$ and $\mathbb{N}$ with infimum ∞.
- The *Viterbi* semiring $\mathbb{V} = ([0, 1], \max, \cdot, 0, 1)$ is isomorphic to $\mathbb{T}$ by means of $x \mapsto e^{-x}$ and thus shares its properties. For interpretations of logic, we think of $\mathbb{V}$ as *confidence scores* (which are different from probabilities).
- The *fuzzy* semiring $([0, 1], \max, \min, 0, 1)$ and, for any nonempty set A, the powerset semiring $(\mathcal{P}(A), \cup, \cap, \emptyset, A)$ are two examples of continuous lattice semirings where addition and multiplication correspond to the lattice operations, since both semirings

[4] That is, $n + \infty = n \cdot \infty = \infty$ for $n \in \mathbb{N}^\infty \setminus \{0\}$ and $0 \cdot \infty = 0$.

are absorptive with idempotent operations. Only the powerset semiring is chain-positive, as the fuzzy semiring has descending chains approaching 0.

- A finite example is the semiring $\mathbb{A} = (\{\mathsf{P} < \mathsf{C} < \mathsf{S} < \mathsf{T} < 0\}, \min, \max, 0, \mathsf{P})$ representing the access levels public, confidential, secret, top-secret and inaccessible. This is an absorptive, chain-positive and continuous lattice semiring (note that the natural order is the opposite of the order used in the definition).
- The Łukasiewicz semiring $\mathbb{L} = ([0,1], \max, \circ, 0, 1)$ with $a \circ b = \max(0, a + b - 1)$ which occurs in many-valued logic (cf. [DK09]) is absorptive and a continuous lattice semiring. Its multiplication differs from the other examples (which only use min, max or standard operations) which makes it interesting for examples.

In many cases, it is desirable to perform a general provenance analysis which one can then specialize to different application semirings via homomorphisms. The semirings we use are mostly based on polynomials and should be universal for (that is, specialize to) a large class of semirings. We refer to them as *provenance semirings* and consider the most relevant ones in detail in chapter 6. For now, we note the following examples.

- The most general example is the polynomial semiring $\mathbb{N}[X]$ which we understand as multivariate polynomials with natural coefficients over a finite set X of variables (and standard polynomial addition and multiplication). $\mathbb{N}[X]$ is the (commutative) semiring freely generated by the set X. An idempotent alternative is the semiring $\mathbb{B}[X]$ with boolean coefficients (which we interpret as coefficients 0 and 1). Neither of these is a cpo semiring, as the chain x, $x + x^2$, $x + x^2 + x^3$, $\dots$ has no supremum. For LFP, we instead need formal power series which we present in section 6.2.
- Another universal semiring is $(\mathsf{PosBool}(X), \vee, \wedge, \bot, \top)$ which consists of all positive boolean formulae in finitely many variables X (these are combinations of variables, $\vee$ and $\wedge$), where we identify logically equivalent expressions. This is an absorptive and continuous lattice semiring, so it is suitable for LFP. Its operations $\vee$ and $\wedge$ correspond to the lattice operations $\sqcup$ and $\sqcap$ of its natural order and $\mathsf{PosBool}(X)$ is the distributive lattice freely generated by the set X. We later present an alternative definition in terms of polynomials.
- Several semirings based on polynomials have been proposed in database theory to capture different forms of provenance analyses, for example $\mathsf{Trio}(X)$, $\mathsf{Why}(X)$, $\mathsf{Sorp}(X)$ and $\mathsf{Lin}(X)$ as defined in [Gre11, DMRT14]. We consider $\mathsf{Why}(X)$ and a generalized version $\mathbb{S}^\infty[X]$ of $\mathsf{Sorp}(X)$ that has also been applied to fixed-point logics [GT19, Mrk18] in sections 6.3 and 6.4.

A first observation is that all of the examples which are cpo semirings are also continuous. Although this is not always the case, as witnessed by examples 4.22 and 4.28 below, we can thus focus primarily on continuous semirings for applications. We also want to emphasize that most of the examples are absorptive semirings.

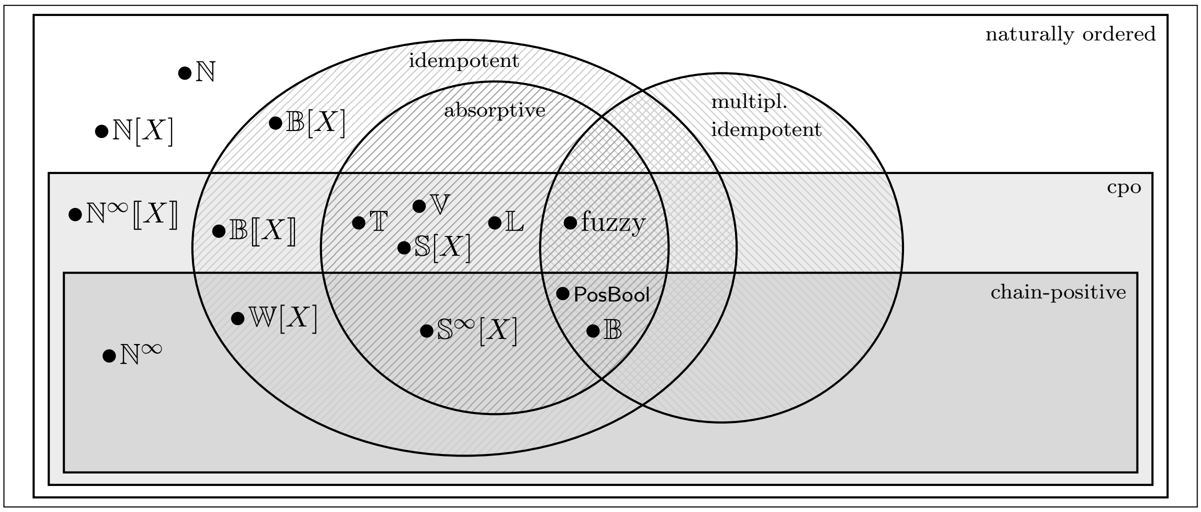

Figure 1: *Interesting semirings and their properties. All cpo semirings in the diagram are also continuous. The semirings not introduced so far are defined in chapter 6.*

4.4 Semiring Properties

For a better classification, we present few general observations about semirings that are useful later on and show how some of the notions relate to each other. An important motivation to study subclasses such as absorptive semirings is *symmetry*: Order theory often allows to dualize statements and is thus very symmetric, whereas semirings are algebraic structures and therefore asymmetric. For example, 0 is the least element and annihilating, but the dual property for 1 does not hold in general. Additionally, multiplication distributes over addition but not the other way around. Requiring idempotence or absorption increases symmetry and makes the respective semirings easier to handle for provenance analysis involving fixed points.

Proposition 4.16. *Let $(S, +, \cdot, 0, 1)$ be a naturally ordered semiring. Then*

(1) *S has the least element $\perp = 0$,*

(2) *$a, b \leq a + b$ (for all $a, b \in S$).*

(3) *S is $+$-positive,*

Proof. For (1), note that $0 + a = a$ and hence $0 \leq a$ for all $a \in S$. Regarding (2), we have $a \leq a + b$ and $b \leq a + b$ by natural order, so $a + b$ is an upper bound for $\{a, b\}$. Then (3) follows from (1) and (2): If $a + b = 0$, then $a, b \leq 0$ and hence $a = 0$ and $b = 0$. □

Proposition 4.17. *A semiring S is idempotent if, and only if, $a \sqcup b = a + b$ for all $a, b \in S$. This implies that every idempotent semiring is naturally ordered.*

Proof. We have shown above that $a+b$ is an upper bound for $\{a, b\}$. Let S be idempotent and let c be any upper bound for $\{a, b\}$, so $a \leq c$ and $b \leq c$. Using the monotonicity of addition, we obtain $a + b \leq c + b \leq c + c = c$. Thus, $a + b$ is the least upper bound and $a \sqcup b = a + b$. For the converse, note that $a + a = a \sqcup a = a$ for all $a \in S$.

To see that every idempotent semiring is naturally ordered, observe that $a \leq b$ implies $a \sqcup b = b$, while $b \leq a$ implies $a \sqcup b = a$, so together $a = b$ (for all $a, b \in S$). $\square$

This observation has an interesting consequence for idempotent cpo semirings. As $a \sqcup b$ always exists, we have suprema of arbitrary finite sets. This allows us to apply a result of [Mar76] that every sup-cpo in which suprema of finite sets exist is a complete lattice.

Corollary 4.18. *Every idempotent cpo semiring is a lattice semiring.*

The following (artificial) example shows that this result does not hold for arbitrary cpo semirings and thus clarifies the relation between cpo semirings and lattice semirings. The situation in the example is in some sense the only possibility for such a counterexample. If the order of a cpo semiring is a lattice, then theorem 4.10 implies that it is a complete lattice. As $\perp$ and $\top$ always provide lower and upper bounds, a counterexample must contain elements a, b without *least* upper bound $a \sqcup b$ (or *greatest* lower bound $a \sqcap b$). This also explains why all of the previous examples are in fact complete lattices.

Example 4.19. The diagram below depicts a canonical example of a cpo which is not a complete lattice (the set $\{1, 2\}$ has the incomparable upper bounds 3, 4 and thus no supremum). A semiring with this order would then be a non-lattice cpo semiring.

$+$	1	2	3	4	5
1	4	3	5	5	5
2		4	5	5	5
3	commutative		5	5	5
4				5	5
5					5

$\cdot$	1	2	3	4	5
1	1	2	3	4	5
2		1	3	4	5
3	commutative		5	5	5
4				5	5
5					5

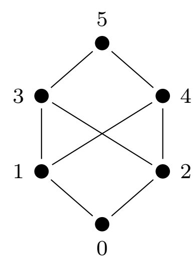

It is not immediately obvious that this example indeed occurs as natural order of a semiring. While addition can easily be defined to yield this order, we must further define a compatible (that is, distributive) multiplication. The tables above show one possible solution for both operations (which was found by smt-solving).

Another question concerns the relation to continuous semirings. A partial answer is given in [Mrk18] which shows that addition is continuous in certain idempotent semirings (the proof given there assumes absorptivity but also works in idempotent semirings).

Definition 4.20. A complete lattice $(P, \leq)$ is *completely distributive* [Ran52] if for every family $(P_i)_{i \in I}$ of subsets $P_i \subseteq P$ indexed by a set I,

$$\bigsqcap \{ \bigsqcup P_i \mid i \in I \} = \bigsqcup \{ \bigsqcap \{ f(i) \mid i \in I \} \mid f \in \mathcal{F} \}$$

where $\mathcal{F}$ is the set of choice functions, i.e., $f : I \to P$ with $f(i) \in P_i$ for all $i \in I$.

Proposition 4.21. *In a completely distributive and idempotent lattice semiring, addition is continuous.*

While most of the semirings of interest to us are continuous, one can also construct non-continuous cpo semirings, as the following example illustrates.

Example 4.22. Let $S = \mathbb{N} \cup \{\omega, \infty\}$ with the order $n < \omega < \infty$ (for all $n \in \mathbb{N}$). Then $(S, \max, \cdot, 0, 1)$ is a semiring if we extend standard multiplication in $\mathbb{N}$ by

$$n \cdot \omega = \omega, \quad \omega \cdot \omega = \infty, \quad a \cdot \infty = \infty, \quad \text{for } n \in \mathbb{N} \setminus \{0\} \text{ and } a \in S \setminus \{0\}$$

One can check by case distinction (for the additional elements ω, ∞) that S is indeed a semiring. By using maximum as addition, S is further idempotent and the natural order coincides with the order $n < \omega < \infty$ used for the definition. This order is well-founded and all subsets without maximal element have the supremum ω, so S is a lattice semiring.

Multiplication in S is *not continuous*, as witnessed by the chain $\mathbb{N}$:

$$\omega \cdot \bigsqcup \mathbb{N} = \omega \cdot \omega = \infty \neq \omega = \bigsqcup \{0, \omega\} = \bigsqcup \omega \cdot \mathbb{N}$$

Note that S is a chain itself and it is thus easy to see that S is completely distributive. Hence S provides a negative analogue of proposition 4.21 regarding multiplication.

A similar example shows that proposition 4.21 requires idempotence. Consider $(S, +, \cdot, 0, 1)$ with $n + \omega = \omega$, $\omega + \omega = \infty$, $a + \infty = \infty$ for $n \in \mathbb{N}$, $a \in S$ and, for multiplication, $a \cdot \omega = a \cdot \infty = \infty$ for $a \in S$ with $a \geq 2$. Following the argument above, we obtain a completely distributive semiring (with the natural order $n < \omega < \infty$) which is not idempotent and in which addition is not continuous ($\omega + \bigsqcup \mathbb{N} = \infty$ while $\bigsqcup \omega + \mathbb{N} = \omega$).

Based on ideas of the proof in [Mar76] we used for corollary 4.18, we can further show that in idempotent semirings, sup-continuity (which is defined on chains) is equivalent to sup-continuity on arbitrary sets.

Theorem 4.23. *Let S and T be idempotent semirings. Let further $f : S \to T$ be a sup-continuous function that preserves addition. Then f preserves suprema of arbitrary sets, so for every $A \subseteq S$:*

$$f(\bigsqcup A) = \bigsqcup f(A)$$

Proof. First note that S and T are lattice semirings by corollary 4.18, so both suprema exist. we first consider finite sets $A = \{a_1, \dots, a_n\}$ for which the statement follows from the additivity of f and the idempotence of both semirings:

$$f(\bigsqcup A) = f(a_1 + \cdots + a_n) = f(a_1) + \cdots + f(a_n) = \bigsqcup f(A)$$

The interesting case are infinite sets A. The intuition is that we can express the supremum of a family $(a_i)_{i<\omega}$ by the chain a_0, $a_0 + a_1$, $a_0 + a_1 + a_2$, ... of partial sums. We apply the same idea to sets of arbitrary cardinality.

Towards a contradiction, let $A \subseteq S$ be an infinite set with *minimal cardinality* such that $f(\bigsqcup A) \neq \bigsqcup f(A)$. Let α be the cardinality of A (that is, the smallest ordinal in bijection with A). Using the bijection between A and α, we can write $A = \{a_\beta \mid \beta < \alpha\}$.

Consider the chain $C = (c_\beta)_{\beta<\alpha}$ with $c_\beta = \bigsqcup\{a_\delta \mid \delta \leq \beta\}$ (in analogy to partial sums). Then C is clearly a chain with $c_\beta \leq \bigsqcup A$ for all $\beta < \alpha$, hence $\bigsqcup C \leq \bigsqcup A$. The other direction holds as well, so C has the same supremum as A:

$$\bigsqcup C = \bigsqcup \Big\{ \bigsqcup\{a_\delta \mid \delta \leq \beta\} \mid \beta < \alpha \Big\} \geq \bigsqcup\{a_\beta \mid \beta < \alpha\} = \bigsqcup A$$

The same holds for the suprema of $f(C)$ and $f(A)$. The crucial step is:

$$\bigsqcup f(C) = \bigsqcup \Big\{ f\big(\bigsqcup\{a_\delta \mid \delta \leq \beta\}\big) \mid \beta < \alpha \Big\} = \bigsqcup \Big\{ \bigsqcup\{f(a_\delta) \mid \delta \leq \beta\} \mid \beta < \alpha \Big\} \qquad (*)$$

We claim that the set $A' = \{a_\delta \mid \delta \leq \beta\}$ has a smaller cardinality than A. Thus A' is either finite (which we have already considered above) or we can exploit the minimality assumption on A to conclude that $f(\bigsqcup A') = \bigsqcup f(A')$. To see that the cardinality is smaller, note that A' is indexed (and thus in bijection) with the ordinal $\beta + 1$. As α is infinite, it must be a limit ordinal and hence $\beta < \alpha$ implies $\beta + 1 < \alpha$, so $\beta + 1$ (and thus A') cannot have cardinality α.

It is then easy to show that $(*)$ equals $\bigsqcup f(A)$ by considering both directions:

$$\bigsqcup f(A) = \bigsqcup \Big\{ \bigsqcup f(A) \mid \beta < \alpha \Big\} \geq (*) \geq \bigsqcup \Big\{ f(a_\beta) \mid \beta < \alpha \Big\} = \bigsqcup f(A)$$

Together with the sup-continuity of f (applied to the chain C), we obtain

$$f\big(\bigsqcup A\big) = f\big(\bigsqcup C\big) = \bigsqcup f(C) = \bigsqcup f(A)$$

which contradicts our assumption on A and thus closes the proof. $\square$

By considering addition and multiplication as functions f, we obtain the following corollary. The reason for the more general formulation of the theorem is that we can later apply it also to homomorphisms.

Corollary 4.24. *Let S be an idempotent sup-continuous semiring. Then S is sup-continuous on arbitrary sets, so for all $b \in S$, $\circ \in \{+, \cdot\}$ and all sets $A \subseteq S$,*

$$b \circ \bigsqcup A = \bigsqcup (b \circ A)$$

An important subclass of idempotent semirings are the *absorptive* ones in which $a+ab = a$ for all elements a, b. The main motivation to study absorption is that it leads to more symmetry. This simplifies reasoning about fixed points and establishes a stronger connection between the algebraic and order theoretic properties of semirings. The reason is that absorption *relates addition and multiplication* such that multiplication decreases elements (thus becoming somewhat dual to addition) and such that the neutral element 1 becomes the greatest element (dually to the least element 0). This further means that 1 is absorbing (dually to 0 being annihilating). In fact, all of these properties are equivalent:

Proposition 4.25. *In a naturally ordered semiring S, the following are equivalent:*

(1) *S is absorptive,*

(2) *S has the greatest element $\top = 1$,*

(3) *1 is the absorbing element of addition, i.e., $1 + a = 1$ for all $a \in S$,*

(4) *multiplication is decreasing, i.e., $a \cdot b \leq a, b$ for all $a, b \in S$.*

Proof. We first prove the equivalence of (1)-(3). If S is absorptive, then $1 + 1 \cdot a = 1$ for all $a \in S$ and thus $\top = 1$. If $\top = 1$, then $1 \leq 1 + a \leq 1$ and thus $1 + a = 1$ for all $a \in S$. Multiplication with $b \in S$ gives absorption.

For (4), we first note that absorption $a + ab = a$ implies $ab \leq a$ for all $a \in S$. For the converse, we consider the elements a and $1 + b$. We have $a \leq a + ab$ by natural order. If multiplication is decreasing, then also $a \geq a \cdot (1 + b) = a + ab$ and together $a = a + ab$. □

The only asymmetric properties left are distributivity and idempotence of addition (which makes addition equivalent to the join). If we require the semiring to be multiplicatively idempotent, then these are resolved as well and we end up with a distributive lattice.

Proposition 4.26. *Let S be an absorptive and multiplicatively idempotent semiring. Then $a \cdot b = a \sqcap b$ for all $a, b \in S$, so S is a distributive lattice with join $+$, meet $\cdot$, least element 0 and greatest element 1.*

Proof. We have already shown that $a \cdot b \leq a, b$ and thus $a \cdot b \leq a \sqcap b$. Now let c be any lower bound, so $c \leq a$ and $c \leq b$. Using monotonicity and idempotence of multiplication, we have $c = c \cdot c \leq a \cdot c \leq a \cdot b$, so $a \cdot b$ is the greatest lower bound of a and b. □

Absorptive semirings are particularly interesting when computing greatest fixed points for which the fixed-point iteration starts at $\top$, which then equals 1, and multiplication offers a way to decrease this value during the iteration. Indeed, many of the above examples are absorptive and we present a universal semiring for the class of all absorptive continuous semirings in section 6.4, the semiring $\mathbb{S}^\infty[X]$ of absorptive polynomials.

4.5 Related Concepts

In [Mrk18], Mrkonjić proposes the notion of *absorptive lattice semirings*[5] for provenance analysis of CTL. In our terminology, these are, unsurprisingly, absorptive lattice semirings together with additional assumptions: The natural order must be a completely distributive lattice, multiplication must be sup-continuous for arbitrary sets (not just for chains) and inf-continuous for ω-chains. Continuity of addition is not required since it is implied by proposition 4.21. The main instance they study is the semiring of absorptive polynomials which is also our primary focus in chapter 6. The application semirings $\mathbb{V}$, $\mathbb{T}$, $\mathbb{L}$ and the fuzzy semiring are further examples of *absorptive lattice semirings*, because they all use the standard order on $[0, 1]$ (or an isomorphic one) which is a complete chain and thus completely distributive (see [Ran52, section 5]).

For the general discussion, we try to keep the restrictions at a minimum and mostly work with cpo semirings or, for some of the deeper results, with continuous semirings. Although the later chapters focus more on absorptive semirings, we also allow non-absorptive semirings such as $\mathbb{N}^\infty$ or $\mathbb{W}[X]$. Regarding continuity, note that we require multiplication to preserve suprema and infima of arbitrary chains whereas [Mrk18] considers infima only over ω-chains. The reason might be that they study CTL instead of LFP. CTL can be embedded into alternation-free LFP for which fully ω-continuous semirings are indeed sufficient, as we shall see in section 5.5.

An open question posed by Mrkonjić is whether some of the additional assumptions are redundant and, in particular, whether every absorptive, naturally ordered semiring is also an *absorptive lattice semiring*. Corollary 4.18 indeed shows that every absorptive cpo semiring is also a lattice semiring. However, there are several examples which constitute a negative answer to this question. A first observation is that there are absorptive semirings which are not chain-complete (and thus no lattice semirings).

[5] We refer to Mrkonjić's definition in *italic letters* to avoid confusion with our terminology.

Example 4.27. Consider the semiring $([0,1] \setminus \{\frac{1}{42}\}, \max, \min, 0, 1)$ which modifies the fuzzy semiring by excluding a single element. As the operations are max and min, they remain well-defined and absorptive. The semiring is not chain-complete, as there are chains approaching $\frac{1}{42}$ which no longer have a least upper bound.

More interesting is the following example which features an absorptive lattice semiring which is not continuous and thus not an *absorptive lattice semiring* in the sense of [Mrk18]. It further shows that the complete distributivity is not redundant.

Example 4.28. Consider the direct semiring product $\mathbb{B} \times \mathbb{V}$ (with operations defined pointwise, e.g., $(b_1, v_1) + (b_2, v_2) = (b_1 \vee b_2, \max(v_1, v_2))$). Both $\mathbb{B}$ and $\mathbb{V}$ are completely distributive, so $\mathbb{B} \times \mathbb{V}$ is a completely distributive lattice semiring as well. It is further continuous due to the continuity of $\mathbb{B}$ and $\mathbb{V}$.

We modify this semiring to be non-continuous by setting $S = \mathbb{B} \times \mathbb{V} \setminus \{(\top, 0)\}$. Note that S is well-defined as $(\top, 0)$ does not result from addition or multiplication of other elements. S is also fully chain-complete: Descending chains in S which would have the infimum $(\top, 0)$ in $\mathbb{B} \times \mathbb{V}$ now have the infimum $(\bot, 0)$ in S; ascending chains (and even arbitrary sets) in S have the same supremum as in $\mathbb{B} \times \mathbb{V}$. This is easy to see in the following depiction of the natural order, hence S is an absorptive lattice semiring.

$$
\begin{array}{ccccccccc}
(\bot, 0) & \leq & \dots & \leq & (\bot, \frac{1}{2}) & \leq & \dots & \leq & (\bot, 1) \\
& \rotatebox{-45}{$\leq$} & \rotatebox{90}{$\leq$} & & \rotatebox{90}{$\leq$} & & \rotatebox{90}{$\leq$} & & \rotatebox{90}{$\leq$} \\
& & \dots & \leq & (\top, \frac{1}{2}) & \leq & \dots & \leq & (\top, 1)
\end{array}
$$

However, *addition is not continuous* and, due to proposition 4.21, S is thus not completely distributive, as witnessed by the chain $C = \{(\top, \frac{1}{n}) \mid n < \omega\}$:

$$
\begin{aligned}
(\bot, 1) + \bigsqcap C = (\bot, 1) + (\bot, 0) = (\bot, 1)& \\
\neq& \\
(\top, 1) = \bigsqcap \{(\top, 1)\} = \bigsqcap \big((\bot, 1) + C\big)&
\end{aligned}
$$

Note that multiplication remains continuous. We have already argued that suprema behave as in $\mathbb{B} \times \mathbb{V}$. The same holds for chains whose infimum in $\mathbb{B} \times \mathbb{V}$ is different from $(\top, 0)$. It remains to consider chains $C \subseteq S$ such that C has the infimum $(\top, 0)$ in $\mathbb{B} \times \mathbb{V}$ and thus $\bigsqcap C = (\bot, 0)$ in S. Such chains have the form $C = (\top, C')$ where $C' \subseteq \mathbb{V}$ with $\bigsqcap C' = 0$. Now consider any $(b, v) \in S$. Then

$$
\bigsqcap \big((b, v) \cdot (\top, C')\big) = \bigsqcap (b, vC') \overset{(*)}{=} (\bot, 0) = (b, v) \cdot (\bot, 0) = (b, v) \cdot \bigsqcap C
$$

For $(*)$, note that v is a constant, so the chain vC' still has the infimum 0 in $\mathbb{V}$ and thus $\bigsqcap(b, vC')$ would have the infimum $(b, 0)$ in $\mathbb{B} \times \mathbb{V}$ which is always $(\bot, 0)$ in S.

This example shows that absorptive lattice semirings are not continuous in general. Even if we require multiplication to be continuous, addition can still be non-continuous if the semiring is not completely distributive.

We can modify $\mathbb{B} \times \mathbb{V}$ in a different way such that addition remains continuous but multiplication does not. The following example illustrates the construction which is dual to the previous example, underlining our claim that absorption leads to more symmetry. Both example 4.29 and example 4.22 suggest that the continuity requirement on multiplication in [Mrk18] is not redundant, but they do not rule out the possibility that continuity could be implied by complete distributivity in absorptive semirings.

Example 4.29. Following the previous example, let $S = \mathbb{B} \times \mathbb{V} \setminus \{(\bot, 1)\}$. Note that, due to symmetry, the natural order on S is dual to the above example. In particular, S is a lattice semiring but not completely distributive.

$$\begin{array}{ccccccccc}
(\bot, 0) & \leq & \dots & \leq & (\bot, \frac{1}{2}) & \leq & \dots & & \\
\rotatebox{90}{$\leq$} & & \rotatebox{90}{$\leq$} & & \rotatebox{90}{$\leq$} & & \rotatebox{90}{$\leq$} & \rotatebox{135}{$\leq$} & \\
(\top, 0) & \leq & \dots & \leq & (\top, \frac{1}{2}) & \leq & \dots & \leq & (\top, 1)
\end{array}$$

Moreover, *multiplication is not continuous* for $C = \{(\bot, 1 - \frac{1}{n}) \mid n \geq 1\}$:

$$\begin{aligned}
(\top, 0) \cdot \bigsqcup C = (\top, 0) \cdot (\top, 1) = (\top, 0)& \\
\neq& \\
(\bot, 0) = \bigsqcup\{(\bot, 0)\} = \bigsqcup\big((\top, 0) \cdot C\big)&
\end{aligned}$$

Addition remains continuous which can be seen by dualizing the argument in the previous example. Together, these two examples show that neither continuity (of either operation) nor complete distributivity hold in all (absorptive) lattice semirings.

We do not consider complete distributivity in the remaining work. Instead, we require continuity of both operations explicitly by working with continuous semirings. Our results indicate that continuity is most likely the property we are actually interested in.

Another related notion worth mentioning is that of *c-semirings* introduced in [BMR97] to generalize constraint satisfaction problems. Their definition is as follows.

Definition 4.30. A *c-semiring* $(S, +, \cdot, 0, 1)$ is a commutative semiring in which addition is defined on (possibly infinite) subsets of S such that:

- $\sum(\{a\}) = a$ for all $a \in S$,
- $\sum(\emptyset) = 0$ and $\sum(S) = 1$,
- $\sum(\bigcup_{i \in I} A_i) = \sum(\{\sum(A_i) \mid i \in I\})$ for families $(A_i)_{i \in I}$ of sets (*flattening property*).

Distributivity then requires that $a \cdot \sum(A) = \sum(a \cdot A)$ for all $a \in S$ and $A \subseteq S$.

Note that the addition in c-semirings is always idempotent, as it is defined on sets. Given a c-semiring, [BMR97] defines the order $a \leq b \iff a + b = b$ which, due to idempotence, coincides with the natural order we consider. They further show that c-semirings are lattice semirings in which the supremum $\bigsqcup A$ is given by the sum $\sum(A)$. It is also shown that multiplication is decreasing, which implies absorption (by proposition 4.25). Distributivity in c-semirings implies that multiplication is sup-continuous, i.e., preserves suprema of chains (addition is trivially sup-continuous due to idempotence).

Coming from our terminology, let S be an absorptive lattice semiring and extend addition to sets by $\sum(A) = \bigsqcup A$ for $A \subseteq S$ (note that this coincides with the addition in S on finite sets). It follows from the properties of suprema that the flattening property is satisfied. The only remaining difference is that we consider continuity only for chains whereas distributivity for c-semirings must hold for arbitrary sets, but this is resolved by corollary 4.24. We can thus state the following equivalence result.

Proposition 4.31. *Every c-semiring is a sup-continuous absorptive lattice semiring.*

Conversely, every sup-continuous absorptive lattice semiring S becomes a c-semiring by setting $\sum(A) = \bigsqcup A$ (which coincides with the addition in S on finite sets).

In particular, every *absorptive lattice semiring* in the sense of [Mrk18] becomes a c-semiring in this way. In the remaining work, we will see that absorptive continuous semirings are the most viable choice for provenance analysis of fixed-point logic. It is interesting to see that a similar notion (although defined in a different way) also arises in other contexts.

4.6 Homomorphisms and Function Semirings

In cpo semirings, it is natural to require that homomorphisms additionally preserve the cpo structure. Due to the definition of the natural order, some properties are always preserved by semiring homomorphisms.

Proposition 4.32. *Let $h : S \to T$ be a semiring homomorphism on cpo semirings S and T. Then*

(1) *h is monotone (preserves the order)*

(2) $h(\bot) = \bot$

(3) $h(\bigsqcup C) \geq \bigsqcup h(C)$ *and* $h(\bigsqcap C) \leq \bigsqcap h(C)$ *for every chain C*

Proof. Monotonicity follows from the additivity of h, claim (2) holds because of $\bot = 0$. Lastly, (3) follows from (1), as $\bigsqcup C \geq c$ and thus $h(\bigsqcup C) \geq h(c)$ for all $c \in C$, hence $h(\bigsqcup C) \geq \bigsqcup h(C)$ (the second statement follows analogously). $\square$

We later see examples violating $h(\top) = \top$ and the other direction in (3), which would make h continuous, and therefore introduce cpo homomorphisms.

Definition 4.33. A *cpo homomorphism* is a function $h : P \to Q$ on cpos P, Q with:

- $h(\bot) = \bot$,
- $h(\top) = \top$,
- h is continuous (preserves suprema and infima of nonempty chains)

A *cpo-semiring homomorphism* is a semiring homomorphism $h : S \to T$ on cpo semirings which is also a cpo homomorphism.

Recall the truth projection $\dagger_S : S \to \mathbb{B}$ from example 2.5 which is a semiring homomorphism if (and only if) S is positive. We can similarly describe when $\dagger_S$ is also a cpo-semiring homomorphism.

Example 4.34. Let S be a positive semiring. Then $\dagger_S$ is a cpo-semiring homomorphism if, and only if, S is chain-positive.

First note that the greatest element $\top \in S$ satisfies $\top \neq 0$. Hence $h(\top) = \top$ and we always have $h(\bot) = \bot$. One can additionally show (by case distinction on $C = \{0\}$) that $h(\bigsqcup C) = \bigsqcup h(C)$ for chains $C \subseteq S$.

If S is chain-positive, then we further see (by case distinction on $0 \in C$) that $\dagger_S$ also preserves infima and is thus continuous. On the other hand, a chain C of positive elements with $\bigsqcap C = 0$ leads to $\dagger_S(\bigsqcap C) = \bot$ but $\bigsqcap \dagger_S(C) = \top$.

This is a first indication that chain-positivity is needed for compatibility with standard truth semantics (which are equivalent to semiring semantics in $\mathbb{B}$).

Note that continuous homomorphisms have to preserve suprema and infima of nonempty chains, but not necessarily of arbitrary sets. However, as mentioned earlier, theorem 4.23 shows that suprema of sets are preserved in idempotent semirings.

Corollary 4.35. *Let $h : S \to T$ be a continuous homomorphism on idempotent semirings. Then h preserves suprema of arbitrary sets, so for all $A \subseteq S$:*

$$h(\bigsqcup A) = \bigsqcup h(A)$$

For the interpretation of logic, we work with S-valuations which are functions $f : A \to S$ for a semiring S and some set A. We denote the set of such functions by S^A and we can turn S^A into a semiring by defining pointwise operations. This simplifies notation and most of the properties transfer from S to S^A.

Proposition 4.36. *Let S be a semiring and A be a nonempty set. Then the set of functions $S^A = \{f : A \to S\}$ is a semiring with pointwise operations ($f, g \in S^A$):*

$$(f + g)(a) = f(a) + g(a)$$
$$(f \cdot g)(a) = f(a) \cdot g(a)$$

and the neutral elements $\mathbf{0} : a \mapsto 0$ and $\mathbf{1} : a \mapsto 1$.

If S is naturally ordered, then the above addition induces a natural order on S^A with $f \leq g$ iff $f(a) \leq g(a)$ for all $a \in A$. If $\bot$ and $\top$ are the least and greatest elements of S, then $\bot : a \mapsto \bot$ and $\top : a \mapsto \top$ are the least and greatest elements of S^A.

Proposition 4.37. *Let S be a semiring and A be a nonempty set. If S is naturally ordered, idempotent, absorptive or multiplicative idempotent, then so is S^A. If S is a cpo, ω-complete, lattice, continuous or chain-positive semiring then so is S^A.*

The proof of these propositions is simply a matter of reducing the statements on S^A to the properties of S. As the operations are defined pointwise, so are natural order, suprema and infima. For instance, the supremum of a chain $C \subseteq S^A$ is given by $\bigsqcup C = \big(a \mapsto \bigsqcup\{f(a) \mid f \in C\}\big)$. Note that not all properties transfer: S^A always has divisors of 0 (given that $|A| \geq 2$) even if S is positive.

5 Semiring Semantics for Fixed-Point Logic

In this chapter, we return to fixed-point logic and show that semiring semantics for LFP are well-defined for cpo semirings. The analysis of these semantics provides positive answers to the previously stated questions which indicate that semiring semantics behave reasonably when extended to fixed-point logic. At the end of the chapter, we discuss how we can relax chain-completeness for certain logical fragments of LFP.

5.1 Definition

Recall the definition of semiring semantics via S-interpretations from definition 3.4,

$$\llbracket [\mathsf{lfp}\, R\,\mathbf{x}.\, \vartheta](\mathbf{y}) \rrbracket_\ell^\alpha = \big(\mathsf{lfp}\, F_\ell^\vartheta\big)(\alpha(\mathbf{y}))$$
$$\llbracket [\mathsf{gfp}\, R\,\mathbf{x}.\, \vartheta](\mathbf{y}) \rrbracket_\ell^\alpha = \big(\mathsf{gfp}\, F_\ell^\vartheta\big)(\alpha(\mathbf{y}))$$

where F_ℓ^ϑ is the update operator induced by the formula $\vartheta(R, x)$. In chapter 3, we relied on a *suitable* semiring for the definition, which we can now make precise.

Theorem 5.1. *Semirings semantics for LFP are well-defined for cpo semirings.*

This result is based on the observation that cpo semirings guarantee least and greatest fixed points of monotone operators. Let us also highlight the converse stated in theorem 4.10 which is a strong indication that cpo semirings are the most general choice[6] for suitable semirings. We have seen that cpo semirings are closely related to lattice semirings (which are perhaps more natural as analogy to the powerset lattice used in standard semantics).

The update operators F_ℓ^ϑ operate on valuations $\pi : A^k \to S$. In order for definition 3.4 to be well-defined, we thus have to define an order on valuations. Given a cpo semiring S, we view the set $\mathsf{Val}_{S,k}$ of valuations as a function semiring. In particular, this induces a pointwise order on valuations, so $\pi_1 \leq \pi_2$ iff $\pi_1(\mathbf{a}) \leq \pi_2(\mathbf{a})$ for all $\mathbf{a} \in A^k$. This also means that suprema and infima are computed pointwise (see section 4.6). Similarly, S-interpretations ℓ form a function semiring and are compared pointwise as well.

With this in mind, we prove theorem 5.1 by showing that monotonicity of addition and multiplication lifts to update operators F_ℓ^ϑ. Like most of the proofs in this chapter, this is shown by induction on the structure of formulae and on the corresponding fixed-point iteration for which we use the following notation.

[6] In section 5.5, we see that fully ω-continuous semirings might be an alternative.

Definition 5.2. Let $\varphi = [\mathsf{lfp}\, R\,\mathbf{x}.\ \vartheta](\mathbf{y})$ be an LFP formula, S a cpo semiring, ℓ an S-interpretation and α a variable valuation. The *fixed-point iteration* is the sequence of S-valuations defined by transfinite recursion as follows.

$$\begin{aligned} \pi_0 &= \bot && \text{where } \bot : \mathbf{a} \mapsto \bot \text{ is the least element of } \mathsf{Val}_k, \\ \pi_{\beta+1} &= F_\ell^\vartheta(\pi_\beta) = [\![\vartheta]\!]^\bullet_{\ell[R/\pi_\beta]} && \text{for ordinals } \beta \in \mathsf{On}, \\ \pi_\lambda &= \bigsqcup\{\pi_\beta \mid \beta < \lambda\} && \text{for limit ordinals } \lambda \in \mathsf{On}. \end{aligned}$$

For $\varphi = [\mathsf{gfp}\, R\,\mathbf{x}.\ \vartheta](\mathbf{y})$, we instead set $\pi_0 = \top$ and $\pi_\lambda = \bigsqcap\{\pi_\beta \mid \beta < \lambda\}$. If we want to make ℓ explicit, we write $\pi_\ell^{(\beta)}$ for π_β.

If F_ℓ^ϑ is monotone, then this definition is well-defined and terminates in the least (or greatest) fixed point (see the proof of the cpo fixed-point theorem 4.8 for details). We split the proof into two parts which together show the monotonicity of F_ℓ^ϑ and $[\![\cdot]\!]_\ell$.

Lemma 5.3. *Let S be a cpo semiring, α be a valuation and $\vartheta(R, \mathbf{x})$ be a formula. If $[\![\vartheta]\!]_\ell^\alpha$ is monotone in ℓ, then the update operator F_ℓ^ϑ is monotone.*

Proof. Let k be the arity of R and let $\pi_1, \pi_2 \in \mathsf{Val}_k$ with $\pi_1 \leq \pi_2$. To simplify notation, let $\pi_1' = F_\ell^\vartheta(\pi_1)$ and $\pi_2' = F_\ell^\vartheta(\pi_2)$.

Due to $\pi_1 \leq \pi_2$, we also have $\ell[R/\pi_1] \leq \ell[R/\pi_2]$. Then $\pi_1' \leq \pi_2'$, as for all $\mathbf{a} \in A^k$:

$$\pi_1'(\mathbf{a}) = [\![\vartheta]\!]^{\alpha[\mathbf{x}/\mathbf{a}]}_{\ell[R/\pi_1]} \leq [\![\vartheta]\!]^{\alpha[\mathbf{x}/\mathbf{a}]}_{\ell[R/\pi_2]} = \pi_2'(\mathbf{a}) \qquad \square$$

Proposition 5.4. *Let S be a cpo semiring. Then $[\![\varphi]\!]_\ell^\alpha$ is monotone with respect to ℓ, so given two S-interpretations ℓ_1 and ℓ_1, the following implication holds for all LFP formulae φ and all variable valuations α:*

$$\ell_1 \leq \ell_2 \quad \Longrightarrow \quad [\![\varphi]\!]^\alpha_{\ell_1} \leq [\![\varphi]\!]^\alpha_{\ell_2}$$

Proof. It suffices to consider formulae in negation normal form. The proof is mostly a straight-forward induction with the most interesting case being lfp- and gfp-formulae.

- If φ is a literal $\varphi = R\,\mathbf{x}$, then $[\![\varphi]\!]^\alpha_{\ell_1} = \ell_1(R\,\alpha(\mathbf{x})) \leq \ell_2(R\,\alpha(\mathbf{x})) = [\![\varphi]\!]^\alpha_{\ell_2}$. The same holds for negative literals $\varphi = \neg R\,\mathbf{x}$ (and for equality atoms).
- If $\varphi = \varphi_1 \vee \varphi_2$, then $[\![\varphi]\!]^\alpha_{\ell_i} = [\![\varphi_1]\!]^\alpha_{\ell_i} + [\![\varphi_2]\!]^\alpha_{\ell_i}$ for $i \in \{1, 2\}$ and the claim follows by induction and monotonicity of $+$. The interpretation of $\wedge$ via $\cdot$ is analogous.

- If $\varphi = \exists x\ \vartheta(x)$, then $[\![\varphi]\!]^{\alpha}_{\ell_i} = \sum_{a \in A} [\![\vartheta]\!]^{\alpha[x/a]}_{\ell_i}$. The sum is finite due to the finite universe A, so the claim again follows by induction and monotonicity of $+$. The case for $\forall$ is analogous.

- If $\varphi = [\mathsf{lfp}\, R\, \mathbf{x}.\ \vartheta](\mathbf{y})$ where R has arity k, we proceed by induction on the fixed-point iterations. By the induction hypothesis and the previous lemma, $F^{\vartheta}_{\ell_1}$ and $F^{\vartheta}_{\ell_2}$ are monotone and hence the fixed-point iterations are well-defined.

  We show by transfinite induction on β that $\pi^{(\beta)}_{\ell_1} \leq \pi^{(\beta)}_{\ell_2}$.

  - $\pi^{(0)}_{\ell_1} = \bot = \pi^{(0)}_{\ell_2}$.
  - For $\beta + 1$, let us abbreviate $\ell_1' = \ell_1[R/\pi^{(\beta)}_{\ell_1}]$ and analogously $\ell_2' = \ell_2[R/\pi^{(\beta)}_{\ell_2}]$. Note that $\ell_1 \leq \ell_2$ together with the induction hypothesis for β imply that $\ell_1' \leq \ell_2'$, so we can apply the outer induction hypothesis for ϑ and obtain:

  $$\pi^{(\beta+1)}_{\ell_1} = F^{\vartheta}_{\ell_1}(\pi^{(\beta)}_{\ell_1}) = [\![\vartheta]\!]^{\bullet}_{\ell_1'} \leq [\![\vartheta]\!]^{\bullet}_{\ell_2'} = F^{\vartheta}_{\ell_2}(\pi^{(\beta)}_{\ell_2}) = \pi^{(\beta+1)}_{\ell_2}$$

  - For limit ordinals λ, we have $\pi^{(\beta)}_{\ell_1} \leq \pi^{(\beta)}_{\ell_2}$ for $\beta < \lambda$ by induction and hence

  $$\pi^{(\lambda)}_{\ell_1} = \bigsqcup\{\pi^{(\beta)}_{\ell_1} \mid \beta < \lambda\} \leq \bigsqcup\{\pi^{(\beta)}_{\ell_2} \mid \beta < \lambda\} = \pi^{(\lambda)}_{\ell_2}$$

  This ends the inner induction. As Val_k is a set, there is a sufficiently large ordinal α with $\pi^{(\alpha+1)}_{\ell_1} = \pi^{(\alpha)}_{\ell_1}$ and $\pi^{(\alpha+1)}_{\ell_2} = \pi^{(\alpha)}_{\ell_2}$. Then $\pi^{(\alpha)}_{\ell_1}$ and $\pi^{(\alpha)}_{\ell_2}$ equal the least fixed points of $F^{\vartheta}_{\ell_1}$ and $F^{\vartheta}_{\ell_2}$, respectively, so we see that $\mathsf{lfp}\, F^{\vartheta}_{\ell_1} \leq \mathsf{lfp}\, F^{\vartheta}_{\ell_2}$. In particular,

  $$[\![\varphi]\!]^{\alpha}_{\ell_1} = \big(\mathsf{lfp}\, F^{\vartheta}_{\ell_1}\big)(\alpha(\mathbf{y})) \leq \big(\mathsf{lfp}\, F^{\vartheta}_{\ell_2}\big)(\alpha(\mathbf{y})) = [\![\varphi]\!]^{\alpha}_{\ell_2}$$

  The proof for gfp-formulae is analogous via the corresponding fixed-point iteration.

This closes the proof and shows that $[\![\varphi]\!]_{\ell}$ is monotone with respect to ℓ. □

Together with the previous lemma, we can conclude that the update operators used in the definition of $[\![\varphi]\!]_{\ell}$ are always monotone and thus least and greatest fixed points always exist. This proves theorem 5.1 and answers the question on suitable semirings.

5.2 Semiring vs. Standard Semantics

Another question concerns the compatibility with standard semantics. For FO, we have seen that we can provide canonical interpretations $\ell_{\mathfrak{A}}$ and relate model-defining interpretations in *positive* semirings to standard semantics. We present analogue statements for LFP under the additional requirement of *chain-positivity*.

A first step is to recall the canonical truth interpretation $\ell_{\mathfrak{A}} : \mathsf{Lit}_A \to \mathbb{B}$ for a given model $\mathfrak{A}$. As for FO, its semantics correspond to standard semantics and we can thus view S-interpretations $\ell : \mathsf{Lit}_A \to S$ as a generalization of standard semantics.

Proposition 5.5. *Let $\mathfrak{A}$ be a τ-structure with universe A, $\varphi(\mathbf{x})$ an LFP formula and $\mathbf{a}$ a tuple (with the same arity as $\mathbf{x}$) of elements from A. Then*

$$\mathfrak{A} \models \varphi(\mathbf{a}) \quad \Longleftrightarrow \quad [\![\varphi(\mathbf{a})]\!]_{\ell_{\mathfrak{A}}} = \top$$

Proof sketch. Straightforward induction on φ. For fixed-point formulae, one shows by induction on the fixed-point iterations that $\mathbf{a} \in R_\beta \iff \pi_\beta(\mathbf{a}) = \top$ (for all $\beta \in \mathsf{On}$). $\square$

It is possible to also consider an analogue of the counting interpretation $\ell_{\#\mathfrak{A}}$ in $\mathbb{N}^\infty$ instead of $\mathbb{N}$, but it is less clear how one would count proofs in the presence of fixed points. We therefore defer the discussion until chapter 6.

The other direction – from a model-defining S-interpretation ℓ to the induced model $\mathfrak{A}_\ell$ – is more involved. For FO, we have required that S is positive to obtain the equivalence of $\mathfrak{A}_\ell \models \varphi$ and $[\![\varphi]\!]_\ell \neq 0$. This is not sufficient for LFP.

Example 5.6. Consider the formula $\varphi_{\text{infpath}}(v) = [\mathsf{gfp}\, R\, x.\ \exists y(Exy \wedge Ry)](v)$ which expresses that, given a directed graph, there is an infinite path from the node v.

We interpret this formula in the Viterbi semiring $\mathbb{V} = ([0,1], \max, \cdot, 0, 1)$ on the graph shown on the right by setting $\ell(Evv) = \frac{1}{2}$ and $\ell(\neg Evv) = 0$. Then ℓ is a model-defining $\mathbb{V}$-interpretation and the induced model satisfies $\mathfrak{A}_\ell \models Evv$ and thus $\mathfrak{A}_\ell \models \varphi_{\text{infpath}}(v)$.

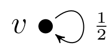

The computation of $[\![\varphi_{\text{infpath}}(v)]\!]_\ell$ in $\mathbb{V}$ yields the following fixed-point iteration:

$$\begin{aligned}
\pi_0(v) &= \top = 1 \\
\pi_1(v) &= [\![Evv]\!]_{\ell[R/\pi_0]} \cdot [\![R\,v]\!]_{\ell[R/\pi_0]} = \tfrac{1}{2} \cdot 1 = \tfrac{1}{2} \\
\pi_2(v) &= [\![Evv]\!]_{\ell[R/\pi_1]} \cdot [\![R\,v]\!]_{\ell[R/\pi_1]} = \tfrac{1}{2} \cdot \tfrac{1}{2} = \tfrac{1}{4} \\
\pi_n(v) &= \left(\tfrac{1}{2}\right)^n \\
\pi_\omega(v) &= 0
\end{aligned}$$

So $[\![\varphi_{\text{infpath}}(v)]\!]_\ell = 0$ even though $\mathfrak{A}_\ell \models \varphi(v)$. The reason is that $\mathbb{V}$ is not chain-positive: All values $\pi_n(v)$ are positive but the infimum is still 0.

Nevertheless, one can argue that the computation in the Viterbi semiring makes sense. If we have to use the fact Evv infinitely often to establish a proof of φ and we only have low confidence in this edge, then it is reasonable to assign an overall confidence score of 0 to φ. So chain-positivity is not a hard requirement for applications.

The example illustrates that the interpretation of gfp-formulae can be inconsistent with the semantics of $\mathfrak{A}_\ell$. To achieve compatibility with standard semantics, we therefore require that the semiring is *chain-positive*, which intuitively means that infima preserve truth (where we again understand truth as non-zero values).

Theorem 5.7. *Let ℓ be a model-defining S-interpretation where S is both positive and chain-positive. For every formula $\varphi(\mathbf{x})$ and every tuple $\mathbf{a}$ from A we then have:*

$$\mathfrak{A}_\ell \models \varphi(\mathbf{a}) \quad \Longleftrightarrow \quad [\![\varphi(\mathbf{a})]\!]_\ell \neq 0$$

Proof. Recall that $\mathfrak{A}_\ell \models L \iff \ell(L) \neq 0$ by definition, so the statement holds for literals φ. This can be lifted to arbitrary formulae by induction on φ. We only have to consider formulae in negation normal form, as φ and $\mathsf{nnf}(\varphi)$ are interpreted in the same way by both $\mathfrak{A}_\ell$ and semiring semantics $[\![\cdot]\!]_\ell$.

- $\varphi = \varphi_1 \wedge \varphi_2$. We need that S is positive and apply the induction hypothesis. The cases for $\vee$, $\exists$, $\forall$ are analogous (recall that the universe is finite).

$$\begin{aligned}[\![\varphi(\mathbf{a})]\!]_\ell \neq 0 &\overset{\text{positive}}{\Longleftrightarrow} [\![\varphi_1(\mathbf{a})]\!]_\ell \neq 0 \text{ and } [\![\varphi_2(\mathbf{a})]\!]_\ell \neq 0 \\ &\overset{\text{IH}}{\Longleftrightarrow} \mathfrak{A}_\ell \models \varphi_1(\mathbf{a}) \text{ and } \mathfrak{A}_\ell \models \varphi_2(\mathbf{a}) \iff \mathfrak{A}_\ell \models \varphi(\mathbf{a})\end{aligned}$$

- $\varphi = [\mathsf{gfp}\, R\mathbf{x}.\, \vartheta](\mathbf{y})$. The proof is again by induction on the fixed-point iterations $(R_\beta)_{\beta \in \mathsf{On}}$ in standard semantics and $(\pi_\beta)_{\beta \in \mathsf{On}}$ in the semiring S. We claim that for all β, we have $\mathbf{a} \in R_\beta \iff \pi_\beta(\mathbf{a}) \neq 0$.

  For $\beta = 0$ this holds by definition. For successor ordinals, $\pi_{\beta+1}(\mathbf{a}) = [\![\vartheta(\mathbf{a})]\!]_{\ell[R/\pi_\beta]}$. Let σ be the valuation with $\sigma(\mathbf{a}) \in \{0,1\}$ and $\sigma(\mathbf{a}) = 0 \iff \pi_\beta(\mathbf{a}) \neq 0$. We need σ for the induction hypothesis on ϑ which requires a model-defining interpretation. As R only occurs positively in ϑ, we can interpret negative literals $\neg R\mathbf{b}$ arbitrarily, so we use the model-defining interpretation $\ell' = \ell[R/\pi_\beta, \neg R/\sigma]$. The induced model $\mathfrak{A}_{\ell'}$ interprets R according to π_β which, by induction, corresponds to R_β. Hence:

$$\begin{aligned}\pi_{\beta+1}(\mathbf{a}) = [\![\vartheta(\mathbf{a})]\!]_{\ell'} \neq 0 &\overset{\text{IH}}{\Longleftrightarrow} \mathfrak{A}_{\ell'} \models \vartheta(\mathbf{a}) \\ &\Longleftrightarrow \mathfrak{A}_\ell \models \vartheta(R_\beta, \mathbf{a}) \iff \mathbf{a} \in R_{\beta+1}\end{aligned}$$

  For limit ordinals, we use that S is chain-positive:

$$\begin{aligned}\pi_\lambda(\mathbf{a}) \neq 0 &\overset{\text{chain-pos}}{\Longleftrightarrow} \pi_\beta(\mathbf{a}) \neq 0 \text{ for all } \beta < \lambda \\ &\overset{\text{IH}}{\Longleftrightarrow} \mathbf{a} \in R_\beta \text{ for all } \beta < \lambda \iff \mathbf{a} \in R_\lambda\end{aligned}$$

  This proves the claim which in turn suffices to prove the case for gfp-formulae. The proof for lfp-formulae is analogous, but does not require chain-positivity. $\square$

These results show a compatibility with standard semantics similar to FO – if we additionally assume that the semiring is chain-positive. Of special interest is the interplay of semiring semantics and negation (as in proposition 2.11) since this includes duality. We consider this in detail in the following section.

5.3 Duality

One reason to analyze duality is that we defined semiring semantics in terms of negation normal form which relies on the equivalence

$$[\mathsf{lfp}\, R\,\mathbf{x}.\ \vartheta](\mathbf{y}) \equiv \neg[\mathsf{gfp}\, R\,\mathbf{x}.\ \overline{\vartheta}](\mathbf{y}) \quad \text{where } \overline{\vartheta} \text{ is the formula} \neg\vartheta[R/\neg R]$$

in standard LFP. This equivalence is based on the general duality of least and greatest fixed points in the standard domain of LFP, the powerset lattice $\mathcal{P}(A^k)$ (cf. [GKL+07]). Let $\overline{B} = A^k \setminus B$ denote the complement of a subset $B \subseteq A^k$ and let $F : \mathcal{P}(A^k) \to \mathcal{P}(A^k)$ be monotone, then we can express duality as

$$\mathsf{lfp}\, F = \overline{\mathsf{gfp}\, \overline{F}} \quad \text{where } \overline{F} \text{ is the mapping } B \mapsto \overline{F(\overline{B})}$$

With semirings, there is no way to express negation explicitly. The reason is that we have multiple truth values, but only use 0 to represent false, hence 0 has no unique negation. However, the compatibility result in theorem 5.7 shows that reasonable assignments of the literals lift to reasonable interpretations of LFP, including negation and duality:

Corollary 5.8. *If ℓ is a model-defining S-interpretation for a positive and chain-positive semiring S, then one of the following two values is 0 and the other is non-zero (for all R, $\mathbf{x}$, ϑ, $\mathbf{a}$):*

$$[\![[\mathsf{lfp}\, R\,\mathbf{x}.\ \vartheta](\mathbf{a})]\!]_\ell\,, \qquad [\![[\mathsf{gfp}\, R\,\mathbf{x}.\ \neg\vartheta[R/\neg R]](\mathbf{a})]\!]_\ell$$

We want to generalize this duality in the context of semirings in the same vein as for the powerset lattice. It is useful to again view a set $B \subseteq A^k$ as a function $f : A^k \to \mathbb{B}$. If $\overline{f} : A^k \to \mathbb{B}$ is the function corresponding to $\overline{B}$, then the complementarity of B and $\overline{B}$ means that for every $\mathbf{a}$, one of $f(\mathbf{a})$ and $\overline{f}(\mathbf{a})$ is $\bot$ and the other $\top$. Alternatively, we can say that $f(\mathbf{a}) \cdot \overline{f}(\mathbf{a}) = 0$ and $f(\mathbf{a}) + \overline{f}(\mathbf{a}) \neq 0$. These two statements are equivalent in $\mathbb{B}$, but we can generalize this idea to arbitrary semirings where they lead to different concepts of complements. So instead of having a unique complement, we propose several ways to express complementarity of two semiring values.

Definition 5.9. Let S be a cpo semiring. We provide three possible definitions of when we consider two values $a, b \in S$ to be *complementary*:

(C1) if $a \cdot b = 0$,

(C2) if $a + b \neq 0$,

(C3) if $a = 0$ and $b \neq 0$, or vice versa.

We lift this to functions $\pi : A \to S$ (for some set A). Fixing one of these definitions, we say that two functions π, σ are *complementary* (in this definition) if $\pi(a)$ and $\sigma(a)$ are complementary (in this definition) for all $a \in A$.

All three notions essentially distinguish between 0 and positive values, in line with our usual interpretation. (C1) expresses consistency whereas (C2) can be seen as completeness condition. Requiring both (C1) and (C2) is equivalent to (C3) in positive semirings.

Recall that for the powerset lattice, we consider $\overline{F} : B \mapsto \overline{F(\overline{B})}$ as the *dual* operator to F. Note that this does not mean that F and $\overline{F}$ are complementary (this would be the case for $\overline{F} : B \mapsto \overline{F(B)}$). Instead, we can formulate the duality of operators on semirings in the following way.

Definition 5.10. Let S be a cpo semiring and $F, G : S^A \to S^A$ two operators on functions (for some set A). Fixing one notion of complementarity, we say that F and G are *dual* if $F(\pi)$ and $G(\sigma)$ are complementary for all complementary functions $\pi, \sigma : A \to S$.

For example, $F(\top)$ and $G(\bot)$ must always be complementary. The duality of operators leads to complementarity of their least and greatest fixed points if we make reasonable assumptions on the semiring. This provides an adaption of the classic duality.

Theorem 5.11. *Let S be a cpo semiring, A a set and fix one notion of complementarity. Then S^A is a cpo semiring and if $F, G : S^A \to S^A$ are monotone dual operators, then $\mathsf{lfp}(F)$ and $\mathsf{gfp}(G)$ are complementary if S has the following properties:*

(C1) *S is positive or continuous,*

(C2),(C3) *S is chain-positive,*

Proof. The set S^A is fully chain-complete as function semiring over S, so we can use induction on the fixed-point iterations $(\pi_\beta)_{\beta \in \mathsf{On}}$ for $\mathsf{lfp}(F)$ and $(\sigma_\beta)_{\beta \in \mathsf{On}}$ for $\mathsf{gfp}(G)$.

By definition, $\pi_0 = \top$ and $\sigma_0 = \bot$ are complementary. For $\beta + 1$, we know by induction that π_β and σ_β are complementary. As F and G are dual, this implies that $\pi_{\beta+1} = F(\pi_\beta)$ and $\sigma_{\beta+1} = G(\sigma_\beta)$ are complementary.

The only interesting case where we need the additional assumptions is for limit ordinals. We have to show that $\pi_\lambda = \bigsqcup\{\pi_\beta \mid \beta < \lambda\}$ and $\sigma_\lambda = \bigsqcap\{\sigma_\beta \mid \beta < \lambda\}$ are complementary. Operations on S^A are defined pointwise, so we consider $\pi_\lambda(a)$ and $\sigma_\lambda(a)$ for an arbitrary $a \in A$. We know by induction that $\pi_\beta(a)$ and $\sigma_\beta(a)$ are complementary (for all $\beta < \lambda$).

Let us first consider (C1) if S is positive. Complementarity then implies $\pi_\beta(a) \cdot \sigma_\beta(a) = 0$ and by positivity, $\pi_\beta(a) = 0$ or $\sigma_\beta(a) = 0$. If there is a $\beta < \lambda$ with $\sigma_\beta(a) = 0$, then certainly $\bigsqcap\{\sigma_\beta(a) \mid \beta < \lambda\} = 0$ and thus $\pi_\lambda(a) \cdot \sigma_\lambda(a) = 0$. Otherwise, $\pi_\beta(a) = 0$ for all $\beta < \lambda$, thus $\bigsqcup\{\pi_\beta(a) \mid \beta < \lambda\} = 0$ and again $\pi_\lambda(a) \cdot \sigma_\lambda(a) = 0$.

For (C2), we similarly have $\pi_\beta(a) + \sigma_\beta(a) \neq 0$, and thus $\pi_\beta(a) \neq 0$ or $\sigma_\beta(a) \neq 0$ (recall that $+$-positivity holds in all cpo semirings). If there is a $\beta < \lambda$ with $\pi_\beta(a) \neq 0$, then certainly $\bigsqcup\{\pi_\beta(a) \mid \beta < \lambda\} \neq 0$. Otherwise, $\sigma_\beta(a) \neq 0$ for all $\beta < \lambda$ and we use chain-positivity to conclude $\bigsqcap\{\sigma_\beta(a) \mid \beta < \lambda\} \neq 0$. In both cases, $\pi_\lambda(a) + \sigma_\lambda(a) \neq 0$.

For (C3), we can apply the same arguments. By complementarity, $\pi_\beta(a) = 0$ and $\sigma_\beta(a) \neq 0$ or vice versa (for each $\beta < \lambda$). Following the two case distinctions above, one of $\pi_\lambda(a)$ and $\sigma_\lambda(a)$ must be 0 and the other $\neq 0$.

It remains to consider (C1) if S is continuous (but not positive):

$$\begin{aligned} \pi_\lambda(a) \cdot \sigma_\lambda(a) &= \bigsqcup\{\pi_\beta(a) \mid \beta < \lambda\} \cdot \bigsqcap\{\sigma_\delta \mid \delta < \lambda\} \\ &= \bigsqcup \left\{\pi_\beta(a) \cdot \bigsqcap\{\sigma_\delta(a) \mid \delta < \lambda\} \;\middle|\; \beta < \lambda\right\} \\ &= \bigsqcup \left\{ \bigsqcap\{\pi_\beta(a) \cdot \sigma_\delta(a) \mid \delta < \lambda\} \;\middle|\; \beta < \lambda\right\} \end{aligned}$$

Fix any $\beta < \lambda$. Then $\pi_\beta(a) \cdot \sigma_\delta(a) = 0$ for $\delta = \beta$ and thus $\bigsqcap\{\pi_\beta(a) \cdot \sigma_\delta(a) \mid \delta < \lambda\} = 0$. As this holds for every β, we have $\pi_\lambda(a) \cdot \sigma_\lambda(a) = \bigsqcup\{0\} = 0$ as claimed. □

While we have mostly focused on positive semirings so far, it is interesting to see that continuous semirings provide an alternative to ensure consistency. This becomes relevant when we introduce polynomials with dual-indeterminates in section 6.6 which have divisors of 0 and are thus not positive. To see that we indeed need the additional requirements, consider the following two examples.

Example 5.12. For (C1), recall the semiring $S = \mathbb{B} \times \mathbb{V} \setminus \{(\bot, 1)\}$ from example 4.29. which is neither continuous nor positive. We define operators $F, G : S \to S$ by

$$F((b,v)) = (b, \tfrac{1+v}{2}), \qquad G((b,v)) = (b, 0)$$

Note that we can lift these operators to S^A (that is, $F, G : S^A \to S^A$) by considering $A = \{a\}$ and identifying functions $\pi : A \to S$ with their value $\pi(a)$.

F and G are monotone and dual to each other. For their duality, note that (b_1, v_1) and

(b_2, v_2) are complementary iff $b_1 b_2 = \bot$ and $v_1 v_2 = 0$. Then also $F((b_1, v_1)) = (b_1, v_1')$ and $G((b_2, v_2)) = (b_2, 0)$ are complementary, as $v_1' \cdot 0 = 0$. The iteration yields (for $n \geq 1$)

$$F^n((\bot, 0)) \cdot G^n((\top, 1)) = (\bot, 1 - \tfrac{1}{2^n}) \cdot (\top, 0) = (\bot, 0)$$

which is still complementary, but the fixed points are not (recall that $(\bot, 1) \notin S$):

$$\mathsf{lfp}(F) \cdot \mathsf{gfp}(G) = (\top, 1) \cdot (\top, 0) = (\top, 0) \neq 0$$

Example 5.13. For definition (C3), consider the Viterbi semiring $\mathbb{V}$ which is not chain-positive. We again use a singleton set $A = \{a\}$ and identify functions $\pi : A \to \mathbb{V}$ with their values $\pi(a) \in \mathbb{V}$. We can then define the monotone operators $F(x) = x$ and $G(y) = \frac{y}{2}$. If x, y are complementary, then so are x and $\frac{y}{2}$, thus F and G are dual.

In the iteration, $F^n(0) = 0$ and $G^n(1) = \frac{1}{2^n}$ are complementary, but for the fixed points,

$$\mathsf{lfp}(F) + \mathsf{gfp}(G) = 0 + 0 = 0$$

Let us apply this general concept of duality back to logic. We have already seen duality for model-defining valuations (if S is positive and chain-positive) in corollary 5.8. It is not an accident that the corollary requires S to be positive – we need this to guarantee that the update operators for ϑ and $\overline{\vartheta} = \neg\vartheta[R/\neg R]$ are dual. In addition, we must require that ℓ maps opposing literals to complementary values, since positive literals in ϑ become negative literals in $\overline{\vartheta}$ (except for the relation symbol R).

By showing duality of the update operators F_ℓ^ϑ and $F_\ell^{\overline{\vartheta}}$, we can apply theorem 5.11 to obtain analogues of proposition 2.11 for LFP. As for FO, these results justify the use of the negation normal form in the definition of semiring semantics.

Proposition 5.14. *Let S be a cpo semiring, $\vartheta(R, \mathbf{x})$ an LFP formula and fix one of the definitions of complementarity. If ℓ is an S-interpretation such that $\ell(L)$ and $\ell(\neg L)$ are complementary for every literal L and additionally,*

(C1) S is positive or continuous,

(C2,3) S is positive and chain-positive,

then F_ℓ^ϑ and $F_\ell^{\overline{\vartheta}}$ with $\overline{\vartheta} = \neg\vartheta[R/\neg R]$ are dual operators.

Proof. We only consider (C1) and (C2), as (C3) is equivalent to their combination in positive semirings. The proof is by induction on ϑ. The case for literals holds by assumption. For the remaining cases, let π and σ be complementary valuations.

- $\vartheta = \vartheta_1 \wedge \vartheta_2$. Then $\mathsf{nnf}(\overline{\vartheta}) = \mathsf{nnf}(\overline{\vartheta}_1) \vee \mathsf{nnf}(\overline{\vartheta}_2)$ and we obtain for (C1):

$$\begin{aligned} F^{\vartheta}(\pi) \cdot F^{\overline{\vartheta}}(\sigma) &= (F^{\vartheta_1}(\pi) \cdot F^{\vartheta_2}(\pi)) \cdot (F^{\overline{\vartheta}_1}(\sigma) + F^{\overline{\vartheta}_2}(\sigma)) \\ &= \underbrace{F^{\vartheta_1}(\pi) \cdot F^{\overline{\vartheta}_1}(\sigma)}_{=0} \cdot F^{\vartheta_2}(\pi) + F^{\vartheta_1}(\pi) \cdot \underbrace{F^{\vartheta_2}(\pi) \cdot F^{\overline{\vartheta}_2}(\sigma)}_{=0} \stackrel{\text{IH}}{=} 0 \end{aligned}$$

For (C2), we have

$$F^{\vartheta}(\pi) + F^{\overline{\vartheta}}(\sigma) = (F^{\vartheta_1}(\pi) \cdot F^{\vartheta_2}(\pi)) + (F^{\overline{\vartheta}_1}(\sigma) + F^{\overline{\vartheta}_2}(\sigma)) \stackrel{(*)}{\neq} 0$$

Note that $(*)$ holds if $F_\ell^{\overline{\vartheta}_1}(\sigma) \neq 0$ or $F_\ell^{\overline{\vartheta}_1}(\sigma) \neq 0$. Otherwise, it follows by induction that $F_\ell^{\vartheta_1}(\pi) \neq 0$ and $F_\ell^{\vartheta_2}(\pi) \neq 0$, so $(*)$ holds since S is positive. The cases for $\vee$, $\exists$ and $\forall$ are analogous.

- $\vartheta = [\mathsf{lfp}\, P\, \mathbf{y}.\, \psi](\mathbf{z})$. Then $\mathsf{nnf}(\overline{\vartheta}) = [\mathsf{gfp}\, P\, \mathbf{y}.\, \mathsf{nnf}(\psi')](\mathbf{z})$ with $\psi' = \neg\psi[R/\neg R, P/\neg P]$. We have:

$$\begin{aligned} F_\ell^{\vartheta,\alpha}(\pi) &= \big(\mathsf{lfp}\, F^{\psi}_{\ell[R/\pi]}\big)(\alpha(\mathbf{z})) \\ F_\ell^{\overline{\vartheta},\alpha}(\sigma) &= \big(\mathsf{gfp}\, F^{\psi'}_{\ell[R/\sigma]}\big)(\alpha(\mathbf{z})) \stackrel{(*)}{=} \big(\mathsf{gfp}\, F^{\neg\psi[P/\neg P]}_{\ell[\neg R/\sigma]}\big)(\alpha(\mathbf{z})) \end{aligned}$$

where $(*)$ holds as R can only occur positively in ϑ and thus also in ψ, so we can replace R by $\neg R$ in ψ' if we also modify ℓ accordingly. As R only occurs positively in ψ and only negatively in $\neg\psi[P/\neg P]$, we can set $\ell' = \ell[R/\pi, \neg R/\sigma]$ such that

$$\begin{aligned} F_\ell^{\vartheta,\alpha}(\pi) &= \big(\mathsf{lfp}\, F^{\psi}_{\ell'}\big)(\alpha(\mathbf{z})) \\ F_\ell^{\overline{\vartheta},\alpha}(\sigma) &= \big(\mathsf{gfp}\, F^{\neg\psi[P/\neg P]}_{\ell'}\big)(\alpha(\mathbf{z})) \end{aligned}$$

which allows us to apply the induction hypothesis for ψ to conclude that $F^{\psi}_{\ell'}$ and $F^{\neg\psi[P/\neg P]}_{\ell'}$ are dual. By theorem 5.11, $F_\ell^{\vartheta,\alpha}(\pi)$ and $F_\ell^{\overline{\vartheta},\alpha}(\sigma)$ are thus complementary.

The proof for gfp-formulae is analogous. □

For (C3), we obtain corollary 5.8 as a consequence of this result and theorem 5.11. For (C1) and (C2), we get the following corollary.

Corollary 5.15. *Let $\ell : \mathsf{Lit}_A \to S$ be an S-interpretation. Then the following holds:*

(1) *If S is positive or continuous, and $\ell(L) \cdot \ell(\neg L) = 0$ for all $L \in \mathsf{Lit}_A$, then this property lifts to all LFP formulae φ, that is: $[\![\varphi]\!]_\ell \cdot [\![\neg\varphi]\!]_\ell = 0$.*

(2) *If S is positive and chain-positive, and $\ell(L) + \ell(\neg L) = 0$ for all $L \in \mathsf{Lit}_A$, then this property lifts to all LFP formulae φ, that is: $[\![\varphi]\!]_\ell + [\![\neg\varphi]\!]_\ell \neq 0$.*

The above considerations show that the definition of interpretations for LFP in terms of cpo semirings is reasonable and can be seen as a generalization of standard semantics. With model-defining interpretations in chain-positive (and positive) semirings, we mainly generalize standard semantics by allowing several values for truth and otherwise retain standard properties. With arbitrary semirings and interpretations, we allow a broader generalization which leads to different properties but is still useful for provenance analysis. Examples are the semiring $\mathbb{V}$ or the dual-indeterminate polynomials in chapter 6.

5.4 Fundamental Property

Towards a general form of provenance, we want to use homomorphisms to switch from general to application specific semirings. Unlike with FO, semiring homomorphisms are not sufficient for this purpose since they do not preserve fixed points.

Example 5.16. Recall example 5.6. For $\varphi(v) = [\mathsf{gfp}\, R\, x.\ \exists y(Exy \wedge Ry)](v)$, we obtained the fixed-point iteration with $\pi_n(v) = (\frac{1}{2})^n$ and the fixed-point $[\![\varphi(v)]\!]_\ell = \pi_\omega(v) = 0$ for the $\mathbb{V}$-interpretation ℓ with $\ell(Evv) = \frac{1}{2}$.

Consider the truth projection homomorphism $\dagger : \mathbb{V} \to \mathbb{B}$. If we compute the fixed point in $\mathbb{V}$ and then apply the homomorphism, we get

$$\dagger([\![\varphi(v)]\!]_\ell) = \dagger(0) = \bot$$

If, on the other hand, we first apply the homomorphism to ℓ and then compute the interpretation (and thus the fixed point) in $\mathbb{B}$, we get

$$[\![\varphi(v)]\!]_{\dagger \circ \ell} = \big(\bigsqcap\{\dagger(\pi_n) \mid n < \omega\}\big)(v) = \big(\bigsqcap\{\top\}\big)(v) = \top$$

So in this example, $[\![\varphi]\!]^\alpha_{\dagger\circ\ell} \neq \dagger([\![\varphi]\!]^\alpha_\ell)$. That is, the homomorphism $\dagger$ does not commute with the interpretation of φ in the Viterbi semiring.

The reason for the different results in the example is that the homomorphism does not preserve the infimum (which in this case happens because $\mathbb{V}$ is not chain-positive). We therefore consider cpo-semiring homomorphisms which are in particular continuous, so they preserve suprema and infima and are thus suitable for our needs.

Theorem 5.17 (fundamental property)**.** *Let $h : S \to T$ be a cpo-semiring homomorphism and $\ell : \mathsf{Lit}_A \to S$ an S-interpretation. Then $h \circ \ell : \mathsf{Lit}_A \to T$ is a T-interpretation and we have $h([\![\varphi]\!]^\alpha_\ell) = [\![\varphi]\!]^\alpha_{h\circ\ell}$ for all LFP formulae φ and valuations α.*

As diagram:

$$\begin{tikzcd} & \mathsf{Lit}_A \arrow[dl, "\ell"'] \arrow[dr, "h \circ \ell"] & \\ S \arrow[rr, "h"] & & T \end{tikzcd} \quad \Longrightarrow \quad \begin{tikzcd} & \mathrm{LFP} \arrow[dl, "\ell"'] \arrow[dr, "h \circ \ell"] & \\ S \arrow[rr, "h"] & & T \end{tikzcd}$$

Proof. We show that for all LFP formulae φ in negation normal form, $h(\llbracket\varphi\rrbracket_\ell^\alpha) = \llbracket\varphi\rrbracket_{h\circ\ell}^\alpha$ holds for all valuations α and S-interpretations ℓ, which suffices to prove the theorem. The proof is by induction on φ.

- For literals $\varphi = R\,\mathbf{x}$, we have $h(\llbracket\varphi\rrbracket_\ell^\alpha) = h(\ell(R\,\alpha(\mathbf{x}))) = (h\circ\ell)(R\,\alpha(\mathbf{x})) = \llbracket\varphi\rrbracket_{h\circ\ell}^\alpha$. The same holds for negative literals $\varphi = \neg R\,\mathbf{x}$.
- For $\varphi = \varphi_1 \wedge \varphi_2$, we use that h is a semiring homomorphism, so we have
$$h(\llbracket\varphi\rrbracket_\ell^\alpha) = h(\llbracket\varphi_1\rrbracket_\ell^\alpha \cdot \llbracket\varphi_2\rrbracket_\ell^\alpha) \overset{\text{hom}}{=} h(\llbracket\varphi_1\rrbracket_\ell^\alpha)\cdot h(\llbracket\varphi_2\rrbracket_\ell^\alpha) \overset{\text{IH}}{=} \llbracket\varphi_1\rrbracket_{h\circ\ell}^\alpha \cdot \llbracket\varphi_2\rrbracket_{h\circ\ell}^\alpha = \llbracket\varphi\rrbracket_{h\circ\ell}^\alpha$$
The proofs for $\vee$ and $\exists$, $\forall$ are analogous.
- For $\varphi = [\mathsf{gfp}\,R\,\mathbf{x}.\ \vartheta](\mathbf{y})$ with R of arity k, we again proceed via the fixed-point iteration. Let $(\pi_\beta)_{\beta\in\mathsf{On}}$ be the iteration for ℓ and let $(\sigma_\beta)_{\beta\in\mathsf{On}}$ be the iteration for $h\circ\ell$. We show by transfinite induction that for all $\beta\in\mathsf{On}$:
$$h(\pi_\beta) = \sigma_\beta, \quad \text{i.e.,} \quad h(\pi_\beta(\mathbf{a})) = \sigma_\beta(\mathbf{a}) \text{ for all } \mathbf{a}\in A^k$$
  - For $\beta = 0$, we have $h(\top) = \top$, as h is a cpo homomorphism.
  - For successor ordinals, we can apply the induction hypothesis. Note that
$$\begin{aligned} \pi_{\beta+1} &= F_\ell^\vartheta(\pi_\beta) = \big(\mathbf{a} \mapsto \llbracket\vartheta\rrbracket_{\ell[R/\pi_\beta]}^{\alpha[\mathbf{x}/\mathbf{a}]}\big) \\ \sigma_{\beta+1} &= F_{h\circ\ell}^\vartheta(\sigma_\beta) = \big(\mathbf{a} \mapsto \llbracket\vartheta\rrbracket_{(h\circ\ell)[R/\sigma_\beta]}^{\alpha[\mathbf{x}/\mathbf{a}]}\big) \overset{(*)}{=} \big(\mathbf{a} \mapsto \llbracket\vartheta\rrbracket_{h\,\circ\,(\ell[R/\pi_\beta])}^{\alpha[\mathbf{x}/\mathbf{a}]}\big) \end{aligned}$$
    In $(*)$, we use the induction hypothesis $h(\pi_\beta) = \sigma_\beta$. So for every $\mathbf{a}$:
$$h(\pi_{\beta+1}(\mathbf{a})) = h(\llbracket\vartheta\rrbracket_{\ell[R/\pi_\beta]}^{\alpha[\mathbf{x}/\mathbf{a}]}) \overset{\text{IH}}{=} \llbracket\vartheta\rrbracket_{h\,\circ\,(\ell[R/\pi_\beta])}^{\alpha[\mathbf{x}/\mathbf{a}]}) = \sigma_{\beta+1}(\mathbf{a})$$
    where we apply the induction hypothesis of the outer induction.
  - For limit ordinals, we use that h is continuous, so for every $\mathbf{a}$:
$$\begin{aligned} h(\pi_\lambda(\mathbf{a})) &= h(\textstyle\bigsqcap\{\pi_\beta(\mathbf{a}) \mid \beta<\lambda\}) \\ &\overset{\text{hom}}{=} \textstyle\bigsqcap\{h(\pi_\beta(\mathbf{a})) \mid \beta<\lambda\} \\ &\overset{\text{IH}}{=} \textstyle\bigsqcap\{\sigma_\beta(\mathbf{a}) \mid \beta<\lambda\} = \sigma_\lambda(\mathbf{a}) \end{aligned}$$

  This closes the proof for gfp-formulae, as for sufficiently large β, we have
$$h(\llbracket\varphi\rrbracket_\ell^\alpha) = h(\pi_\beta(\alpha(\mathbf{y}))) = \sigma_\beta(\alpha(\mathbf{y})) = \llbracket\varphi\rrbracket_{h\circ\ell}^\alpha$$
  The proof for lfp-formulae is analogous. $\square$

This result leads to the question which semiring homomorphisms are additionally cpo homomorphisms. We have already seen in example 4.34 that the truth projection is continuous (and thus a cpo-semiring homomorphism) precisely for the chain-positive positive semirings. Another example is the isomorphism $\mathbb{T} \to \mathbb{V}$, $x \mapsto e^{-x}$ that is continuous as composition of continuous functions (on the real numbers). So there certainly are natural homomorphisms that preserve semiring semantics $[\![\varphi]\!]_\ell$.

To see why the fundamental property is useful, we show how it can be utilized to provide a shorter proof of theorem 5.7 which states that $\mathfrak{A}_\ell \models \varphi(\mathbf{a}) \iff [\![\varphi(\mathbf{a})]\!]_\ell \neq 0$ holds for all model-defining S-interpretations ℓ where S is both positive and chain-positive.

Alternative proof of theorem 5.7.
We have already argued that $\dagger_S \circ \ell = \ell_{\mathfrak{A}_\ell}$ (see section 2.2) and hence

$$[\![\varphi(\mathbf{a})]\!]_{\dagger_S \circ \ell} = \top \iff \mathfrak{A}_\ell \models \varphi(\mathbf{a})$$

by proposition 5.5. As $\dagger_S$ is a cpo-semiring homomorphism, we then have:

$$\mathfrak{A}_\ell \models \varphi(\mathbf{a}) \iff [\![\varphi(\mathbf{a})]\!]_{\dagger_S \circ \ell} = \top \iff \dagger_S([\![\varphi(\mathbf{a})]\!]_\ell) = \top \iff [\![\varphi(\mathbf{a})]\!]_\ell \neq 0 \qquad \square$$

5.5 Logical Fragments and Continuity

The semiring semantics for full LFP require the semiring to have both suprema and infima for chains of arbitrary cardinality. In the following section, we discuss how this can be relaxed for two logical fragments: The positive fragment without greatest fixed points and the alternation-free fragment which basically avoids nested fixed points.

Definition 5.18. An LFP formula is *positive* if it is in negation normal form and contains no gfp-subformulae. The *positive* fragment (posLFP) consists of all positive formulae.

The interpretation of greatest fixed points leads to several complications, most notably the notion of chain-positive semirings, which can be avoided when we only consider the positive fragment. Instead of discussing these advantages in general, we want to bring an interesting class of semirings into effect which we have not yet considered in detail, the continuous semirings (recall that all suitable application semirings we consider are also continuous). This is especially useful for posLFP where, due to the simpler formulae, we can work with ω-continuous semirings (cf. [GT19]). Indeed, ω-continuous semirings have already been proposed for provenance analysis of positive LFP and similar logics (e.g., [GKT07, GT19, Mrk18]).

Theorem 5.19. *The lifting of S-interpretations to posLFP according to definition 3.4 is well-defined if S is an ω-continuous semiring.*

To see why suprema of ω-chains suffice to express least fixed points, we observe that the update operators F^ϑ are ω-continuous if ϑ is positive. This allows us to apply Kleene's fixed-point theorem 4.6. We can formulate the slightly stronger result:

Proposition 5.20. *Let S be an ω-continuous semiring and φ a positive formula. Then $[\![\varphi]\!]_\ell$ is ω-continuous in ℓ, so for every ω-chain $\mathcal{L} = (\ell_i)_{i<\omega}$ of S-interpretations,*

$$[\![\varphi]\!]_{\bigsqcup \mathcal{L}} = \bigsqcup \{ [\![\varphi]\!]_{\ell_i} \mid i < \omega \}$$

The proof is a straight-forward induction based on the ω-continuity of the semiring operations. We defer the proof to a similar statement for the alternation-free fragment. Let us first give an example which illustrates that such a general continuity statement does not hold for full LFP.

Example 5.21. Recall the Łukasiewicz semiring $\mathbb{L} = ([0,1], \max, \circ, 0, 1)$ with $a \circ b = \max(0, a+b-1)$ as introduced in section 4.3. Let us consider $\varphi = \exists y\, [\mathsf{gfp}\, R\,x.\, \vartheta](y)$ with $\vartheta = R\,x \wedge P\,x$ over the singleton universe $A = \{a\}$.

We define the (ascending) ω-chain $\mathcal{L} = (\ell_n)_{n<\omega}$ of $\mathbb{L}$-interpretations by

$$\ell_n(Pa) = 1 - \tfrac{1}{n} \text{ (for } n \geq 2), \qquad \ell_0(Pa) = \ell_1(Pa) = 0$$

such that $\left(\bigsqcup \mathcal{L}\right)(Pa) = 1$. The fixed-point iteration for $[\![\varphi]\!]_{\bigsqcup \mathcal{L}}$ yields $\pi_n(a) = 1$ (for all $n < \omega$), so we obtain $[\![\varphi]\!]_{\bigsqcup \mathcal{L}} = 1$.

For $[\![\varphi]\!]_{\ell_{10}}$ with $\ell_{10}(Pa) = 0.9$, we instead get:

$$\begin{aligned}
\pi_0(a) &= 1 \\
\pi_1(a) &= 1 \circ 0.9 = 1 - 0.1 = 0.9 \\
\pi_2(a) &= 0.9 \circ 0.9 = 0.9 - 0.1 = 0.8 \\
\\
\pi_\omega(a) &= 0
\end{aligned}$$

So the value decreases by 0.1 in every step until it reaches 0. For ℓ_n (with $n \geq 2$), we similarly have a decrease by $\frac{1}{n}$ in every step, so 0 is reached after n (and thus finitely many) steps. The fixed point is thus always 0 for every ℓ_n and we obtain

$$\bigsqcup \{ [\![\varphi]\!]_{\ell_n} \mid n < \omega \} = \bigsqcup \{0\} = 0 \neq 1 = [\![\varphi]\!]_{\bigsqcup \mathcal{L}}$$

So although $\mathbb{L}$ is continuous, $[\![\varphi]\!]_\ell$ is not continuous (or ω-continuous) in ℓ.

Another advantage of posLFP is that ω-continuous homomorphisms are sufficient to preserve $[\![\varphi]\!]_\ell$. As we will see in chapter 6, this yields a much simpler situation than using cpo-semiring homomorphisms for full LFP.

The second fragment we consider is the alternation-free fragment. Intuitively, this allows both least and greatest fixed points as long as they are not nested (or, if they are nested, they must be independent of each other). Note that we allow alternation of several lfp-formulae (or several gfp-formulae). More background can be found in [GKL+07].

Definition 5.22. An LFP formula is *alternating* if its negation normal form contains two subformulae $\varphi = [\mathsf{lfp}\, R\, \mathbf{x}.\, \hat{\varphi}](\mathbf{y})$ and $\varphi' = [\mathsf{gfp}\, R'\, \mathbf{x}'.\, \hat{\varphi}'](\mathbf{y}')$ such that φ' is a subformula of $\hat{\varphi}$ and R occurs free in $\hat{\varphi}'$, or vice versa (φ is a subformula of $\hat{\varphi}'$ and R' occurs free in $\hat{\varphi}$). The *alternation-free* fragment contains all formulae which are not alternating.

We first observe that the previous example features an alternation-free formula, so we cannot obtain the same continuity result as for posLFP. The problem in the example is that the ℓ_n differ in the interpretation of P which occurs free in the formula. This cannot happen in alternation-free LFP if the values ℓ_n result from an outer fixed-point formula, as we would then have an alternation between the outer formula (defining P) and the inner formula (using R). We can formulate a slightly restricted continuity result which shows that we still only need ω-chains. As we allow both least and greatest fixed points, we now have to work work with fully ω-continuous semirings.

Theorem 5.23. *The lifting of S-interpretations to the alternation-free fragment of LFP according to definition 3.4 is well-defined if S is a fully ω-continuous semiring.*

To prove the theorem, we show the following continuity statement which is tailored to alternation-free formulae by restricting the chains of S-interpretations.

Lemma 5.24. *Let φ be an alternation-free formula, S a fully ω-continuous semiring and ℓ an S-interpretation. Let further $R_1, \dots, R_n$ relation symbols (for some $n < \omega$). We then have the following statement (and its dual):*

Let $(\pi_1^{(i)})_{i<\omega}, \dots, (\pi_n^{(i)})_{i<\omega}$ be ascending (descending) ω-chains of S-valuations. Then $\mathcal{L} = (\ell_i)_{i<\omega}$ with $\ell_i = \ell[R_1/\pi_1^{(i)}, \dots, R_n/\pi_n^{(i)}]$ is an ascending (descending) ω-chain of S-interpretations. If $R_1, \dots, R_n$ do not occur free within any gfp-subformula (lfp-subformula) of φ, then

$$[\![\varphi]\!]_{\bigsqcup\mathcal{L}} = \bigsqcup\{[\![\varphi]\!]_{\ell_i} \mid i < \omega\} \quad \Big(\textit{or, dually, } [\![\varphi]\!]_{\bigsqcap\mathcal{L}} = \bigsqcap\{[\![\varphi]\!]_{\ell_i} \mid i < \omega\}\Big)$$

Moreover, if the above statement holds for a formula $\vartheta(P, \mathbf{x})$ where P does not occur free within any gfp*-subformula (*lfp*-subformula) of ϑ, then the update operator F_ℓ^ϑ is continuous (for every ℓ) in the sense that for every ascending (descending) ω-chain C,*

$$F_\ell^{\vartheta,P}(\bigsqcup C) = \bigsqcup\{F_\ell^{\vartheta,P}(\pi) \mid \pi \in C\} \quad \Big(\textit{or, dually, } F_\ell^{\vartheta,P}(\bigsqcap C) = \bigsqcap\{F_\ell^{\vartheta,P}(\pi) \mid \pi \in C\}\Big)$$

Proof. For the proof, we only show the statements about ascending ω-chains, as the argument for the dual statements is symmetric. Let us first prove the second part of the lemma, so let ℓ be arbitrary and consider an ascending ω-chain $C = (\pi_i)_{i<\omega}$. We set $\ell_i = \ell[P/\pi_i]$ and $\mathcal{L} = (\ell_i)_{i<\omega}$. Then $\mathcal{L}$ satisfies the conditions of the lemma and we have $\ell[P/\bigsqcup C] = \bigsqcup\{\ell[P/\pi_i] \mid i < \omega\} = \bigsqcup \mathcal{L}$, as the supremum is defined pointwise. By the assumptions on ϑ, we obtain

$$F_\ell^{\vartheta,P}(\bigsqcup C) = [\![\vartheta]\!]^\bullet_{\ell[P/\bigsqcup C]} = [\![\vartheta]\!]^\bullet_{\bigsqcup\mathcal{L}} = \bigsqcup\{[\![\vartheta]\!]^\bullet_{\ell_i} \mid i < \omega\} = \bigsqcup\{F_\ell^{\vartheta,P}(\pi_i) \mid i < \omega\}$$

The proof of the first part is by induction on the negation normal form of φ.

- For literals φ, the statement is trivial.
- For $\varphi = \varphi_1 \wedge \varphi_2$, the statement follows by induction as we can use the continuity of multiplication and the monotonicity of $[\![\varphi]\!]_\ell$ (with respect to ℓ):

$$\begin{aligned}
[\![\varphi]\!]_{\bigsqcup\mathcal{L}} = [\![\varphi_1]\!]_{\bigsqcup\mathcal{L}} \cdot [\![\varphi_1]\!]_{\bigsqcup\mathcal{L}} &\overset{\text{IH}}{=} \bigsqcup\{[\![\varphi_1]\!]_{\ell_i} \mid i < \omega\} \cdot \bigsqcup\{[\![\varphi_2]\!]_{\ell_j} \mid j < \omega\} \\
&\overset{\text{cont}}{=} \bigsqcup\{[\![\varphi_1]\!]_{\ell_i} \cdot [\![\varphi_2]\!]_{\ell_j} \mid i, j < \omega\} \\
&\overset{\text{mon}}{=} \bigsqcup\{[\![\varphi_1]\!]_{\ell_i} \cdot [\![\varphi_2]\!]_{\ell_i} \mid i < \omega\} = \bigsqcup\{[\![\varphi]\!]_{\ell_i} \mid i < \omega\}
\end{aligned}$$

  The same argument holds for $\vee$, $\exists$ and $\forall$.
- For $\varphi = [\mathsf{gfp}\, P\mathbf{x}.\, \vartheta](\mathbf{y})$, we know that $R_1, \ldots, R_n$ do not occur in ϑ. So $[\![\varphi]\!]_{\bigsqcup\mathcal{L}} = [\![\varphi]\!]_{\ell_i} = [\![\varphi]\!]_\ell$ (for all $i < \omega$) and the statement follows.
- For $\varphi = [\mathsf{lfp}\, P\mathbf{x}.\, \vartheta](\mathbf{y})$, we show by induction on the fixed-point iteration that

$$\pi^{(\beta)}_{\bigsqcup\mathcal{L}} = \bigsqcup\{\pi^{(\beta)}_{\ell_i} \mid i < \omega\}$$

  holds for every $\beta \le \omega$. For $\beta = 0$, this is trivial. For successor ordinals, we use the induction hypothesis for β (IH_2) and the outer induction hypothesis on ϑ (IH_1):

$$\begin{aligned}
\bigsqcup\big\{\pi^{(\beta+1)}_{\ell_i} \bigm| i < \omega\big\} = \bigsqcup\big\{[\![\vartheta]\!]^\bullet_{\ell_i[P/\pi^{(\beta)}_{\ell_i}]} \bigm| i < \omega\big\} \\
&\overset{\text{IH}_1}{=} [\![\vartheta]\!]^\bullet_{\bigsqcup\{\ell_i[P/\pi^{(\beta)}_{\ell_i}] \mid i<\omega\}} \\
&\overset{(*)}{=} [\![\vartheta]\!]^\bullet_{(\bigsqcup\mathcal{L})[P/\bigsqcup\{\pi^{(\beta)}_{\ell_i} \mid i<\omega\}]} \\
&\overset{\text{IH}_2}{=} [\![\vartheta]\!]^\bullet_{(\bigsqcup\mathcal{L})[P/\pi^{(\beta)}_{\bigsqcup\mathcal{L}}]} = \pi^{(\beta+1)}_{\bigsqcup\mathcal{L}}
\end{aligned}$$

where $(*)$ holds as the supremum of $\mathcal{L}$ is computed pointwise. When we apply IH_1, note that $\ell_i[P/\pi_{\ell_i}^{(\beta)}]$ defines an ascending ω-chain that satisfies the assumptions of the lemma: P does not occur within gfp-subformulae of φ, as φ is alternation-free.

For the limit ordinal ω, we use the fact that we can swap suprema:

$$\begin{aligned}\pi_{\bigsqcup\mathcal{L}}^{(\omega)} = \bigsqcup\{\pi_{\bigsqcup\mathcal{L}}^{(\beta)} \mid \beta < \omega\} &\stackrel{\mathrm{IH}}{=} \bigsqcup\Big\{\bigsqcup\{\pi_{\ell_i}^{(\beta)} \mid i < \omega\} \mid \beta < \omega\Big\} \\ &= \bigsqcup\Big\{\bigsqcup\{\pi_{\ell_i}^{(\beta)} \mid \beta < \omega\} \mid i < \omega\Big\} = \bigsqcup\{\pi_{\ell_i}^{(\omega)} \mid i < \omega\}\end{aligned}$$

This ends the induction on the iteration. For the statement about φ, note that the induction hypothesis on ϑ together with the second part of the lemma imply that F^ϑ is ω-continuous. So by Kleene's fixed-point theorem 4.6, $\pi_{\bigsqcup\mathcal{L}}^{(\omega)}$ and $\pi_{\ell_i}^{(\omega)}$ (for every $i < \omega$) are already the least fixed points and the statement follows. □

The above lemma is rather technical, but has a number of consequences. First, it implies theorem 5.23, as all update operators are fully ω-continuous such that Kleene's fixed-point theorem (and its dual version) apply and guarantee the existence of all fixed points. As the positive fragment is a subset of the alternation-free fragment, this also proves theorem 5.19 (if we drop the requirements on infima). Without greatest fixed points, the lemma can be simplified (we no longer need the assumptions on $R_1, \dots, R_n$ if there are no gfp subformulae), which leads to the continuity statement of proposition 5.20.

Open Question. *These results raise the question whether fixed-point iterations $(\pi_\beta)_{\beta\in\mathsf{On}}$ induced by LFP formulae always reach the fixed-point at π_ω in continuous semirings. This would imply that fully ω-continuous semirings are suitable for full LFP.*

If semiring semantics $[\![\varphi]\!]_\ell$ are continuous in ℓ, then all update operators F_ℓ^ϑ are continuous and Kleene's fixed-point theorem answers this question. Example 5.21 makes clear that this is not the case in general, so we have to seek other ways to answer this question. We provide a positive answer for absorptive semirings in section 6.4, following our analysis of absorptive polynomials $\mathbb{S}^\infty[X]$. The general case remains open.

Background: To understand why the continuity of $[\![\varphi]\!]_\ell$ fails and to provide connections to order theory beyond the context of logic, let us briefly discuss a related, more general question. Let S be a continuous semiring. Let $f : S \times S \to S$ be a function which is sup-continuous in the first and inf-continuous in the second argument. We define

- $g_y : x \mapsto f(x, y)$ for every $y \in S$
- $h_x : y \mapsto f(x, y)$ for every $x \in S$
- $G : y \mapsto \mathsf{lfp}(g_y)$,
- $H : x \mapsto \mathsf{gfp}(h_x)$

Then g_y is always sup-continuous and h_x always inf-continuous. The question now is: Does the continuity of f further imply that G is inf-continuous (and that H is sup-continuous)?

Example 5.25. Example 5.21 can be translated into this setting by choosing $S = \mathbb{L}$, $f(x, y) = x \circ y$ and considering the ascending ω-chain $x_n = 1 - \frac{1}{n}$ (for $n \geq 2$):

$$H(\textstyle\bigsqcup_n x_n) = H(1) = \mathsf{gfp}(h_1) = \mathsf{gfp}(y \mapsto y \circ 1) = 1$$

$$\bigsqcup_n H(x_n) = \bigsqcup_n \mathsf{gfp}(h_{x_n}) = \bigsqcup_n \mathsf{gfp}(y \mapsto y \circ x_n) = 0$$

This provides a negative answer for H. $f(x, y)$ represents the formula ϑ with two relation symbols, one corresponding to the first argument x and the other corresponding to y.

Example 5.26. We can further give a negative answer for G by considering the semiring $\mathbb{S}^\infty[a]$ of univariate absorptive polynomials ($\mathbb{S}^\infty[X]$ is introduced in section 6.4).

Consider the function f and the descending ω-chain $(y_n)_{n<\omega}$ defined by

$$f(x, y) = \max\left(y, \tfrac{x}{a}\right)$$

$$y_n = a^n, \quad \text{for } n < \omega$$

The diagram on the right depicts the steps of the fixed-point iteration $g^k_{y_n}(0)$ depending on n.

$$\begin{array}{c|cccccccc|c}
 & & & & n \longrightarrow & & & & \\
\hline
 & 1 & a^1 & a^2 & a^3 & a^4 & a^5 & \cdots & a^\infty \\
 & 1 & 1 & a^1 & a^2 & a^3 & a^4 & \cdots & a^\infty \\
 & 1 & 1 & 1 & a^1 & a^2 & a^3 & \cdots & a^\infty \\
k \downarrow & 1 & 1 & 1 & 1 & a^1 & a^2 & \cdots & a^\infty \\
 & 1 & 1 & 1 & 1 & 1 & a^1 & \cdots & a^\infty \\
 & 1 & 1 & 1 & 1 & 1 & 1 & \cdots & a^\infty \\
 & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & & \vdots \\
\hline
 & 1 & 1 & 1 & 1 & 1 & 1 & \cdots & \lightning
\end{array}$$

Here, $\frac{x}{a}$ stands for division by a. That is, $\frac{a^{n+1}}{a} = a^n$, $\frac{a^\infty}{a} = a^\infty$ and we additionally set $\frac{1}{a} = 1$. Note that univariate absorptive polynomials consist of only one monomial (due to absorption). For the empty polynomial 0, we additionally set $\frac{0}{a} = 0$.

Then f is inf-continuous in y (for a fixed x). Due to the ascending chain condition of $\mathbb{S}^\infty[a]$, it is further easy to see that f is sup-continuous in x. As the picture indicates, G is not inf-continuous (we reason about $\mathsf{lfp}(g_y)$ via the fixed-point iteration $g^k_y(\bot)$):

$$G(\textstyle\bigsqcap_n y_n) = G(a^\infty) = \mathsf{lfp}(x \mapsto f(x, a^\infty)) = \bigsqcup\{\bot, a^\infty, a^\infty, a^\infty, \dots\} = a^\infty$$

$$\bigsqcap_n G(y_n) = \bigsqcap_n \mathsf{lfp}(x \mapsto f(x, a^n)) = \bigsqcap_n \bigsqcup\{\bot, a^n, a^{n-1}, \dots, a, 1\} = \bigsqcap_n 1 = 1$$

Unlike the first example, f does not correspond to a formula, but the chain $(y_n)_{n<\omega}$ is the result of repeated multiplication with a and thus occurs as a fixed-point iteration of a formula in $\mathbb{S}^\infty[a]$.

These counterexamples to the generalized question make clear that our open question has to be attacked by either focusing on certain semirings or by incorporating properties of logical formulae (and the corresponding fixed-point iterations) into the argument. Our answer for absorptive semirings in section 6.4 indeed relies on both absorption and, indirectly, on the structure of formulae by using the fundamental property.

6 Provenance Semirings

This chapter discusses semirings of particular interest for provenance analysis, mainly polynomial semirings in which we can use variables to track the influence of certain literals. From an algebraic point of view, polynomial semirings are interesting due to their universal property. That is, they capture computations in a large class of semirings (via homomorphisms induced by polynomial evaluation) while at the same time having a reasonable representation. Our main focus is on the chain-positive semiring $\mathbb{S}^\infty[X]$ of (generalized) absorptive polynomials which we show to be universal among all absorptive continuous semirings (in terms of cpo-semiring homomorphisms). This makes absorptive polynomials a promising choice for provenance analysis and also leads to an answer of the earlier stated open question in the context of absorption. We close this chapter by extending the idea of dual-indeterminate polynomials from [GT17a] to polynomial cpo semirings in order to improve the interpretation of negation.

As a reference point, consider the situation for FO as shown in figure 2 below. We have no requirements on the semiring (other than being positive for some results) and can thus work with polynomials $\mathbb{N}[X]$. This is the (commutative) semiring freely generated by the set X, so we can compute provenance information in a very general way. As indicated in the figure, we can specialize this via homomorphisms to the universal idempotent semiring $\mathbb{B}[X]$ and further to semirings which are additionally absorptive, multiplicative idempotent or both (these semirings are introduced in this chapter).

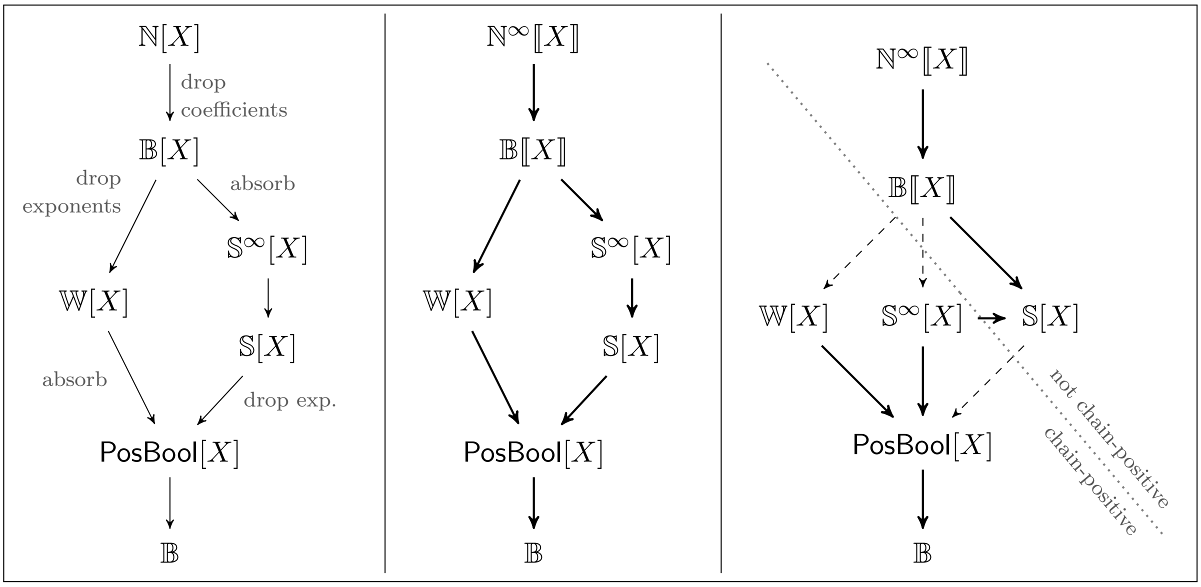

Figure 2: *Polynomial semirings for FO (left), posLFP (center) and LFP (right) discussed in this chapter. Arrows indicate canonical surjective homomorphisms. Thick arrows are ω-continuous (center) or cpo homomorphisms (right) while dashed ones are not continuous.*

For fixed-point logics, we need cpo semirings and thus have to work with formal power series instead of polynomials. If we only consider least fixed-points, then the situation is comparable to FO as the canonical homomorphism happen to be ω-continuous and thus preserve least fixed-points (see center diagram in figure 2). This also means that formal power series $\mathbb{N}^\infty[\![X]\!]$ are the most general ω-continuous semiring.

For greatest fixed points, the situation is more complex (diagram on the right). We can use the same semirings as for posLFP, but we need cpo-semiring homomorphisms to preserve all fixed points (see e.g. example 5.16). This leads to two kinds of polynomial semirings between which the canonical homomorphisms are not continuous: The chain-positive (but less general) ones, and the semirings which are not chain-positive (mainly formal power series). Without chain-positivitiy, these semirings do not preserve truth (in particular, $\dagger_S$ is not continuous) and are thus not favorable for general provenance computation. We therefore have to make additional restrictions such as absorption or leaving out exponents to obtain universal semirings for the analysis of full LFP.

6.1 Counting

The first semiring we consider is $\mathbb{N}^\infty$ as analogue of $\mathbb{N}$ for FO. The addition of ∞ makes $\mathbb{N}^\infty$ a lattice semiring suitable for provenance computation. In FO, $\mathbb{N}$ can be used to count the number of proofs and we can make similar observations for posLFP.

Example 6.1. In the introductory example 3.5 we have considered the reachability formula $\varphi_{\text{path}}(u,v) = [\mathsf{lfp}\, R\, x.\ x = v \lor \exists y(Exy \land R\, y)](u)$. We have seen that provenance computation yields the number of paths from u to v (or, equivalently, the number of proofs of $\varphi_{\text{path}}(u,v)$). The two graphs we considered are shown below.

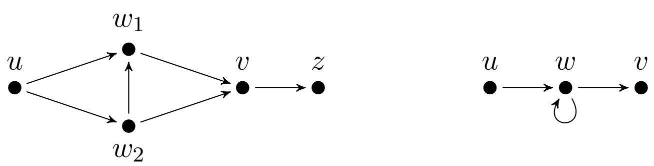

Writing valuations π as tuples $(\pi(u), \pi(w_1), \pi(w_2), \pi(v), \pi(z))$ or $(\pi(u), \pi(v), \pi(w))$, we have seen that the fixed-point iteration for the left graph terminates at $\pi_4 = (3, 1, 2, 1, 0)$. On the second graph, the iteration does not terminate after finitely many steps, as we have $\pi_n = (n-2, n-1, 1)$. The least fixed-point in $\mathbb{N}^\infty$ is thus $\pi_\omega = (\infty, \infty, 1)$.

In both cases, each entry is equal to the number of paths (possibly empty and allowing repetitions) from the corresponding node to v. This example suggests that, for lfp-formulae, we can use semiring semantics in $\mathbb{N}^\infty$ to count proofs (for now, we rely on an intuitive notion of proofs, we later make this precise).

For greatest fixed points, the provenance information in $\mathbb{N}^\infty$ is less useful as witnessed by the example below. One reason is that both operations are increasing (that is, $a+b \geq a, b$ and $a \cdot b \geq a, b$ for $a, b \neq 0$). As the fixed-point iteration starts with ∞, it therefore remains at ∞ (unless the value of the previous iteration is multiplied by 0).

Example 6.2. Recall the formula $\varphi(v) \coloneqq \varphi_{\text{infpath}}(v) = [\mathsf{gfp}\, R\, x.\, \exists y(Exy \wedge Ry)](v)$ from example 5.6. Consider the interpretation ℓ induced by the following graph.

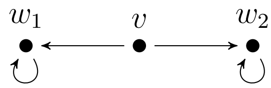

There are two different infinite paths from v and thus two different proofs of $\varphi_1(v)$, so we might expect $[\![\varphi(v)]\!]_\ell = 2$. Instead, we get $[\![\varphi(v)]\!]_\ell = \infty$. To see this, consider the fixed-point iteration (where we write valuations π as $(\pi(w_1), \pi(v), \pi(w_2))$):

$$\pi_0 = (\infty, \infty, \infty), \qquad \pi_1 = (1 \cdot \infty, 1 \cdot \infty + 1 \cdot \infty, 1 \cdot \infty) = (\infty, \infty, \infty)$$

The fixed-point iteration starts at ∞ and immediately terminates, as every node has at least one outgoing edge. The same happens for alternating formulae such as

$$\varphi_2(v) = [\mathsf{gfp}\, X\, x.\, [\mathsf{lfp}\, Y\, x.\, \exists y(Exy \wedge ((Py \wedge Xy) \vee Yy)](x)](v)$$

Here, $\varphi_2(v)$ says that there is a path from v which visits infinitely many nodes labeled with P. If we extend ℓ by $\ell(Pw_1) = 1$ and $\ell(Pw_2) = \ell(Pv) = 0$, then we only consider w_1 to be labeled with P and hence there is only one path witnessing the truth of $\varphi_2(v)$. However, we again get $[\![\varphi_2(v)]\!]_\ell = \infty$.

These examples show that provenance computations in $\mathbb{N}^\infty$ are not very informative for greatest fixed points, so $\mathbb{N}^\infty$ cannot be used to count proofs in LFP.

This example motivates the search for semirings in which multiplication is not increasing. One possibility are polynomial semirings, such as formal power series, in which xy is incomparable to x and y. The alternative are absorptive semirings in which multiplication is always decreasing.

6.2 Formal Power Series

Polynomials with coefficients in $\mathbb{N}$ are not chain-complete for two reasons: They lack a maximal coefficient and a supremum of the chain x, $x + x^2$, $x + x^2 + x^3$, ... To address the first issue, we use $\mathbb{N}^\infty$ instead of $\mathbb{N}$ for coefficients. For the second, we use formal power series which are like polynomials but admit infinitely many monomials.

As we introduce a number of polynomial semirings below, let us first fix notation. The set X of variables is always finite. In explanations and examples, we often use the symbols x, y, z as well as a, b, c for variables. Regardless of the formal definition, we use the common notation for monomials, e.g. x^2y^3 and write polynomials or formal power series as (possibly infinite) sums of monomials.

Definition 6.3. A monomial m in the variables X with exponents from a semiring E is a function $m : X \to E$. The set of monomials with exponents in E is denoted by $\mathcal{M}_E$. Monomial multiplication is defined as usual by adding exponents: $(m_1 \cdot m_2)(x) = m_1(x) + m_2(x)$. We write $\mathcal{M}$ and $\mathcal{M}_\infty$ for the sets $\mathcal{M}_{\mathbb{N}}$ and $\mathcal{M}_{\mathbb{N}^\infty}$, respectively, and 1 for the empty monomial with $x \mapsto 0$ (for all $x \in X$).

Definition 6.4. Let S be a semiring. A formal power series with coefficients in S is a function $f : \mathcal{M} \to S$. The set of formal power series over S forms the semiring $S[\![X]\!]$ with pointwise addition and the usual polynomial multiplication:

$$\big(f \cdot g\big)(m) \;=\; \sum_{m = m_1 m_2} f(m_1) \cdot g(m_2)$$

Following this definition, we can view $S[\![X]\!]$ as a function semiring (apart from its multiplication, which is independent of the order anyway) and it follows that $S[\![X]\!]$ is a lattice semiring whenever the same holds for the coefficient semiring S. Additionally, $S[\![X]\!]$ is continuous whenever S is continuous (see [DK09]; they only consider continuity with respect to suprema, but the argument for infima is analogous).

Proposition 6.5. *If S is continuous or a lattice semiring, then so is $S[\![X]\!]$.*

We are interested in the formal power series $\mathbb{N}^\infty[\![X]\!]$ and the idempotent version $\mathbb{B}[\![X]\!]$ that results from $\mathbb{N}^\infty[\![X]\!]$ by dropping coefficients. The above proposition implies that both are ω-continuous and thus suitable for posLFP. Indeed, $\mathbb{N}^\infty[\![X]\!]$ has already been proposed for provenance analysis of positive fixed-point logics such as datalog [GKT07] and positive fragments of PDL, CTL [Mrk18] and LFP [GT19].

Example 6.6. As an example for posLFP, we again consider graph reachability via $\varphi_{\text{path}}(u, v) = [\mathsf{lfp}\, R x.\, \underbrace{x = v \lor \exists y(Exy \land Ry)}_{\vartheta(x)}](u)$ on the following two graphs:

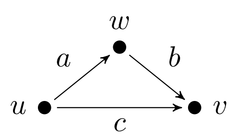

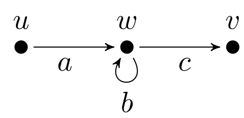

For the graph on the left, we get the following iteration (each column is one valuation):

$$\begin{array}{rcccccccc} \pi(u): & 0 & & 0 & & c & & c+ab \\ \pi(w): & 0 & \xmapsto{F^\vartheta} & 0 & \xmapsto{F^\vartheta} & b & \xmapsto{F^\vartheta} & b \\ \pi(v): & 0 & & 1 & & 1 & & 1 \end{array}$$

which quickly terminates and yields $[\![\varphi_{\text{path}}(u,v)]\!]_\ell = c + ab$ as expected. For the second graph, we obtain an infinite iteration,

$$\begin{array}{rccccccccccc} \pi(u): & 0 & & 0 & & 0 & & ac & & ac+abc \\ \pi(w): & 0 & \mapsto & 0 & \mapsto & c & \mapsto & c+bc & \mapsto & c+bc+b^2c & \cdots \\ \pi(v): & 0 & & 1 & & 1 & & 1 & & 1 \end{array}$$

which leads to the infinite power series $[\![\varphi_{\text{path}}(u,v)]\!]_\ell = \sum_{n<\omega} ab^n c$. Each monomial $ab^n c$ correspond to one possible path that cycles n times at node w.

For both examples, we can consider the variable assignment $a, b, c \mapsto 1$ which induces an ω-continuous homomorphism $h : \mathbb{N}^\infty[\![X]\!] \to \mathbb{N}^\infty$ corresponding to polynomial evaluation. Applying h to the result in $\mathbb{N}^\infty[\![X]\!]$ then yields the same value as the computation in $\mathbb{N}^\infty$ (when we interpret all edges by 1). For the first example, we indeed have $h(ab+c) = 1+1 = 2$ and for the second graph, we have $h(\sum_{n<\omega} ab^n c) = \sum_{n<\omega} h(ab^n c) = \sum_{n<\omega} 1 = \infty$.

These examples show that $\mathbb{N}^\infty[\![X]\!]$ is well-suited for provenance analysis of posLFP.

As seen for the homomorphism h above, another reason why formal power series $\mathbb{N}^\infty[\![X]\!]$ are interesting is that they inherit the universality of $\mathbb{N}[X]$, as stated in [GKT07].

Proposition 6.7. *Let T be an ω-continuous semiring and let $v : X \to T$ be a variable assignment. Then there is a uniquely defined ω-continuous homomorphism $h : \mathbb{N}^\infty[\![X]\!] \to T$ that extends v (i.e., $h(x) = v(x)$ for all $x \in X$).*

The homomorphism can be defined as $h : f \mapsto \sum_{m \in \mathcal{M}} f(m) \cdot h(m)$, where the factors $f(m)$ are interpreted as repeated addition (possibly infinite) in the target semiring T and the countable sum can be defined via suprema in ω-continuous semirings (see [GKT07]). As h is ω-continuous, it preserves semiring semantics for posLFP.

Formal power series are, however, not suitable as universal semirings for the interpretation of full LFP. The reason is that they are not chain-positive, so the truth projection † is not a continuous homomorphism and formal power series are thus not compatible with standard semantics. The reason why the ω-continuous homomorphism h defined above is not continuous is that infima do not commute with the countable sum over all monomials.

Another potential issue is that formal power series may be infinite and thus difficult to represent (although [GKT07] shows that finiteness and individual coefficients can be computed for datalog provenance).

Example 6.8. Consider φ_{infpath} on one of the graphs from the previous example:

$$\varphi(u) = \varphi_{\text{infpath}}(u) = [\mathsf{gfp}\, R\, x.\ \exists y(Exy \wedge Ry)](u)$$

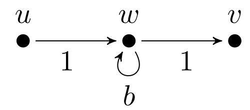

To simplify the presentation, we only track the self-loop at w and consider $\mathbb{B}[\![X]\!]$ with $X = \{b\}$ instead of $\mathbb{N}^\infty[\![X]\!]$. The following argument also applies to the more general setting. The top element of $\mathbb{B}[\![b]\!]$ is the infinite power series $1 + b + b^2 + b^3 + \dots$ and the fixed-point iteration yields the following values for the node w:

$$\begin{aligned}
\pi_1(w) &= 1 + b + b^2 + b^3 + \cdots + b^n + \dots \\
\pi_2(w) &= \quad\ b + b^2 + b^3 + \cdots + b^n + \dots \\
\pi_3(w) &= \qquad\quad b^2 + b^3 + \cdots + b^n + \dots \\
\\
\pi_{n+1}(w) &= \qquad\qquad\qquad\qquad\ \ b^n + \dots
\end{aligned}$$

Consider any monomial m with the exponent $m(b) = k$. Then the last occurrence of m is in $\pi_{k+1}(w)$ and the infimum is thus 0. Hence also $[\![\varphi(u)]\!]_\ell = [\![\varphi(w)]\!]_\ell = 0$.

This is in conflict with standard semantics. The assignment $b \mapsto \top$ lifts to a homomorphism $h : \mathbb{B}[\![b]\!] \to \mathbb{B}$ such that $h \circ \ell$ is model-defining for the above graph. As the graph has an infinite path, we get $[\![\varphi(u)]\!]_{h\circ\ell} = \top$ although $h([\![\varphi(u)]\!]_\ell) = h(0) = \bot$. This shows in particular that the homomorphism h (which is simply the truth projection $\dagger$) is not continuous, even though it is defined in a natural way by a variable assignment.

In $\mathbb{N}^\infty[\![X]\!]$, we obtain $[\![\varphi(u)]\!]_\ell = 0$ by the same argument. This leads to a similar conflict with $\mathbb{N}^\infty$: The assignment $b \mapsto 1$ induces a homomorphism h, but we have $[\![\varphi(u)]\!]_{h\circ\ell} = \infty$ in $\mathbb{N}^\infty$ and $[\![\varphi(u)]\!]_\ell = 0$ in $\mathbb{N}^\infty[\![X]\!]$.

These considerations show that formal power series are no longer universal in the presence of greatest fixed points. All of these issues are connected to the missing chain-positivity which is here exhibited by repeatedly applying multiplication.

6.3 Why-Semiring

In previous examples, we have seen several issues with greatest fixed points related to chain-positivity or caused by an increasing multiplication. In the reachability example for $\mathbb{N}^\infty$, we have further argued that we are often mostly interested in shortest paths. We now present two possible solutions: With the why-semiring, we only consider which facts (literals) are needed for a proof, but we do not count how often they are used. So it does not matter how often we loop through the cycle, we only remember whether we used the cycle at all. The second solution is to use the absorptive polynomials presented in the next section which only represent shortest proofs.

The why-semiring is motivated by why-provenance for databases. In [Gre11], it has been defined as $\mathsf{Why}(X) = (\mathcal{P}(\mathcal{P}(X)), \cup, \Cup, \emptyset, \{\emptyset\})$ where $P_1 \Cup P_2 = \{m_1 \cup m_2 \mid m_1 \in P_1, m_2 \in P_2\}$ is pairwise union. We give an equivalent definition in terms of polynomials. A slightly more general semiring considered in [Gre11] is $\mathsf{Trio}(X)$ which additionally allows coefficients from $\mathbb{N}$. A chain-complete version could use $\mathbb{N}^\infty$ instead of $\mathbb{N}$, but would exhibit similar behavior and issues as $\mathbb{N}^\infty$ and $\mathbb{W}[X]$, so we do not consider it here.

Definition 6.9. The why-semiring $\mathbb{W}[X]$ is the set of all functions $P : \mathcal{M}_\mathbb{B} \to \mathbb{B}$ which we interpret as polynomials, together with standard polynomial addition and multiplication.

Following the definition of $\mathbb{W}[X]$ via $\mathcal{P}(\mathcal{P}(X))$, we view monomials from $\mathcal{M}_\mathbb{B}$ as subsets of X and functions P as subsets of $\mathcal{M}_\mathbb{B}$ for convenience.

$\mathbb{W}[X]$ results from $\mathbb{B}[X]$ or $\mathbb{B}[\![X]\!]$ by dropping exponents, as $x^2 = x$ for all $x \in X$. However, $\mathbb{W}[X]$ is *not* multiplicative idempotent in general,[7] as $(x+y)^2 = x + xy + y$. An important property of the why-semiring is its finiteness which makes it chain-complete and continuous (and thus, being idempotent, a lattice semiring).

Proposition 6.10. *The why-semiring $\mathbb{W}[X]$ is a continuous lattice semiring.*

Moreover, all chains are finite and thus have maximal (or minimal) elements. As homomorphisms are order-preserving, this implies that every homomorphism is continuous and we can formulate the following universality result. It further implies that $[\![\varphi]\!]_\ell$ is continuous in ℓ, as there are only finitely many $\mathbb{W}[X]$-interpretations.

Proposition 6.11. *Let S be an idempotent cpo semiring and let $f : X \to S$ be a variable assignment with $f(x)^2 = f(x)$ for all $x \in X$. Then f extends to a unique homomorphism $h : \mathbb{W}[X] \to T$. This homomorphism is continuous.*

In particular, this applies to all semirings which are both idempotent and multiplicatively idempotent. Note that this result does not suffice to preserve fixed-points. For this, we require h to be a cpo homomorphism, so we additionally need $h(\top) = \top$. In $\mathbb{W}[X]$, the neutral element 1 is the empty monomial while $\top$ is the set of all monomials, so $h(1) = 1$ does not imply $h(\top) = \top$. And we can indeed define homomorphisms which do not preserve fixed-points.

[7] It should be noted that this observation contradicts the claim in [GT17b] that $\mathsf{Why}(X)$ results from $\mathbb{B}[X]$ by making multiplication idempotent.

Example 6.12. The mapping $x \mapsto 1$ (for all $x \in X$) induces the trivial homomorphism $h : \mathbb{W}[X] \to \mathbb{W}[X]$ with $h(p) = 1$ for all $p \neq 0$ (we can view the image of h as an embedding of $\mathbb{B}$ into $\mathbb{W}[X]$; this is possible for every idempotent semiring).

Let $X = \{x\}$ and consider $\varphi_{\text{infpath}}(v)$ on a graph consisting of a single node v with a self-loop (as in example 5.6). For the $W[X]$-interpretation with $\ell(Evv) = x$, we then get the iteration $\pi_0(v) = \top = 1 + x$, $\pi_1(v) = x \cdot (1 + x) = x$, $\pi_2(v) = x \cdot x = x$ and hence $[\![\varphi]\!]_\ell = x$. Applying the homomorphism h, we obtain $h([\![\varphi]\!]_\ell) = h(x) = 1$. If we instead apply h first, we get $(h \circ l)(Evv) = 1$ and then $[\![\varphi]\!]_{h\circ\ell} = \top = 1 + x$.

The why-semiring is also interesting in its own right. If we assign variables to the facts (literals) we want to track, we may think of a monomial $[\![\varphi]\!]_\ell$ as the set of tracked facts occurring in a proof. For posLFP this indeed works, so we learn *why* the formula holds (but not *how* exactly the proof looks). For greatest fixed-points, the situation is, yet again, not so clear. The reason can be found in the fixed-point iteration: We start with the set of all monomials and then rule those out whose variables (or rather, the associated facts) do not suffice to prove φ. We do however keep monomials with superfluous variables even though they might not correspond to proofs, as the following example shows.

Example 6.13. Let us first consider the positive formula φ_{path} on the following graph:

$$\varphi_{\text{path}}(u, v) = [\mathsf{lfp}\, R\,x.\ x = v \vee \exists y(Exy \wedge R\,y)](u)$$

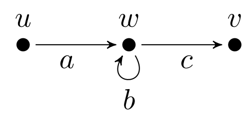

Consider $X = \{a, b, c\}$ and the model-defining $\mathbb{W}[X]$-interpretation indicated in the picture (e.g., $\ell(Euw) = a$, $\ell(\neg Euw) = 0$). We then get the fixed-point iteration

$$\begin{array}{lllllllll}
\pi(u): & 0 & & 0 & & 0 & & ac & & ac + abc \\
\pi(w): & 0 & \mapsto & 0 & \mapsto & c & \mapsto & c + bc & \mapsto & c + bc \\
\pi(v): & 0 & & 1 & & 1 & & 1 & & 1
\end{array}$$

and obtain the overall value $[\![\varphi_{\text{path}}(u, v)]\!]_\ell = ac + abc$. The monomial ac corresponds to the path $u \to w \to v$ witnessing the truth of $\varphi_{\text{path}}(u, v)$, while the monomial abc corresponds to all paths which additionally cycle (finitely often) at w. So we learn from the provenance computation that we can prove the formula by using the edges $\{a, c\}$ or by using the edges $\{a, b, c\}$ as we would expect.

Now consider $\varphi_{\text{infpath}}(u) = [\mathsf{gfp}\, R\,x.\ \exists y(Exy \wedge Ry)](u)$ on the same graph. There is only one infinite path using the edges $\{a, b\}$. However, the fixed-point iteration yields

$$\begin{array}{lllllll}
\pi(u): & \top & & a\top & & ab\top + ac\top & & ab\top \\
\pi(w): & \top & \mapsto & b\top + c\top & \mapsto & b\top & \mapsto & b\top \\
\pi(v): & \top & & 0 & & 0 & & 0
\end{array}$$

where $\top$ is the set of all monomials. Hence $[\![\varphi_{\text{infpath}}(u)]\!]_\ell = ab\top = ab + abc$. The monomial ab indeed corresponds to the infinite path, but there is no infinite path for the monomial abc. So for greatest fixed points, not all monomials correspond to proofs.

The reason is that multiplication in the gfp-iteration results in expressions of the form $m\top$ where m corresponds to the actual proof of φ. The monomials in $m\top$ all contain enough variables for a proof of φ, but also contain superfluous variables (such as abc). The actual provenance information is thus carried by the shortest monomials m.

For least fixed-points, we only add monomials during the iteration that actually correspond to proofs, so such problems cannot occur. For full LFP however, the example makes clear that $\mathbb{W}[X]$ is not the semiring we are looking for. What we need is a way to avoid the unwanted longer monomials – this is achieved by absorption.

6.4 Absorptive Polynomials

The alternative to the why-semiring is to introduce absorption among monomials. Then only the shortest monomials remain, so $[\![\varphi]\!]_\ell$ only contains information about the minimal witnesses for the truth of φ. Absorptive polynomials have been introduced as the semiring $\mathsf{Sorp}(X)$ in [DMRT14] where it was defined as a quotient of $\mathbb{N}[X]$ (modulo absorption). We refer to this semiring as $\mathbb{S}[X]$. More important for us is the definition of *generalized* absorptive polynomials $\mathbb{S}^\infty[X]$ in [GT19] and [Mrk18] which complete $\mathbb{S}[X]$ by allowing the exponent ∞, thereby achieving chain-completeness. As we are only concerned with $\mathbb{S}^\infty[X]$ in this work, we use the term *absorptive polynomials* for the semiring $\mathbb{S}^\infty[X]$.

A central property of absorptive polynomials is their universality. The definition as quotient of $\mathbb{N}[X]$ modulo absorption in [DMRT14] means that $\mathsf{Sorp}(X)$ is the most general absorptive semiring. For provenance of LFP, however, we need universality in terms of cpo-semiring homomorphisms to preserve fixed-points. We show that in this respect, $\mathbb{S}^\infty[X]$ is the most general absorptive continuous semiring.

We have already seen another motivation to consider absorptive semirings: In $\mathbb{N}^\infty$, both operations are increasing. This leads to an asymmetry between least and greatest fixed points and causes the latter to contain less information. If we instead require that $a, b \geq ab$ (for all elements a, b), then this already implies absorption (see proposition 4.25). In other words, absorption goes hand in hand with a higher degree of symmetry which in turn leads to better results for greatest fixed points.

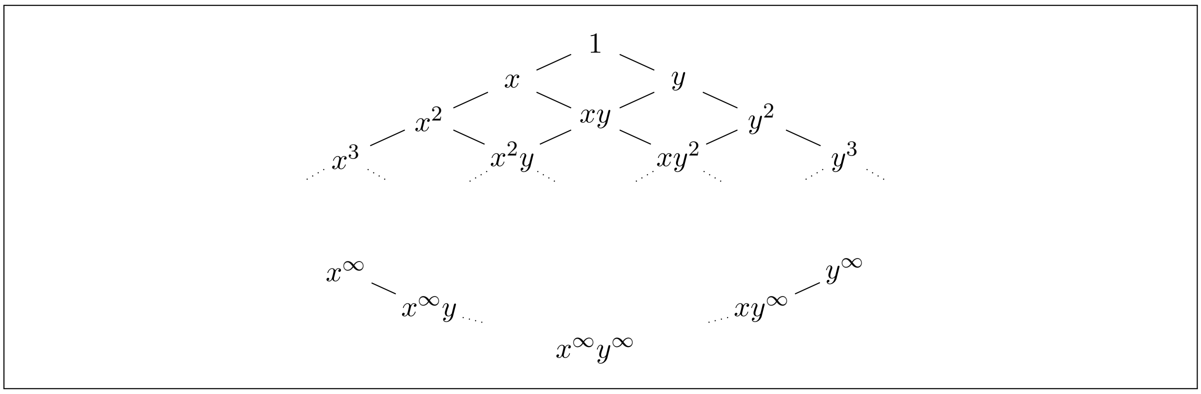

Figure 3: *Hasse diagram of the lattice of monomials over* $\{x, y\}$ *ordered by absorption.*

Definition

We follow the definition by Grädel and Tannen [GT19]. The starting point is the absorption order on monomials which intuitively means that *shorter* monomials (with smaller exponents) absorb longer ones. Absorptive polynomials are then polynomials (without coefficients) in which we only keep *shortest* monomials.

Definition 6.14. The *absorption order* on $\mathcal{M}_\infty$ is the partial order defined by

$$m_1 \geq m_2 \quad \iff \quad m_1(x) \leq m_2(x) \text{ for all } x \in X$$

If $m_1 \geq m_2$, we say that m_1 *absorbs* the monomial m_2.

Note that we have $m_1 \geq m_2$ if the exponents in m_1 are *smaller* than m_2, so m_1 is *shorter*. For example, ab^2 absorbs a^3b^2 and ab^∞, but not a^2b. The empty monomial 1 absorbs all other monomials. As the absorption order is defined variable-wise and $\mathbb{N}^\infty$ is a complete lattice, the monomials $\mathcal{M}_\infty$ form a complete lattice under absorption (see figure 3).

Remark: The monomials $\mathcal{M}_\infty$ can also be described as the function semiring $(\mathbb{N}^\infty)^X$. Following the above definition, we see that the absorption order is the *inverse* of the natural order on $(\mathbb{N}^\infty)^X$. This is the reason why the *infimum* of a family of monomials $(m_i)_{i \in I}$ can be expressed variable-wise by taking the *supremum* of the exponents $m_i(x)$ for each variable x. For example, the chain $(m_i)_{i<\omega} = x, xy, xy^2, \dots, xy^i, \dots$ has the infimum xy^∞ where $1 = \bigsqcup_i m_i(x)$ and $\infty = \bigsqcup_i m_i(y) = \bigsqcup_i i$.

Absorptive polynomials are *antichains* of monomials, as we only keep the absorption-maximal monomials after each operation, e.g. $ab^2 + ab = ab$. While $\mathbb{S}^\infty[X]$ is infinite, it is a simple yet crucial observation that antichains of monomials and ascending chains are always finite [GT19]. One way to see this is by Dickson's lemma as formulated in [FFSS11] which states that every infinite sequence $(\mathbf{a}_i)_{i<\omega}$ of elements from $\mathbb{N}^k$ (for some finite k and using the product order for k-tuples) contains an ascending pair $\mathbf{a}_i \leq \mathbf{a}_j$ with $i < j$. Following the proof of Dickson's lemma, it is easy to see that the same holds for $(\mathbb{N}^\infty)^k$. This translates to our setting when we view $\mathcal{M}_\infty$ as $(\mathbb{N}^\infty)^{|X|}$ and note that absorption order is the inverse of the standard product order.

Proposition 6.15. *$\mathcal{M}_\infty$ is a complete lattice under the absorption order. Moreover, all antichains and ascending chains in $\mathcal{M}_\infty$ are finite.*

Addition and multiplication are defined via the standard operations on polynomials, except that we afterwards apply absorption. To this end, we write $\mathsf{maximals}(M)$ for the set of absorption-maximal monomials in the $M \subseteq \mathcal{M}_\infty$, which always yields an antichain.

Definition 6.16. An *absorptive polynomial* (with variables in X and exponents in $\mathbb{N}^\infty$) is an antichain in $\mathcal{M}_\infty$. The set of absorptive polynomials forms the semiring $\mathbb{S}^\infty[X]$ with operations defined by

$$\begin{aligned} P + Q &= \mathsf{maximals}(P \cup Q) \\ P \cdot Q &= \mathsf{maximals}(\{m_1 \cdot m_2 \mid m_1 \in P, m_2 \in Q\}) \end{aligned}$$

The semiring $\mathbb{S}[X]$ is defined in the same way based on $\mathcal{M}$ instead of $\mathcal{M}_\infty$.

One can easily verify that this defines a semiring. We write an absorptive polynomial as sum of its monomials, e.g. $P = a^2b + b^\infty$. The neutral elements are the empty antichain, which we write as 0, and the antichain consisting only of the monomial 1.

The semiring $\mathbb{S}[X]$ of absorptive polynomials with exponents from $\mathbb{N}$ instead of $\mathbb{N}^\infty$ is not chain-positive, as the infimum of $x, x^2, x^3, x^4, \dots$ is 0. We can thus view the additional monomials x^∞ as a completion of $\mathbb{S}[X]$ to achieve chain-positivity. The canonical homomorphism from formal power series $\mathbb{N}^\infty[\![X]\!]$ or $\mathbb{B}[\![X]\!]$ to $\mathbb{S}^\infty[X]$ is given by $f \mapsto \mathsf{maximals}(\{m \mid f(m) \neq 0\})$ and is not continuous:

Example 6.17. Recall the problematic example 6.8 for formal power series $\mathbb{B}[\![b]\!]$. In $\mathbb{S}^\infty[b]$, the iteration has a similar but simpler form (values on the right):

$$\begin{aligned} \pi_1(w) &= 1 + b + b^2 + b^3 + \dots + b^n + \dots &&\mapsto \quad 1 \\ \pi_2(w) &= \phantom{1 +{}} b + b^2 + b^3 + \dots + b^n + \dots &&\mapsto \quad b \\ \pi_3(w) &= \phantom{1 + b +{}} b^2 + b^3 + \dots + b^n + \dots &&\mapsto \quad b^2 \\ \pi_{n+1}(w) &= b^n + \dots &&\mapsto \quad b^n \\ \pi_\omega(w) &= 0 &&\not\mapsto \quad b^\infty \end{aligned}$$

Due to the completion of $\mathbb{S}^\infty[X]$, we now obtain the infimum b^∞ instead of 0.

Properties

A key property is the finiteness of absorptive polynomials due to proposition 6.15. This further implies that $\mathbb{S}^\infty[X]$ is *countable* as a set of finite subsets of the countable set $\mathcal{M}_\infty$. The natural order can be described via the absorption order on monomials:

Proposition 6.18. *For $P, Q \in \mathbb{S}^\infty[X]$,*

$$P \leq Q \quad \iff \quad \text{for every } m_1 \in P \text{ there is } m_2 \in Q \text{ with } m_1 \leq m_2.$$

Proof. If $P \leq Q$, then $P + P' = Q$ for some $P' \in \mathbb{S}^\infty[X]$ and thus $Q = \mathsf{maximals}(P \cup P')$. As $P \cup P'$ is finite, for every $m \in P$ there is $m' \in \mathsf{maximals}(P \cup P') = Q$ with $m \leq m'$.

For the other direction, consider $P + Q$. As every $m_1 \in P$ is absorbed by some $m_2 \in Q$, we have $P + Q = \mathsf{maximals}(P \cup Q) = \mathsf{maximals}(Q) = Q$ and thus $P \leq Q$. □

Most importantly, $\mathbb{S}^\infty[X]$ is an absorptive lattice semiring [GT19] and thus suitable for the interpretation of fixed-point logics. Absorption is not surprising due to our construction. For the lattice property, we need an additional observation.

Lemma 6.19. *Let $M \subseteq \mathcal{M}_\infty$. Then for every monomial $m \in M$ there exists $m' \in \mathsf{maximals}(M)$ with $m \leq m'$.*

Proof sketch. Towards a contradiction, assume that that there is no such m'. Then m is not maximal, so there is an $m_1 \geq m$. By repeating this argument, we obtain an infinite ascending chain $m \leq m_1 \leq m_2 \leq \dots$ which contradicts proposition 6.15. □

Theorem 6.20. *$\mathbb{S}^\infty[X]$ is an absorptive lattice semiring.*

Proof. First note that $\mathbb{S}^\infty[X]$ is naturally ordered as a consequence of proposition 6.18. For absorption, we have to show $P + PQ = P$. Every monomial in PQ is of the form $m_1 \cdot m_2$ with $m_1 \in P$ and $m_2 \in Q$. Such a monomial is absorbed by $m_1 \in P$ and thus $P + PQ \leq P$. The other direction holds by natural order.

To show that $\mathbb{S}^\infty[X]$ is a complete lattice, we show that suprema of arbitrary sets exist:

$$\bigsqcup S = \mathsf{maximals}(\bigcup S), \qquad \text{where } S \subseteq \mathbb{S}^\infty[X].$$

Let $P \in S$. Then $P \subseteq \bigcup S$ and thus $P \leq \mathsf{maximals}(\bigcup S)$ by the previous lemma. Hence $\mathsf{maximals}(\bigcup S)$ is an upper bound for S.

To see that it is the least upper bound, let Q be any upper bound for S, so $Q \geq P$ for all $P \in S$. For each $m \in \mathsf{maximals}(\bigcup S)$ there is a $P \in S$ with $m \in P$ and hence $m \leq P \leq Q$. It follows that $\mathsf{maximals}(\bigcup S) \leq Q$. $\square$

The above proof provides a simple characterization for suprema of arbitrary sets. For chains, the situation is even simpler as $\mathbb{S}^\infty[X]$ fulfills the *ascending chain condition* (ACC). That is, each ascending chain reaches its supremum after finitely many steps.

Proposition 6.21. *Let $(P_i)_{i<\omega}$ be an ascending ω-chain in $\mathbb{S}^\infty[X]$. Then there is a $k < \omega$ such that $P_j = P_k = \bigsqcup_i P_i$ for all $j \geq k$.*

Proof. The supremum has the finite form $\bigsqcup_i P_i = m_1 + \cdots + m_n$ for some $n < \omega$. As $\bigsqcup_i P_i = \mathsf{maximals}(\bigcup_i P_i)$, for every $1 \leq j \leq n$ there is an index i_j such that $m_j \in P_{i_j}$. Then $P_k = \bigsqcup_i P_i$ with $k = \max\{i_j \mid 1 \leq j \leq n\}$. $\square$

Descending chains of absorptive polynomials may be infinite. To simplify the analysis of infima, we first observe that it suffices to consider ω-chains since $\mathbb{S}^\infty[X]$ is countable.

Lemma 6.22. *Let S be a cpo semiring and $C \subseteq S$ a countable chain. Then there is an ω-chain $(x_i)_{i<\omega}$ such that $x_i \in C$ for all $i < \omega$ and further $\bigsqcap C = \bigsqcap_i x_i$.*

Moreover, if $f : S \to T$ is a monotone function into a cpo semiring T, then additionally $\bigsqcap f(C) = \bigsqcap_i f(x_i)$. Analogue statements hold for suprema.

Proof. We only show the statement involving f, as it implies the first (using the identity on S) and only consider infinite C (otherwise the statement is trivial).

Fix a bijection $g : \omega \to C$ and recursively define

$$x_0 = g(0), \qquad x_{i+1} = \begin{cases} g(i+1), & \text{if } g(i+1) \leq x_i \\ x_i, & \text{otherwise} \end{cases}$$

Then $(x_i)_{i<\omega}$ is a descending ω-chain. We have $x_i \in C$ by definition and thus $\bigsqcap_i f(x_i) \geq \bigsqcap f(C)$. Conversely, for every $c \in C$ there is an i with $g(i) = c$ and thus $c \geq x_i$. By monotonicity, we have $f(c) \geq f(x_i) \geq \bigsqcap_i f(x_i)$ and hence $\bigsqcap f(C) \geq \bigsqcap_i f(x_i)$. $\square$

In [Mrk18] it is shown (as requirement of the notion of *absorptive lattice semirings* they consider) that the natural order on $\mathbb{S}^\infty[X]$ is completely distributive and that the operations satisfy some notion of continuity. We adapt this result to our setting:

Theorem 6.23. $\mathbb{S}^\infty[X]$ *is a completely distributive, continuous lattice semiring.*

Proof. Only continuity is left to prove. It is already shown in [Mrk18] that addition is continuous (due to proposition 4.21) and that multiplication satisfies

$$Q \cdot \bigsqcup S = \bigsqcup Q \cdot S, \qquad \text{for every set } S \subseteq \mathbb{S}^\infty[X] \text{ and } Q \in \mathbb{S}^\infty[X],$$

$$Q \cdot \bigsqcap_{i<\omega} P_i = \bigsqcap_{i<\omega} Q \cdot P_i, \qquad \text{for every } \omega\text{-chain } (P_i)_{i<\omega} \subseteq \mathbb{S}^\infty[X] \text{ and } Q \in \mathbb{S}^\infty[X].$$

The last statement can be lifted to arbitrary chains $C \subseteq \mathbb{S}^\infty[X]$ by the previous lemma which yields an ω-chain $(P_i)_{i<\omega}$ with $\bigsqcap C = \bigsqcap_i P_i$ and $\bigsqcap Q \cdot C = \bigsqcap_i Q \cdot P_i$. □

Characterization of Infima

We have seen in theorem 6.20 that suprema of arbitrary sets can easily be characterized. In this section, we seek a characterization for infima of ω-chains which lays the foundation for later results. The main idea is to consider chains of monomials instead of polynomials as they are much easier to handle.

For the remainder of this section, fix a descending ω-chain of polynomials $(P_i)_{i<\omega}$ in $\mathbb{S}^\infty[X]$ and let $P_\omega = \bigsqcap_i P_i$. We consider special chains of monomials:

Definition 6.24. A descending ω-chain $\mathfrak{m} = (\mathfrak{m}_i)_{i<\omega}$ in $\mathcal{M}_\infty$ is a *monomial chain through* $(P_i)_{i<\omega}$ if $m_i \in P_i$ for all $i < \omega$. We denote the set of such monomial chains by $\mathfrak{M}$.

In order to characterize P_ω, we introduce a graph representation for $(P_i)_{i<\omega}$ to visualize all possible monomial chains. An example is shown in figure 4.

Definition 6.25. Given $(P_i)_{i<\omega}$, the *chain graph* is the directed graph $G = (V, E)$ with

- $V = \{(i, m) \mid i < \omega, m \in P_i\}$,
- E consists precisely of the edges $((i, m), (i + 1, m'))$ with $m' \leq m$.

Lemma 6.26. *The following statements hold in any chain graph:*

(1) *For each node (i, m), there is a path from $(0, m_0)$ to (i, m) for some $m_0 \in P_0$. Every node (i', m') on this path satisfies $m \leq m'$.*

(2) *The infinite paths from nodes $(0, m)$ correspond to monomial chains through $(P_i)_{i<\omega}$.*

(3) *The chain graph is finitely branching.*

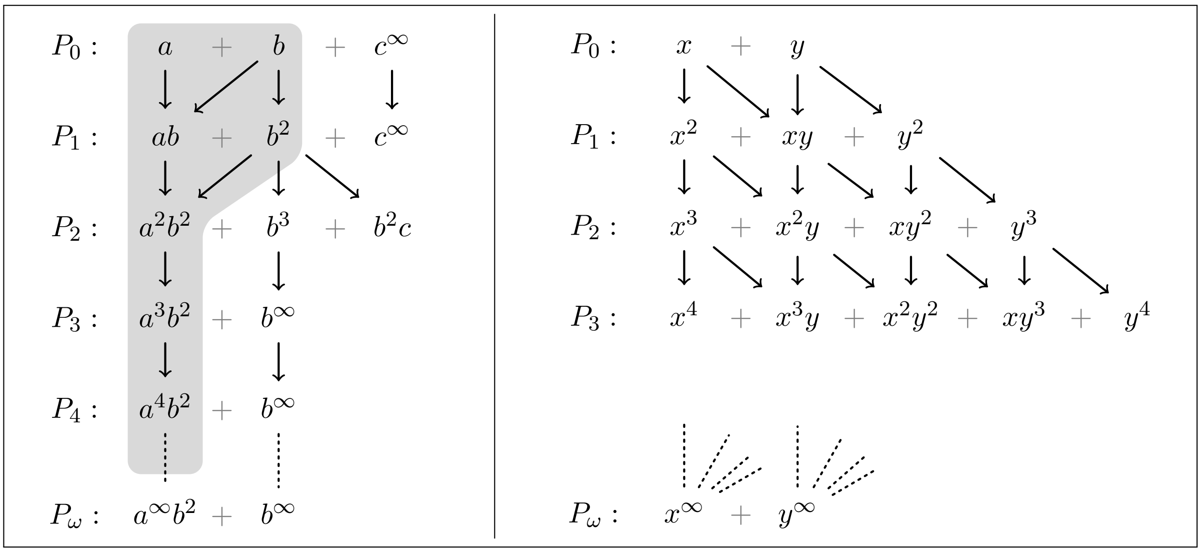

Figure 4: *Two examples of descending polynomial chains, the corresponding chain graph and infimum. On the left, the gray area indicates the subgraph for $a^\infty b^2$ (in the proof of theorem 6.27). On the right, the number of monomials increases unboundedly.*

Proof. Claim (1) is by induction on i. The base case for $i = 0$ is trivial. For $(i+1, m)$, we have $m \in P_{i+1} \leq P_i$. So there is an $m' \in P_i$ with $m \leq m'$ and thus an edge from (i, m') to $(i+1, m)$. By induction, a path to (i, m') exists and can be extended to $(i+1, m)$. Claims (2) and (3) hold by construction and since absorptive polynomials are finite. $\square$

We are now ready to prove the characterization for infima of chains in $\mathbb{S}^\infty[X]$. The proof makes use of the well-known Kőnig's lemma which states that every node in a finitely branching, connected, infinite graph is part of an infinite path.

Theorem 6.27 (characterization of infima).

$$\bigsqcap_{i<\omega} P_i = \bigsqcup \Big\{ \bigsqcap_{i<\omega} \mathfrak{m}_i \mid \mathfrak{m} \in \mathfrak{M} \Big\} = \mathsf{maximals}\Big(\Big\{ \bigsqcap_{i<\omega} \mathfrak{m}_i \mid \mathfrak{m} \in \mathfrak{M} \Big\} \Big)$$

Proof of theorem 6.27. By definition, $\mathfrak{m}_i \leq P_i$ and thus $\bigsqcap_i \mathfrak{m}_i \leq \bigsqcap_i P_i$ for every $\mathfrak{m} \in \mathfrak{M}$. Hence $\bigsqcap_i P_i \geq \bigsqcup \{ \bigsqcap_i \mathfrak{m}_i \mid \mathfrak{m} \in \mathfrak{M} \}$.

For the other direction, consider the infimum $P_\omega = \bigsqcap_i P_i$. We claim that for every monomial $m_\omega \in P_\omega$, there is a monomial chain $\mathfrak{m} \in \mathfrak{M}$ with $\bigsqcap_i \mathfrak{m}_i \geq m_\omega$. This implies the other direction and thus the theorem.

To prove the claim, fix a monomial $m_\omega \in P_\omega$ and let $G = (V, E)$ be the chain graph for $(P_i)_{i<\omega}$. Consider the subgraph G' induced by $V' = \{(i, m) \mid m \geq m_\omega\} \subseteq V$. For each i, we have $P_\omega \leq P_i$ and hence there is a monomial $m \in P_i$ with $m_\omega \leq m$. By the previous lemma, each of the corresponding nodes $(i,m) \in V'$ lies on a path in G' from a node $(0,m_0)$ with $m_0 \in P_0$. As P_0 is finite, there must be an $m_0 \in P_0$ that is the origin of infinitely many such paths. Then the component of $(0,m_0)$ in G' is infinite.

By Kőnig's lemma, there is an infinite path from $(0,m_0)$ in G' which corresponds to a monomial chain $\mathfrak{m} \in \mathfrak{M}$. By definition of G', we have $\mathfrak{m}_i \geq m_\omega$ for all $i < \omega$ and thus $\bigsqcap_i \mathfrak{m}_i \geq m_\omega$. This proves the claim, ending the proof. □

Remark 1: We only need finitely many monomial chains to define the infimum due to the finiteness of P_ω. So there are $\mathfrak{m}^{(1)}, \dots, \mathfrak{m}^{(k)} \in \mathfrak{M}$ such that

$$\bigsqcap_{i<\omega} P_i = \bigsqcup \{ \bigsqcap_{i<\omega} \mathfrak{m}_i \mid \mathfrak{m} \in \mathfrak{M}_{\text{fin}} \} = \bigsqcap_{i<\omega} \mathfrak{m}_i^{(1)} + \dots + \bigsqcap_{i<\omega} \mathfrak{m}_i^{(k)}$$

Remark 2: This result is similar in shape to the complete distributivity of $\mathbb{S}^\infty[X]$. Let $(P_i)_{i<\omega}$ be a family (not necessarily a chain) of polynomials. Since polynomials are antichains of monomials (which we can, under slight abuse of notation, view as polynomials), we have $P_i = \bigsqcup P_i$ for all i. Complete distributivity then yields

$$\bigsqcap_{i<\omega} P_i = \bigsqcap_{i<\omega} \bigsqcup P_i = \bigsqcup_{f \in \mathcal{F}} \bigsqcap_{i<\omega} f(i)$$

where $\mathcal{F}$ is the set of choice functions with $f(i) \in P_i$ for all $i < \omega$. Each function f defines what we might call a *monomial family through* $(P_i)_{i<\omega}$. If the P_i form a chain, then each $\mathfrak{m} \in \mathfrak{M}$ induces the function $f(i) = \mathfrak{m}_i$, but there are also functions which do not correspond to monomial chains. The key insight of the above theorem is thus that for *chains* of polynomials, we only need *chains* of monomials to describe the infimum. Monomial chains have a very simple structure (this becomes apparent when we view $\mathcal{M}_\infty$ as function semiring) which makes this insight useful for us.

The characterization via monomial chains enables us to prove further statements about infima. A first consequence of our proof is chain-positivity.

Corollary 6.28. $\mathbb{S}^\infty[X]$ *is chain-positive.*

Proof. By lemma 6.22, it suffices to consider ω-chains $(P_i)_{i<\omega}$. As $P_i \neq 0$ by assumption, there is at least one node (i,m) in the chain graph with $m \in P_i$ for each $i < \omega$. As P_0 is finite, there is $m_0 \in P_0$ such that infinitely many of these nodes are reachable from $(0,m_0)$. By Kőnig's lemma, there is an infinite path from $(0,m_0)$ which corresponds to a monomial chain $\mathfrak{m} \in \mathfrak{M}$. By the above theorem, $P_\omega \geq \bigsqcap_i \mathfrak{m}_i$ and thus $P_\omega \neq 0$. □

Universality

A central property of $\mathbb{S}^\infty[X]$ is the following universality result which, combined with the fundamental property, has several consequences and in particular shows that provenance computation in $\mathbb{S}^\infty[X]$ is sufficient to capture computations in any absorptive continuous semiring. The remainder of this section is devoted to its proof.

Theorem 6.29 (universality). *Let T be an absorptive continuous semiring and let further $f : X \to T$ be a variable assignment. Then there is a* uniquely *defined* cpo-semiring homomorphism $h : \mathbb{S}^\infty[X] \to T$ *that extends f (i.e., $h(x) = f(x)$ for all $x \in X$).*

We first observe that, similar to $x^\infty = \prod_n x^n$ for $x \in X$, we can define an *infinitary power* a^∞ for every $a \in T$, as the powers $(a^n)_{n<\omega}$ form a descending chain due to absorption.

Definition 6.30. The *infinitary power* of $a \in T$ is defined as $a^\infty = \prod_{n<\omega} a^n$.

Lemma 6.31. *Let $a \in T$. Then $a \cdot a^\infty = a^\infty$.*

Proof. This follows from the continuity of T, as $a \cdot \prod_n a^n = \prod_n (a \cdot a^n) = \prod_n a^n$. □

Let us start by defining the homomorphism h. On monomials, h is uniquely defined by f as we require h to be multiplicative and continuous:

- For $x \in X$, we require $h(x) = f(x)$ and further $h(1) = 1$.
- By induction and multiplicity, it follows that $h(x^n) = f(x)^n$ for all $n < \omega$.
- By continuity, $x^\infty = \prod_n x^n$ implies $h(x^\infty) = h(\prod_n x^n) = \prod_n h(x^n)$.
- Multiplicity requires $h(x^a y^b) = h(x^a) \cdot h(y^b)$ for $x, y \in X$ and $a, b \in \mathbb{N}^\infty$.
- We thus have to set $h(m) := \prod_{x\in X} f(x)^{m(x)}$ for monomials $m \in \mathcal{M}_\infty$.

This uniquely defines h on monomials and we can verify multiplicity:

$$h(m_1) \cdot h(m_2) = \prod_{x\in X} f(x)^{m_1(x)} \cdot \prod_{x\in X} f(x)^{m_2(x)} \overset{(*)}{=} \prod_{x\in X} f(x)^{m_1(x)+m_2(x)} = h(m_1 \cdot m_2)$$

where $(*)$ holds as T is commutative using the above lemma if $m_1(x) = \infty$ or $m_2(x) = \infty$. Note that without the continuity requirement, h would not be uniquely defined as we could also define a homomorphism with $h(x^\infty) = 0$.

Lemma 6.32. *h preserves the absorption order on monomials.*

Proof. For $m \in \mathcal{M}_\infty$, we can write $h(m) = \prod_{x \in X} f(x)^{m(x)}$. If $m_1 \leq m_2$, then $m_1(x) \geq m_2(x)$ and thus $f(x)^{m_1(x)} \leq f(x)^{m_2(x)}$ by absorption in T, for all $x \in X$. As X is finite and multiplication in T is monotone, this implies $h(m_1) \leq h(m_2)$. $\square$

The definition on monomials uniquely lifts to polynomials:

- Additivity requires $h(m_1 + m_2) = h(m_1) + h(m_2)$ for monomials m_1, m_2.
- We thus have to define $h(P) \coloneqq \sum_{m \in P} h(m)$ for $P \in \mathbb{S}^\infty[X]$.

To prove that h is a homomorphism, it remains to show that h is additive and multiplicative. As T is idempotent, $a \leq b$ implies $a + b = b$ for $a, b \in T$. Together with the previous lemma, $m_1 \leq m_2$ implies $h(m_1) + h(m_2) = h(m_2) = h(m_1 + m_2)$ for monomials m_1, m_2. If m_1 and m_2 are incomparable, then $h(m_1) + h(m_2) = h(m_1 + m_2)$ by definition. By induction, this yields $h(m_1 + \cdots + m_k) = h(m_1) + \cdots + h(m_k)$ for all monomials $m_1, \ldots, m_k$ and $k < \omega$. This implies that h is indeed a homomorphism:

$$h(P+Q) = h\Big(\sum_{m \in P} m + \sum_{m' \in Q} m'\Big) = \sum_{m \in P} h(m) + \sum_{m' \in Q} h(m') = h(P) + h(Q)$$

$$\begin{aligned} h(P \cdot Q) = h\Big(\sum_{m \in P, m' \in Q} m \cdot m'\Big) &= \sum_{m \in P, m' \in Q} h(m) \cdot h(m') \\ &= \Big(\sum_{m \in P} h(m)\Big) \cdot \Big(\sum_{m' \in Q} h(m')\Big) = h(P) \cdot h(Q) \end{aligned}$$

The crucial observation for theorem 6.29 is that h is continuous. This is easy to see for suprema due to the ascending chain condition. In fact, we can directly show that h preserves suprema of arbitrary sets S which are essentially described by the union $\bigcup S$.

Proposition 6.33. *Let $h : \mathbb{S}^\infty[X] \to T$ be a semiring homomorphism where T is an absorptive cpo semiring. Then for every set $S \subseteq \mathbb{S}^\infty[X]$,*

$$\bigsqcup h(S) = h\Big(\bigsqcup S\Big)$$

Proof. First note that T is a lattice semiring, so $\bigsqcup h(S)$ is defined. Moreover, recall that h preserves addition and thereby also the natural order. The direction $h(\bigsqcup S) \geq \bigsqcup h(S)$ is then trivial by monotonicity (same argument as in proposition 4.32).

For the other direction, note that $\bigsqcup S = m_1 + \cdots + m_k$ for some monomials $m_1, \ldots, m_k$. Consider one monomial m_i. Since $\bigsqcup S = \mathsf{maximals}\big(\bigcup S\big)$, there is a $P \in S$ with $m_i \in P$. Then $\bigsqcup h(S) \geq h(P) \geq h(m_i)$. This holds for each $1 \leq i \leq k$ and implies $\bigsqcup h(S) \geq h(m_1) + \cdots + h(m_k) = h(m_1 + \cdots + m_k) = h(\bigsqcup S)$. $\square$

For infima, the proof is more complicated. We first show that h is continuous on monomials and then turn our attention to ω-chains of polynomials. We start with a simple fact on ω-chains in continuous semirings.

Lemma 6.34 (splitting lemma)**.** *Let T be a continuous semiring, $\circ \in \{+, \cdot\}$ and let $(a_i)_{i<\omega}$ and $(b_i)_{i<\omega}$ be descending ω-chains in T. Then*

$$\bigsqcap_{i<\omega} a_i \circ b_i = \bigsqcap_{i<\omega} a_i \;\circ\; \bigsqcap_{j<\omega} b_j$$

Proof. We have the following equality, where $(*)$ holds by continuity of T:

$$\bigsqcap_i a_i \circ b_i \overset{(1)}{=} \bigsqcap_i \bigsqcap_j a_i \circ b_j \overset{(*)}{=} \bigsqcap_i (a_i \circ \bigsqcap_j b_j) \overset{(*)}{=} \bigsqcap_i a_i \;\circ\; \bigsqcap_j b_j$$

We prove both directions of (1). Fix i, j and let $k = \max(i, j)$. Then $a_i \circ b_j \geq a_k \circ b_k \geq \bigsqcap_k a_k \circ b_k$ by monotonicity of $\circ$. As i, j are arbitrary, this proves $\bigsqcap_i \bigsqcap_j a_i \circ b_j \geq \bigsqcap_k a_k \circ b_k$.

For the other direction, we have $a_i \circ b_i \geq a_i \circ \bigsqcap_j b_j$ for every i by monotonicity of $\circ$. By continuity, $a_i \circ b_i \geq \bigsqcap_j a_i \circ b_j$ for every i, and thus $\bigsqcap_i a_i \circ b_i \geq \bigsqcap_i \bigsqcap_j a_i \circ b_j$. □

Proving that h preserves infima of monomial chains is easy due to the fixed number of variables. This enables us to split a chain of monomials into chains for all variables.

Proposition 6.35. *For every ω-chain $(m_i)_{i<\omega}$ of monomials, $h(\bigsqcap_i m_i) = \bigsqcap_i h(m_i)$.*

Proof. Let $m_\omega = \bigsqcap_i m_i$. As the infimum is computed variable-wise, we have $m_\omega(x) = \bigsqcup_i m_i(x)$ for all $x \in X$. Then

$$\begin{aligned} \bigsqcap_i h(m_i) = \bigsqcap_i \prod_{x \in X} h(x)^{m_i(x)} &\overset{(1)}{=} \prod_{x\in X} \bigsqcap_i h(x)^{m_i(x)} \\ &\overset{(2)}{=} \prod_{x \in X} h(x)^{\bigsqcup m_i(x)} = \prod_{x\in X} h(x)^{m_\omega(x)} = h(m_\omega) \end{aligned}$$

where (1) holds by the splitting lemma. For (2), we fix an $x \in X$ and make a case distinction. If $\bigsqcup_i m_i(x) = n$ for some $n < \omega$, then there is an i such that $m_j(x) = n$ for all $j \geq n$ and thus $\bigsqcap_i h(x)^{m_i(x)} = h(x)^n$. If $\bigsqcup_i m_i(x) = \infty$, then the values $m_i(x)$ become arbitrarily large and thus $\bigsqcap_i h(x)^{m_i(x)} = \bigsqcap_i h(x)^i = h(x)^\infty$. □

For polynomials, more work is required. While the number of monomials is always finite, it is not fixed and cannot be limited in general (see figure 4 for an example). For the remainder of this section, let $(P_i)_{i<\omega}$ be a descending ω-chain and let $P_\omega = \bigsqcap_i P_i$ be its infimum. Theorem 6.27 expresses P_ω in terms of monomial chains, but this result does

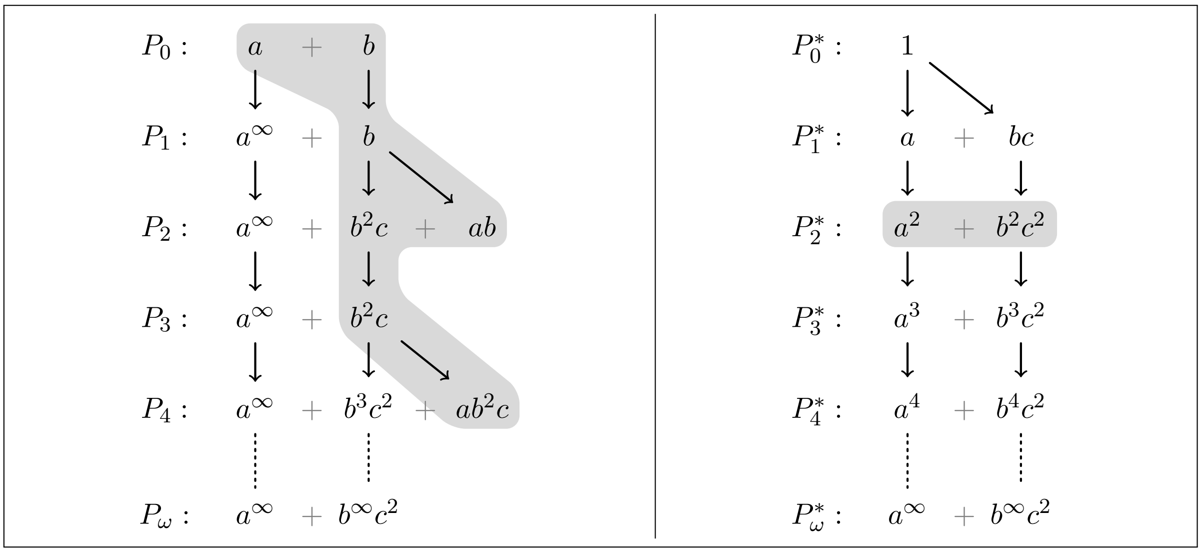

Figure 5: *An example of a polynomial chain and the corresponding canonical chain. On the left, the gray area indicates the uncovered monomials for P_2^* (in the proof of proposition 6.38).*

not apply to $\prod_i h(P_i)$. Instead, we would like to find a decomposition into finitely many monomial chains such that for all i,

$$P_i = \mathfrak{m}_i^{(1)} + \cdots + \mathfrak{m}_i^{(k)} \quad \text{and thus} \quad h(P_i) = h(\mathfrak{m}_i^{(1)}) + \cdots + h(\mathfrak{m}_i^{(k)})$$

An application of the splitting lemma would then suffice to prove continuity. As figure 4 shows, we do not have a decomposition into finitely many chains in general: The number of monomials might grow unboundedly and there may be monomials (such as b^2c and c^∞ in the picture on the left) that are not part of any monomial chain in $\mathfrak{M}$.

We therefore introduce a second, canonical ω-chain $(P_i^*)_{i<\omega}$ which by construction is composed of finitely many monomial chains. This chain has the same infimum $P_\omega^* = P_\omega$ and we show that, when applying h, we still have $\prod_i h(P_i^*) \geq \prod_i h(P_i)$ which suffices to prove theorem 6.29. An example is shown in figure 5.

Definition 6.36. Let $m \in \mathcal{M}_\infty$ be a monomial. The *canonical monomial chain* $(m_n^*)_{n<\omega}$ induced by m is defined as follows, for all $x \in X$:

$$m_n^*(x) = \min(n, m(x))$$

The *canonical polynomial chain* $(P_n^*)_{n<\omega}$ for $P_\omega = m_1 + \cdots + m_k$ is defined by

$$P_n^* = (m_1)_n^* + \cdots + (m_k)_n^*$$

It is clear from this definition that $\prod_n m_n^* = m$ and thus $\prod_n P_n^* = P_\omega$. We make two further observations which are crucial for the following proof.

Lemma 6.37. *The canonical monomial chain has the following properties:*

(1) *If $m, v \in \mathcal{M}_\infty$ with $m \leq v$, then $m^*_n \leq v^*_n$ for all $n < \omega$.*

(2) *If $m = \bigsqcap_i m_i$ for an ω-chain $(m_i)_{i<\omega}$ of monomials, then $\forall n\ \exists i$: $m^*_n \geq m_i$.*

Proof. Recall that $m \leq v$ holds if $m(x) \geq v(x)$ for each $x \in X$. Hence $m \leq v$ implies $\min(n, m(x)) \geq \min(n, v(x))$ and thus $m^*_n \leq v^*_n$ for all $n < \omega$.

For (2), fix $n < \omega$ and $x \in X$. As the infimum m is computed variable-wise, we have $m(x) = \bigsqcup_i m_i(x)$ and proceed by case distinction.

- If $m(x) = \infty$, there must be an i with $m_i(x) \geq n$ and thus $m^*_n(x) = n \leq m_i(x)$.
- If $m(x) = c \in \mathbb{N}$, then we must have $m_i(x) = c$ for some sufficiently large i and hence $m^*_n(x) = c \leq m_i(x)$.

So for each x, we can find an i_x such that $m^*_n(x) \leq m_{i_x}(x)$. As X is finite, we can set $i = \max\{i_x \mid x \in X\}$ and obtain $m^*_n \geq m_i$. ☐

Given a monomial chain $(m_i)_{i<\omega}$, the lemma relates this chain to the canonical monomial chain of its infimum. The next step is to lift this observation to polynomial chains.

Proposition 6.38. *The canonical polynomial chain $(P^*_n)_{n<\omega}$ for P_ω satisfies*

$$\forall n\ \exists i : \quad P^*_n \geq P_i$$

Proof. Fix n and assume towards a contradiction that there is no i with $P^*_n \geq P_i$.

Consider the chain graph for $(P_i)_{i<\omega}$. We call a node (i, m) *covered* if $P^*_n \geq m$ and *uncovered* otherwise. For each i, there must be an $m \in P_i$ such that (i, m) is uncovered (as otherwise $P^*_n \geq P_i$). By lemma 6.26, there is a path from $(0, m_0)$ to (i, m) for some $m_0 \in P_0$ and for all nodes (i', m') on this path, we have $m \leq m'$. But then all nodes on the path are uncovered, as otherwise $P^*_n \geq m'$ would imply $P^*_n \geq m$.

As P_0 is finite, there must be an $m_0 \in P_0$ that lies on infinitely many of these paths. If we consider the subgraph induced by the uncovered nodes, then the component of m_0 must thus be infinite and by Kőnig's lemma, we obtain a monomial chain $\mathfrak{m} \in \mathfrak{M}$ that only contains uncovered monomials. Let $\mathfrak{m}_\omega = \bigsqcap_i \mathfrak{m}_i$ be its infimum.

By the characterization of infima in theorem 6.27, we know that $P_\omega \geq \bigsqcap_i \mathfrak{m}_i$, so there is a monomial $v \in P_\omega$ with $v \geq \mathfrak{m}_\omega$. Consider the canonical monomial chains $(v^*_n)_{n<\omega}$ and $(\mathfrak{m}^*_n)_{n<\omega}$ for v and $\mathfrak{m}_\omega$, respectively. By lemma 6.37 (2), there is an i such that $\mathfrak{m}^*_n \geq \mathfrak{m}_i$ and part (1) of the lemma implies $v^*_n \geq \mathfrak{m}^*_n$. And since v^*_n occurs in P^*_n, we thus have $P^*_n \geq v^*_n \geq \mathfrak{m}^*_n \geq \mathfrak{m}_i$ which is a contradiction to $\mathfrak{m}_i$ being uncovered. ☐

This relation between the chain $(P_i)_{i<\omega}$ and the canonical chain $(P_n^*)_{n<\omega}$ yields the connection between $\bigsqcap_i h(P_i)$ and $\bigsqcap_n h(P_n^*)$ we need:

Corollary 6.39.

$$\bigsqcap_{n<\omega} P_n^* \geq \bigsqcap_{i<\omega} P_i, \qquad \textit{and further,} \qquad \bigsqcap_{n<\omega} h(P_n^*) \geq \bigsqcap_{i<\omega} h(P_i).$$

The above corollary is a direct consequence of the proposition, as h is monotone. It is the main piece needed for the universality theorem which we can now finally prove, putting everything together.

Proof of the universality theorem 6.29.

We have already defined $h : \mathbb{S}^\infty[X] \to T$ and shown that this definition is uniquely induced by $f : X \to T$. Moreover, we have seen that h is a homomorphism, so what remains is to show that h is continuous.

Let $C \subseteq \mathbb{S}^\infty[X]$ be a chain. By proposition 6.33, $\bigsqcup h(C) = h(\bigsqcup C)$. For infima, one direction is trivial by monotonicity of h (see proposition 4.32):

$$\bigsqcap h(C) \geq h(\bigsqcap C)$$

For the other direction, consider an ω-chain $(P_i)_{i<\omega}$ with infimum $P_\omega = v_1 + \cdots + v_k$. We use the above corollary, the splitting lemma and the observation that h preserves infima of monomial chains from proposition 6.35 to conclude:

$$\begin{aligned}
\bigsqcap_{i<\omega} h(P_i) \leq \bigsqcap_{n<\omega} h(P_n^*) &= \bigsqcap_{n<\omega} h\big((v_1)_n^*\big) + \cdots + h\big((v_k)_n^*\big) \\
&\overset{\text{split}}{=} \bigsqcap_{n<\omega} h\big((v_1)_n^*\big) + \cdots + \bigsqcap_{n<\omega} h\big((v_k)_n^*\big) \\
&\overset{6.35}{=} h\big(\bigsqcap_{n<\omega}(v_1)_n^*\big) + \cdots + h\big(\bigsqcap_{n<\omega}(v_k)_n^*\big) \\
&= h\big(v_1\big) + \cdots + h\big(v_k\big) = h(P_\omega) = h(\bigsqcap_{i<\omega} P_i)
\end{aligned}$$

This result lifts to arbitrary chains by lemma 6.22, so h is continuous. $\square$

Before we consider consequences of the universality theorem, we show that it is as broad as possible. It is relatively easy to see that the target semiring T has to be absorptive for h to be well-defined. The following example shows that T must further be continuous, as otherwise h is not guaranteed to be continuous.

Example 6.40. Recall the absorptive non-continuous semiring $T = \mathbb{B} \times \mathbb{V} \setminus \{(\top, 0)\}$ from example 4.28. Consider the homomorphism $h : \mathbb{S}^\infty[a, b] \to T$ induced by the variable assignment $a \mapsto (\perp, 1)$, $b \mapsto (\top, \frac{1}{2})$. Then h is not continuous, as witnessed by the descending ω-chain $a + b^n$ (for $n < \omega$):

$$\prod_n h(a + b^n) = \prod_n (\perp, 1) + (\top, \tfrac{1}{2})^n = \prod_n (\top, 1) = (\top, 1)$$

$$h\Big(\prod_n a + b^n\Big) = h(a + b^\infty) = (\perp, 1) + (\top, \tfrac{1}{2})^\infty = (\perp, 1)$$

where $(\top, \frac{1}{2})^\infty = \prod_n (\top, \frac{1}{2^n}) = (\perp, 0)$, as we exclude the element $(\top, 0)$.

Continuity of Semiring Semantics

Based on the results of [Mrk18], we have seen that $\mathbb{S}^\infty[X]$ is a continuous semiring. With the universality theorem, we can now address our open question from section 5.5 by showing that in $\mathbb{S}^\infty[X]$, $[\![\varphi]\!]_\ell$ is continuous in ℓ. We first need auxiliary statements on the infinitary power which mostly follow from the splitting lemma 6.34.

Lemma 6.41. *Let S be an absorptive continuous semiring. The infinitary power has the following properties, for $a, b \in S$:*

(1) $a \cdot a^\infty = a^\infty$ *(and thus $a^n \cdot a^k = a^{n+k}$ for all $n, k \in \mathbb{N}^\infty$),*

(2) $(ab)^\infty = a^\infty b^\infty$ *(and thus $(a^n)^\infty = a^{n\cdot\infty}$ for all $n \in \mathbb{N}^\infty$),*

(3) $(a + b)^\infty = a^\infty + b^\infty$,

(4) $a \leq b$ *implies* $a^\infty \leq b^\infty$ *(monotonicity),*

(5) $h(a^\infty) = h(a)^\infty$ *for continuous homomorphisms* $h : S \to T$,

(6) $(\prod_i x_i)^\infty = \prod x_i^\infty$ *for descending ω-chains $(x_i)_{i<\omega}$ in S.*

Proof. Property (1) was already shown in lemma 6.31, (2) holds by the splitting lemma 6.34 and (4) follows from (3) since $a \leq b \iff a + b = b$ (in idempotent semirings). Statement (5) follows from the properties of h, in (6) we can swap the two infima:

$$\prod_{i<\omega} x_i^\infty = \prod_{i<\omega}\prod_{n<\omega} x_i^n = \prod_{n<\omega}\prod_{i<\omega} x_i^n \overset{6.34}{=} \prod_{n<\omega}\Big(\prod_{i<\omega} x_i\Big)^n = \Big(\prod_{i<\omega} x_i\Big)^\infty$$

Only (3) remains to prove. We clearly have $(a + b)^n \geq a^n + b^n$ (for all $n < \omega$) and hence $(a + b)^\infty \geq a^\infty + b^\infty$. For the other direction, fix n and consider $(a + b)^{2n} =$

$\sum_{i=0}^{2n}\binom{2n}{i}a^{2n-i}b^i$. Each summand is absorbed by either a^n (if $i \leq n$) or by b^n (if $i \geq n$), hence $a^n + b^n \geq (a+b)^{2n} \geq (a+b)^\infty$ and the claim follows. □

In $\mathbb{S}^\infty[X]$, we further have an analogue of property (6) for arbitrary suprema.

Lemma 6.42. *For each set $S \subseteq \mathbb{S}^\infty[X]$, we write $S^\infty = \{P^\infty \mid P \in S\}$ and have:*

$$\left(\bigsqcup S\right)^\infty = \bigsqcup S^\infty$$

Proof. The direction $(\bigsqcup S)^\infty \geq \bigsqcup S^\infty$ is trivial by monotonicity of the infinitary power. For the other direction, let $\bigsqcup S = m_1 + \cdots + m_k$ for a finite number of monomials $m_1, \ldots, m_k$. By the above lemma, $(\bigsqcup S)^\infty = m_1^\infty + \cdots + m_k^\infty$.

Fix one such monomial m_i. As $\bigsqcup S = \mathsf{maximals}(\bigcup S)$, there is a $P \in S$ with $m_i \in P$. Hence $m_i \leq P$ and thus $m_i^\infty \leq P^\infty \leq \bigsqcup S^\infty$. As this holds for each $1 \leq i \leq k$, we can conclude $m_1^\infty + \cdots + m_k^\infty \leq \bigsqcup S^\infty$. □

Remark: The above lemma does not generalize to absorptive continuous semirings, even if we require S to be a chain. For an example, consider any chain $(a_i)_{i<\omega} \subseteq \mathbb{V}$ approaching 1 with $a_i < 1$. Then $(\bigsqcup_i a_i)^\infty = 1^\infty = 1$ while $\bigsqcup_i (a_i^\infty) = \bigsqcup_i 0 = 0$.

Theorem 6.43. *In $\mathbb{S}^\infty[X]$, $[\![\varphi]\!]_\ell$ is continuous in ℓ. That is, for every formula φ and every chain $\mathcal{L}$ of $\mathbb{S}^\infty[X]$-interpretations,*

$$[\![\varphi]\!]_{\bigsqcup \mathcal{L}} = \bigsqcup\{[\![\varphi]\!]_\ell \mid \ell \in \mathcal{L}\} \quad \text{and} \quad [\![\varphi]\!]_{\bigsqcap \mathcal{L}} = \bigsqcap\{[\![\varphi]\!]_\ell \mid \ell \in \mathcal{L}\}.$$

Proof. We show that $[\![\varphi]\!]_{\bigsqcap_i \ell_i} = \bigsqcap_i [\![\varphi]\!]_{\ell_i}$ for descending ω-chains $(\ell_i)_{i<\omega}$. The statement for $\mathcal{L}$ follows from lemma 6.22 (note that $[\![\varphi]\!]_\ell$ is monotone in ℓ), suprema are analogous.

We consider the *most general* interpretation ℓ^*. Let $X^* = \{x_L \mid L \in \mathsf{Lit}_A\}$ and define $\ell^*(L) = x_L$. Let further $[\![\varphi]\!]_{\ell^*} = m_1 + \cdots + m_k$. Given any $\mathbb{S}^\infty[X]$-interpretation ℓ, we can apply the universality theorem to the variable assignment $x_L \mapsto \ell(L)$ and obtain a cpo-semiring homomorphism $h_\ell : \mathbb{S}^\infty[X^*] \to \mathbb{S}^\infty[X]$ with $\ell = h_\ell \circ \ell^*$.

Let $\ell_\omega = \bigsqcap_i \ell_i$ to ease notation. We claim that $h_{\ell_\omega}(m) = \bigsqcap_i h_{\ell_i}(m)$ for every monomial m in $\mathbb{S}^\infty[X^*]$. To see this, let $x_L \in X$ and $a \in \mathbb{N}^\infty$. We then have:

$$h_{\ell_\omega}(x_L^a) \overset{(1)}{=} h_{\ell_\omega}(x_L)^a = \ell_\omega(L)^a = \Big(\bigsqcap_{i<\omega} \ell_i(L)\Big)^a \overset{(2)}{=} \bigsqcap_{i<\omega} \ell_i(L)^a \overset{(1)}{=} \bigsqcap_{i<\omega} h_{\ell_i}(x_L^a)$$

where (1) and (2) hold by lemma 6.41. If we consider suprema instead of infima, then (2) needs the second lemma 6.42. This is the reason why the proof works for $\mathbb{S}^\infty[X]$, but does not generalize to arbitrary absorptive continuous semirings. The claim for m follows by applying the splitting lemma. We can now prove the theorem:

$$\begin{aligned}
[\![\varphi]\!]_{\ell_\omega} = h_{\ell_\omega}([\![\varphi]\!]_{\ell^*}) &= h_{\ell_\omega}(m_1) + \dots + h_{\ell_\omega}(m_k) \\
&= \prod_{i<\omega} h_{\ell_i}(m_1) + \dots + \prod_{i<\omega} h_{\ell_i}(m_k) \\
&= \prod_{i<\omega} h_{\ell_i}(m_1) + \dots + h_{\ell_i}(m_k) = \prod_{i<\omega} h_{\ell_i}([\![\varphi]\!]_{\ell^*}) = \prod_{i<\omega} [\![\varphi]\!]_{\ell_i} \qquad \square
\end{aligned}$$

This continuity result allows us to apply Kleene's fixed-point theorem 4.6. In particular, all fixed-point iterations $(\pi_\beta)_{\beta \in \mathsf{On}}$ of (least and greatest) fixed-point formulae terminate at π_ω (or earlier), answering our open question in section 5.5 for the semiring $\mathbb{S}^\infty[X]$.

Corollary 6.44. *Let ℓ be a $\mathbb{S}^\infty[X]$-interpretation and $\varphi = [\mathsf{fp}\, R x.\, \vartheta](\mathbf{y})$. Let F_ℓ^ϑ and $(\pi_\beta)_{\beta \in \mathsf{On}}$ be the corresponding update operator and fixed-point iteration. Then F_ℓ^ϑ is continuous and the iteration reaches its fixed point at π_ω (i.e., $\pi_{\omega+1} = \pi_\omega$).*

We now want to generalize this answer to absorptive semirings. Recall that the continuity of $[\![\varphi]\!]_\ell$ does not hold in all absorptive continuous semirings as witnessed by the semiring $\mathbb{L}$ in example 5.21. We can nevertheless state the following consequence.

Corollary 6.45. *Semiring semantics $[\![\varphi]\!]_\ell$ are well-defined for fully ω-continuous absorptive semirings and fixed-point iterations $(\pi_\beta)_{\beta \in \mathsf{On}}$ of LFP formulae always reach the fixed point at π_ω (i.e., $\pi_{\omega+1} = \pi_\omega$).*

In order to establish this result, let us review the proofs of both the universality theorem 6.29 and the fundamental property 5.17. Regarding the universality result, where did we rely on T being continuous? Most of the reasoning was concerned with $\mathbb{S}^\infty[X]$, so the only places are in the definition of the infinitary power (which we used to define h) and in the splitting lemma, both of which only talk about ω-chains. We also need the continuity of T for lemma 6.22, but this is only used to lift our reasoning from ω-chains to arbitrary chains. By carefully following our proof, we can thus formulate a similar universality for ω-continuous semirings:

Theorem 6.46. *Let T be an absorptive, fully ω-continuous semiring and let $f : X \to T$ be an assignment of the variables. Then there is a* uniquely *defined* fully ω-continuous homomorphism $h : \mathbb{S}^\infty[X] \to T$ *that extends f.*

Now recall the proof of the fundamental property in section 5.4. The continuity of the homomorphism h is only needed for fixed-point formulae, to be precise only for limit ordinal steps in the fixed-point iteration. Given a homomorphism $h : S \to T$, let $\varphi = [\mathsf{fp}\, R\,\mathbf{x}.\, \vartheta](\mathbf{y})$ and let $(\pi_\beta)_{\beta \in \mathsf{On}}$ and $(\sigma_\beta)_{\beta \in \mathsf{On}}$ be the corresponding iterations in S and T, respectively. The proof shows by induction that $h(\pi_\beta) = \sigma_\beta$.

If S, T and h are not continuous but only fully ω-continuous, then the iterations are certainly well-defined for $\beta < \omega + \omega$ and we still have $h(\pi_\beta) = \sigma_\beta$ (in fact, this even holds for larger β due to the argument in lemma 6.22, but this is not relevant here). In particular, this applies to $\beta = \omega$ and $\beta = \omega + 1$ which suffices to prove corollary 6.45.

Proof of corollary 6.45. Given a a formula φ, a fully ω-continuous absorptive semiring T and a T-interpretation ℓ, we again consider the most general $\mathbb{S}^\infty[X^*]$-interpretation ℓ^* as in the proof of theorem 6.43. By the above universality, there is a fully ω-continuous homomorphism $h : \mathbb{S}^\infty[X^*] \to T$ such that $\ell = h \circ \ell^*$.

The only interesting case is if $\varphi = [\mathsf{fp}\, R\,\mathbf{x}.\, \vartheta](\mathbf{y})$. Let $(\pi_\beta)_{\beta \in \mathsf{On}}$ be the fixed-point iteration for $[\![\varphi]\!]_{\ell^*}$ in $\mathbb{S}^\infty[X^*]$ and let $(\sigma_\beta)_{\beta < \omega+\omega}$ be the first steps of the fixed-point iteration for $[\![\varphi]\!]_{h \circ \ell^*}$ in T. Note that the iteration σ_β is well-defined, as T is fully ω-complete.

Following the proof of the fundamental property 5.17, we see that $h(\pi_\beta) = \sigma_\beta$ for all $\beta < \omega + \omega$. We further know from corollary 6.44 that $\pi_\omega = \pi_{\omega+1}$. But then, $\sigma_{\omega+1} = h(\pi_{\omega+1}) = h(\pi_\omega) = \sigma_\omega$, hence σ_ω is the (least or greatest) fixed-point of F_ℓ^ϑ. $\square$

These results answer our open question for absorptive semirings and underline the importance of the universality theorem for $\mathbb{S}^\infty[X]$.

As a conclusion of this section, we can say that absorptive polynomials are arguably the most interesting semiring for provenance analysis of LFP. The reason is that they are not only compatible with standard semantics, but also universal among all absorptive continuous semirings which covers most of the application semirings presented earlier. Absorption of course limits the information to shortest proofs (or minimal witnesses), but we have seen for $\mathbb{W}[X]$ that this is most likely what we want in the presence of greatest fixed-points. Compared to $\mathbb{W}[X]$, the exponents in $\mathbb{S}^\infty[X]$ can provide additional information which, depending on the situation, can be useful.

Example 6.47. Recall the following example we considered for $\mathbb{W}[X]$ and $\mathbb{B}[\![X]\!]$:

$$\varphi_{\text{infpath}}(u) = [\mathsf{gfp}\, R\, x.\ \exists y(Exy \wedge Ry)](u)$$

$$\begin{tikzcd} \overset{u}{\bullet} \arrow[r, "a"'] & \overset{w}{\bullet} \arrow[r, "c"'] \arrow[loop below, "b"] & \overset{v}{\bullet} \end{tikzcd}$$

In $\mathbb{W}[X]$, we obtained the monomial abc (not corresponding to an infinite path) and in $\mathbb{B}[\![X]\!]$, the formula evaluated to 0. In $\mathbb{S}^\infty[X]$, we instead get a reasonable result:

$$\begin{array}{lllllllll} \pi(u): & 1 & & a & & ab+ac & & ab^2+abc & & ab^\infty \\ \pi(w): & 1 & \mapsto & b+c & \mapsto & b^2+bc & \mapsto & b^3+b^2c & \cdots & b^\infty \\ \pi(v): & 1 & & 0 & & 0 & & 0 & & 0 \end{array}$$

Not only do we see that there is only one infinite path from u, but we further learn that edge a is used only once whereas edge b appears infinitely often on this path. Interpreting φ_{infpath} in $\mathbb{S}^\infty[X]$ thus provides information about reachable cycles: The edges with exponent ∞ form the cycle, the other edges indicate the path to the cycle.

The universal property now allows us to evaluate φ_{infpath} in other absorptive semirings by simply instantiating the variables. For example, assume that we have high confidence in edges a and c, but low confidence in b. We can then consider the homomorphism $h : \mathbb{S}^\infty[X] \to \mathbb{V}$ with $h(a) = h(c) = 1$ and $h(b) = \frac{1}{2}$ and obtain the following result:

$$[\![\varphi_{\text{infpath}}(u)]\!]_{h\circ\ell} = 1 \cdot (\tfrac{1}{2})^\infty = 0$$

$$\begin{tikzcd} \overset{u}{\bullet} \arrow[r, "1"'] & \overset{w}{\bullet} \arrow[r, "1"'] \arrow[loop below, "1/2"] & \overset{v}{\bullet} \end{tikzcd}$$

Another possibility is to use the access control semiring $\mathbb{A}$. If we think of the vertices as different pages of a website and the edges as links between them (some of which can only be accessed by members or administrators with certain access privileges), we can easily compute the access level a user needs to spend his lifetime clicking on these links:

$$[\![\varphi_{\text{infpath}}(u)]\!]_{h\circ\ell} = \max(\mathsf{C}, \mathsf{S}^\infty) = \max(\mathsf{C}, \mathsf{S}) = \mathsf{S}$$

$$\begin{tikzcd} \overset{u}{\bullet} \arrow[r, "\mathsf{C}"'] & \overset{w}{\bullet} \arrow[r, "\mathsf{P}"'] \arrow[loop below, "\mathsf{S}"] & \overset{v}{\bullet} \end{tikzcd}$$

For some applications, we do not need the additional information of exponents as in the access control example. We can then use the finite semiring $\mathsf{PosBool}[X]$ that results from $\mathbb{S}^\infty[X]$ by dropping exponents.

6.5 PosBool-Semiring

The $\mathsf{PosBool}[X]$ semiring is the simplest polynomial semiring we consider. We have already introduced it in section 4.3 in terms of positive boolean formulae (up to logical equivalence). It also results from $\mathbb{W}[X]$ by applying absorption [Gre11]. Coming from $\mathbb{S}^\infty[X]$, we can further characterize $\mathsf{PosBool}[X]$ as the set of absorptive polynomials without exponents.

Definition 6.48. The semiring $\mathsf{PosBool}[X]$ is the set of absorptive polynomials with variables in X and exponents in $\mathbb{B}$, with operations defined as for $\mathbb{S}^\infty[X]$. We view monomials $m \in \mathcal{M}_\mathbb{B}$ as sets $\{x \in X \mid m(x) = \top\}$.

To see that this is isomorphic to our earlier definition, note that we can identify absorptive polynomials such as $xy + xz$ with positive formulae in disjunctive normal form such as $(x \wedge y) \vee (x \wedge z)$. Absorption corresponds to the logical equivalence $x \vee (x \wedge y) \equiv x$. Since $\mathsf{PosBool}[X]$ is isomorphic to positive boolean formulae *up to logical equivalence*, this makes semiring semantics less dependent on the syntactic presentation of a formula. That is, $\varphi \vee \varphi$, $\varphi \wedge \varphi$ and φ all yield the same value when interpreted over $\mathsf{PosBool}[X]$, whereas their interpretations in $\mathbb{W}[X]$ or $\mathbb{S}^\infty[X]$ may differ. In particular, $\mathsf{PosBool}[X]$ is multiplicatively idempotent, for example $(x+y)^2 = x + xy + y = x + y$ due to absorption.

We have already mentioned that $\mathsf{PosBool}[X]$ is the distributive lattice (with $+, \cdot$ as join and meet) freely generated by the set X. We can thus formulate a universality result similar to the previous ones. Note that $\mathsf{PosBool}[X]$ is finite, so the continuity is trivial.

Proposition 6.49. *Let S be an absorptive and multiplicative idempotent cpo semiring. Every function $f : X \to S$ extends to a unique homomorphism $h : \mathsf{PosBool}[X] \to T$ that coincides with f on X. This homomorphism is a cpo-semiring homomorphism.*

Intuitively, $\mathsf{PosBool}[X]$ combines the abstractions of $\mathbb{W}[X]$ and $\mathbb{S}^\infty[X]$ for provenance analysis: It only counts which literals are used in a proof (not how often) and it only considers shortest proofs (due to absorption). This makes $\mathsf{PosBool}[X]$ less informative than the other two semirings but also easier to handle. If we only want to know whether a formula depends on a certain literal, then the information in $\mathsf{PosBool}[X]$ is sufficient and we do not need the additional information (and complexity) of $\mathbb{W}[X]$ or $\mathbb{S}^\infty[X]$. We consider an example of provenance analysis in $\mathsf{PosBool}[X]$ in section 7.2 where we use it to compute winning strategies in parity games.

6.6 Polynomials in Dual Indeterminates

The polynomial semirings introduced in this chapter can be used for *tracking provenance* where we track certain literals in the computation of $[\![\varphi]\!]_\ell$ by mapping them to (unique) variables. We sometimes not only want to *track* given facts, but instead leave *open* whether a certain fact is true or not. This form of *reverse provenance analysis* can give information about *all* models of a formula instead of fixing a particular model. We can always achieve this by setting $\ell(R\mathbf{a}) = x$ and $\ell(\neg R\mathbf{a}) = y$, but then we also obtain monomials such as xy which correspond to opposing literals and thus bear no connection to actual models.

In [GT17a], Grädel and Tannen therefore proposed polynomials in dual indeterminates for reverse analysis of FO which we adapt for LFP. That is, we use variables X for positive literals and a dual set $\bar{X} = \{\bar{x} \mid x \in X\}$ for negative literals (e.g., $\ell(R\mathbf{a}) = x$ and $\ell(\neg R\mathbf{a}) = \bar{x}$). Instead of working in $\mathbb{S}^\infty[X \cup \bar{X}]$, where we still have conflicting monomials such as $x\bar{x}$, we define a quotient semiring $\mathbb{S}^\infty[X, \bar{X}]$ in which $x\bar{x} = 0$.

In this section, we formally introduce dual-indeterminate polynomials (dual polynomials for short) as quotient semirings and show that the resulting semirings are suitable for fixed-point logics. Many previous results for compatibility with standard semantics and duality require positive semirings, whereas $x\bar{x} = 0$ introduces divisors of 0. We thus devise an alternative duality result for dual polynomials.

Definition

In the following, we understand $\mathcal{M}$, $\mathcal{M}_\infty$ and all monomials to be defined over the variables $X \cup \bar{X}$ instead of X. Most of the statements hold for several semirings; we write $\mathbb{P}[X]$ for either of the polynomial semirings $\mathbb{W}[X]$, $\mathbb{S}^\infty[X]$, $\mathsf{PosBool}[X]$ or the formal power series $\mathbb{N}^\infty[\![X]\!]$, $\mathbb{B}[\![X]\!]$.

Definition 6.50. A monomial $m \in \mathcal{M}_\infty$ is *conflicting* if there is an $x \in X$ such that both $m(x) \neq 0$ and $m(\bar{x}) \neq 0$. Otherwise m is *non-conflicting*.

For $P \in \mathbb{P}[X \cup \bar{X}]$, we define $\widehat{P} \in \mathbb{P}[X \cup \bar{X}]$ by

$$\begin{aligned} \widehat{P} &= \{m \in P \mid m \text{ is non-conflicting}\} \subseteq P, && \text{for } \mathbb{S}^\infty, \mathbb{W}, \mathsf{PosBool} \\ \widehat{P}(m) &= 0 \text{ if } m \text{ is conflicting, and } \widehat{P}(m) = P(m) \text{ otherwise}, && \text{for power series} \end{aligned}$$

For provenance with dual indeterminates, we want to work only with the polynomials $\widehat{P}$. We thus define the following congruence relation.

Definition 6.51. We define $\sim$ on $\mathbb{P}[X \cup \bar{X}]$ by

$$P \sim Q \quad \Longleftrightarrow \quad \widehat{P} = \widehat{Q}$$

We write $[P]_\sim$ for the equivalence class containing P.

Proposition 6.52. *$\sim$ is a congruence relation on $\mathbb{P}[X \cup \bar{X}]$.*

Proof. By definition, $\sim$ is an equivalence relation. For the congruence property, let $P \sim P'$ and $Q \sim Q'$, so $\widehat{P} = \widehat{P'}$ and $\widehat{Q} = \widehat{Q'}$. We first consider $\mathbb{S}^\infty[X \cup \bar{X}]$. For addition,

$$\widehat{P+Q} = \mathsf{maximals}\{m \in P \cup Q \mid m \text{ non-conflicting}\} = \mathsf{maximals}(\widehat{P} \cup \widehat{Q}) = \widehat{P} + \widehat{Q}$$

With the analogue for $\widehat{P' + Q'}$, it follows that $\widehat{P+Q} = \widehat{P'+Q'}$ and thus $P+Q \sim P'+Q'$. If m_1 is conflicting, then so is $m_1 m_2$ (for every monomial m_2). We thus have

$$\begin{aligned}\widehat{PQ} &= \mathsf{maximals}\{m_1 m_2 \mid m_1 \in P, m_2 \in Q, m_1 m_2 \text{ non-conflicting}\}\\ &= \mathsf{maximals}\{m_1 m_2 \mid m_1 \in \widehat{P}, m_2 \in \widehat{Q}, m_1 m_2 \text{ non-conflicting}\}\end{aligned}$$

Together with the analogue for $\widehat{P'Q'}$, we obtain $PQ \sim P'Q'$ and $\sim$ is a congruence. The proof for $\mathsf{PosBool}[X \cup \bar{X}]$ is analogous.

For $\mathbb{N}^\infty[\![X \cup \bar{X}]\!]$ the operations are defined monomial-wise, so for non-conflicting m,

$$\begin{aligned}(P+Q)(m) &= P(m) + Q(m) = \widehat{P}(m) + \widehat{Q}(m)\\ (PQ)(m) &= \sum_{m = m_1 m_2} P(m_1) \cdot Q(m_1) \overset{(*)}{=} \sum_{m = m_1 m_2} \widehat{P}(m_1) \cdot \widehat{Q}(m_1)\end{aligned}$$

where $(*)$ holds as $m = m_1 m_2$ implies that m_1 and m_2 are also non-conflicting. Together with the corresponding statements for $P' + Q'$ and $P'Q'$, it follows that $\sim$ is a congruence. The proofs for $\mathbb{W}[X \cup \bar{X}]$ and $\mathbb{B}[\![X \cup \bar{X}]\!]$ are analogous. $\square$

We can then use $\sim$ to define the desired quotient semirings[8] (cf. [Gol99, chapter 8]).

Definition 6.53. The semirings $\mathbb{P}[X, \bar{X}]$ are the quotients $\mathbb{P}[X \cup \bar{X}]/_\sim$, that is

$$\mathbb{P}[X, \bar{X}] = \{[P]_\sim \mid P \in \mathbb{P}[X \cup \bar{X}]\}, \qquad [P]_\sim \circ [Q]_\sim = [P \circ Q]_\sim \text{ (for } \circ \in \{+, \cdot\})$$

We refer to these semirings as dual(-indeterminate) polynomials.

[8] Equivalently, we could use the ideal generated by the monomials $x\bar{x}$ to define the quotients. Note however that, unlike in rings, not every ideal defines a quotient semiring.

Note that $\sim$ is a particularly simple congruence in the sense that we can unambiguously identify each equivalence class $[P]_\sim$ with its representative $\widehat{P}$. If the context is clear, we may thus write $\widehat{P}$ for $[P]_\sim$ to simplify notation (as in the next definition).

While we can still work with model-defining $\mathbb{P}[X, \overline{X}]$-interpretations, we need a different notion to unleash the full power of dual polynomials. Grädel and Tannen introduce *model-compatible* interpretation for reverse provenance analysis [GT17a].

Definition 6.54. A $\mathbb{P}[X, \bar{X}]$-interpretation ℓ is *model-compatible* if for each positive literal $R\mathbf{a} \in \mathsf{Lit}_A$, one of the following holds:

(1) $\ell(R\mathbf{a}) = x$ and $\ell(\neg R\mathbf{a}) = \bar{x}$, for some $x \in X$,

(2) $\ell(R\mathbf{a}) = 0$ and $\ell(\neg R\mathbf{a}) = 1$,

(3) $\ell(R\mathbf{a}) = 1$ and $\ell(\neg R\mathbf{a}) = 0$.

Although they satisfy $\ell(L) \cdot \ell(\neg L) = 0$ and $\ell(L) + \ell(\neg L) \neq 0$, model-compatible interpretations are not model-defining in general. Instead, such interpretations only partially define models (literals that are mapped to $0, 1$) but leave other facts open (literals mapped to $x, \bar{x}$), so they are compatible with several models. Before we consider an example, let us first make a few general observations about dual polynomials.

Properties

In the following, we establish that dual polynomials remain continuous (and in particular suitable for fixed-point logic). It is further clear from the definition that $\mathbb{P}[X, \overline{X}]$ is idempotent or absorptive whenever $\mathbb{P}[X \cup \overline{X}]$ has the respective property.

Lemma 6.55. *$\mathbb{P}[X, \bar{X}]$ is naturally ordered and has the following properties:*

(1) *If $P, Q \in \mathbb{P}[X \cup \bar{X}]$ with $P \leq Q$, then $[P]_\sim \leq [Q]_\sim$.*

(2) *If $[P]_\sim \leq [Q]_\sim$, then $\widehat{P} \leq \widehat{Q}$.*

Proof. For (1), $P + P' = Q$ implies $[P]_\sim + [P']_\sim = [Q]_\sim$ and thus $[P]_\sim \leq [Q]_\sim$.

For (2), $[P]_\sim + [P']_\sim = [Q]_\sim$ implies $[P + P']_\sim = [Q]_\sim$ and thus $\widehat{P + P'} = \widehat{Q}$. We have shown that $\widehat{P + P'} = \widehat{P} + \widehat{P'}$ in the proof of proposition 6.52 and thus $\widehat{P} \leq \widehat{Q}$.

Statement (2) implies that $\mathbb{P}[X, \bar{X}]$ is naturally ordered. $\square$

Proposition 6.56. *Let $C \subseteq \mathbb{P}[X \cup \bar{X}]$ be a chain and $[C]_\sim = \{[P]_\sim \mid P \in C\}$. Then*

$$\bigsqcup [C]_\sim = \left[\bigsqcup C\right]_\sim \quad \textit{and} \quad \bigsqcap [C]_\sim = \left[\bigsqcap C\right]_\sim$$

Proof. Note that $[C]_\sim$ is a chain by the above lemma. We first show that $[\bigsqcup C]_\sim$ is an upper bound: We have $\bigsqcup C \geq P$ for all $P \in C$ and thus $[\bigsqcup C]_\sim \geq [P]_\sim$ by the lemma. Now let $[Q]_\sim$ be any upper bound for the chain $[C]_\sim$. We set $\widehat{C} = \{\widehat{P} \mid P \in C\}$, which is a chain by the above lemma, and claim that $\bigsqcup \widehat{C} = \widehat{\bigsqcup C}$. The claim implies the proposition: We have $[Q]_\sim \geq [P]_\sim$ for all $P \in C$. Then $\widehat{Q} \geq \widehat{P}$ by the lemma and hence $\widehat{Q} \geq \bigsqcup \widehat{C} = \widehat{\bigsqcup C}$ which implies $[Q]_\sim \geq [\bigsqcup C]_\sim$. The proof for infima is completely analogous and requires $\bigsqcap \widehat{C} = \widehat{\bigsqcap C}$.

We show the two claims for each semiring. For formal power series, the claim holds as suprema/infima are computed monomial-wise. In $\mathbb{W}[X \cup \bar{X}]$ and $\mathsf{PosBool}[X \cup \bar{X}]$, the claim is trivial due to the finiteness of these semirings. The only interesting case is $\mathbb{S}^\infty[X \cup \bar{X}]$. The claim for suprema is trivial due to the ascending chain property. For infima, we employ the characterization in theorem 6.27.

Let $(P_i)_{i<\omega}$ be an ω-chain in $\mathbb{S}^\infty[X \cup \bar{X}]$. Then $\bigsqcap_i P_i = \bigsqcup\{\bigsqcap_i \mathfrak{m}_i \mid \mathfrak{m} \in \mathfrak{M}\}$. Consider any $\mathfrak{m} \in \mathfrak{M}$. If $\mathfrak{m}_i$ is conflicting for some i, then so is $\bigsqcap_i \mathfrak{m}_i \leq \mathfrak{m}_i$. Conversely, if $\bigsqcap_i \mathfrak{m}_i$ is conflicting, then there is an i such that $\mathfrak{m}_i$ is conflicting. Following this observation,

$$\widehat{\bigsqcap_{i<\omega} P_i} = \bigsqcup \left\{ \widehat{\bigsqcap_{i<\omega} \mathfrak{m}_i} \mid \mathfrak{m} \in \mathfrak{M} \right\} = \bigsqcup \left\{ \bigsqcap_{i<\omega} \mathfrak{m}_i \mid \mathfrak{m} \in \widehat{\mathfrak{M}} \right\} = \bigsqcap_{i<\omega} \widehat{P}_i$$

where $\widehat{\mathfrak{M}} \subseteq \mathfrak{M}$ is the set of monomial chains such that $\mathfrak{m}_i$ is non-conflicting for each i – these are precisely the monomial chains through $(\widehat{P}_i)_{i<\omega}$. This result lifts to arbitrary chains by lemma 6.22 (note that $\widehat{\cdot}$ is monotone, i.e., $P \leq Q$ implies $\widehat{P} \leq \widehat{Q}$). □

Corollary 6.57. $\mathbb{P}[X, \bar{X}]$ *is a continuous semiring.*

Proof. We first show that $\mathbb{P}[X, \bar{X}]$ is a cpo semiring. Let C be any chain and consider $\widehat{C} = \{\widehat{P} \mid [P]_\sim \in C\}$. Then $\widehat{C}$ is a chain in $\mathbb{P}[X \cup \bar{X}]$ and $C = [\widehat{C}]_\sim$. By the previous proposition, $\bigsqcup[\widehat{C}]_\sim$ and $\bigsqcap[\widehat{C}]_\sim$ exist.

For continuity, we need to show $[Q]_\sim \cdot \bigsqcup C = \bigsqcup([Q]_\sim \cdot C)$ (and the analogues for addition/infima). As $\mathbb{P}[X \cup \bar{X}]$ is continuous, we again use the chain $\widehat{C}$:

$$\begin{aligned}
[Q]_\sim \cdot \bigsqcup C &= [Q]_\sim \cdot \bigsqcup [\widehat{C}]_\sim \overset{6.56}{=} [Q]_\sim \cdot \left[\bigsqcup \widehat{C}\right]_\sim = \left[Q \cdot \bigsqcup \widehat{C}\right]_\sim \\
&\overset{\text{cont.}}{=} \left[\bigsqcup Q \cdot \widehat{C}\right]_\sim \overset{6.56}{=} \bigsqcup [Q \cdot \widehat{C}]_\sim = \bigsqcup [Q]_\sim \cdot [\widehat{C}]_\sim = \bigsqcup [Q]_\sim \cdot C
\end{aligned}$$

□

These results show that working with dual polynomials is possible for all semirings discussed in this chapter. Given that suprema and infima are preserved in the sense of proposition 6.56, it is not surprising that switching from $\mathbb{P}[X \cup \overline{X}]$ to $\mathbb{P}[X, \overline{X}]$ also preserves $[\![\varphi]\!]_\ell$. That is, instead of computing with dual indeterminates we may as well compute in $\mathbb{P}[X \cup \overline{X}]$ and afterwards discard all conflicting monomials.

Corollary 6.58. *Let ℓ be a $\mathbb{P}[X \cup \overline{X}]$-interpretation. This induces the interpretation*

$$[\ell]_\sim : \mathsf{Lit}_A \to \mathbb{P}[X, \overline{X}],\ R\mathbf{a} \mapsto \big[\ell(R\mathbf{a})\big]_\sim$$

Then for every formula φ,

$$[\![\varphi]\!]_{[\ell]_\sim} = [[\![\varphi]\!]_\ell]_\sim$$

Proof sketch. Induction on the negation normal form of φ. For literals, the claim holds by definition of $[\ell]_\sim$. For $\wedge$, $\vee$, $\exists$, $\forall$, it follows by induction using the definition of the operations in $\mathbb{P}[X, \overline{X}]$. For fixed-point formulae, consider the fixed-point iterations $(\pi_\beta)_{\beta \in \mathsf{On}}$ and $(\sigma_\beta)_{\beta \in \mathsf{On}}$ for $[\![\varphi]\!]_{[\ell]_\sim}$ and $[\![\varphi]\!]_\ell$, respectively. One can show by induction on β that $\pi_\beta = [\sigma_\beta]_\sim$. For $\beta = 0$, note that if $\top$ is the top element of $\mathbb{P}[X \cup \overline{X}]$, then $[\top]_\sim$ is the top element of $\mathbb{P}[X, \overline{X}]$. For limit ordinals, the claim follows by proposition 6.56. □

The universality theorems can also be adapted to dual polynomials. If the variable assignment $f : X \cup \overline{X} \to T$ respects dual variables such that $f(x) \cdot f(\bar{x}) = 0$, then the induced homomorphism $h : \mathbb{P}[X \cup \overline{X}]$ satisfies $h(x\bar{x}) = 0$ and thus factors through the quotient $\mathbb{P}[X, \overline{X}]$. We state the universality for absorptive dual polynomials, results for the other semirings can be adapted in the same way.

Theorem 6.59. *Let T be an absorptive continuous semiring and let $f : X \cup \overline{X} \to T$ be an assignment of the variables such that $f(x) \cdot f(\bar{x}) = 0$ for all $x \in X$.*

Then there is a uniquely *defined* cpo-semiring homomorphism $h : \mathbb{S}^\infty[X, \overline{X}] \to T$ *that extends f (that is, $h(x) = f(x)$ and $h(\bar{x}) = f(\bar{x})$ for all $x \in X$).*

Duality

Most of the duality results from section 5.3 require positive semirings. Dual polynomials are not positive due to $x \cdot \bar{x} = 0$ so these results do not apply. As an example, consider the $\mathbb{S}^\infty[X, \overline{X}]$-valuation ℓ over a singleton universe $A = \{a\}$ with

$$\ell : \begin{array}{c} Ra \mapsto x \\ \neg Ra \mapsto 0 \end{array} \quad \text{and} \quad \ell : \begin{array}{c} Pa \mapsto 0 \\ \neg Pa \mapsto \bar{x} \end{array}$$

Then ℓ is complementary for every notion of complementary discussed in section 5.3. However, for $\varphi(a) = Ra \wedge \neg Pa$ we have $[\![\varphi]\!]_\ell + [\![\neg\varphi]\!]_\ell = x\bar{x} + (0+0) = 0$.

The quotient semirings $\mathbb{P}[X, \bar{X}]$ are continuous, so the consistency result still holds. That is, if ℓ is consistent then $[\![\varphi]\!]_\ell \cdot [\![\neg\varphi]\!]_\ell = 0$ for all formulae φ. To establish stronger duality results, we introduce a new notion of complementarity which is based on the observation that we can define complements in $\mathsf{PosBool}[X, \bar{X}]$.

Definition 6.60. Let $P \in \mathsf{PosBool}[X, \bar{X}]$. We define the *complement* of P by

$$\overline{P} = \prod_{m \in P} \overline{m}, \qquad \text{and} \qquad \overline{m} = \sum_{x \in m} \bar{x} + \sum_{\bar{x} \in m} x$$

In particular, we have $\overline{0} = 1$ and $\overline{1} = 0$.

As an example, $\overline{x\bar{y} + z} = (\bar{x} + y) \cdot \bar{z} = \bar{x}\bar{z} + y\bar{z}$. Recall that we can view $\mathsf{PosBool}[X, \bar{X}]$ as set of (normalized) positive boolean formulae. In this sense, $\overline{P}$ corresponds to $\neg P$ (to be precise, to the negation normal form of $\neg P$ where we replace $\neg x$ by $\bar{x}$ and $\neg\bar{x}$ by x to obtain a positive formula), so it is not surprising that we can define such a complement and that it satisfies similar laws as negation:

$$\overline{P+Q} = \overline{P} \cdot \overline{Q}, \quad \overline{PQ} = \overline{P} + \overline{Q}, \quad \overline{\overline{P}} = P, \quad P \cdot \overline{P} = 0$$

Note that, although we call $\overline{P}$ the complement of P, this is not a complement in the sense of a boolean algebra, as $P + \overline{P} = 1$ does not hold in general (e.g., $x + \bar{x} \neq 1$).

Given this complement, we view P and Q as complementary if $\overline{P} = Q$. Instead of defining a similar complement directly in $\mathbb{S}^\infty[X, \bar{X}]$, we use the canonical homomorphism to $\mathsf{PosBool}[X, \bar{X}]$ and make the following generalized definition.

Definition 6.61. Let S be a semiring and $h : S \to \mathsf{PosBool}[X, \bar{X}]$ a cpo-semiring homomorphism. We consider $a, b \in S$ to be complementary if $\overline{h(a)} = h(b)$.

Similar to previous duality results, this complementarity lifts from ℓ to $[\![\varphi]\!]_\ell$.

Proposition 6.62. *Let S and h be as in definition 6.61. If ℓ is a complementary S-valuation, then $[\![\varphi]\!]_\ell$ and $[\![\neg\varphi]\!]_\ell$ are complementary for all formulae φ.*

Proof. Induction on the negation normal form of φ. We only show the interesting cases.

- For $\varphi = \varphi_1 \wedge \varphi_2$, the claim holds by induction and the properties of h:

$$\begin{aligned}\overline{h([\![\varphi]\!]_\ell)} = \overline{h([\![\varphi_1]\!]_\ell) \cdot h([\![\varphi_2]\!]_\ell)} &= \overline{h([\![\varphi_1]\!]_\ell)} + \overline{h([\![\varphi_2]\!]_\ell)} \\ &\overset{\text{IH}}{=} h([\![\neg\varphi_1]\!]_\ell) + h([\![\neg\varphi_2]\!]_\ell) = h([\![\neg\varphi]\!]_\ell)\end{aligned}$$

- For $\varphi = [\mathsf{gfp}\, R\mathbf{x}.\ \vartheta](\mathbf{y})$, we proceed by induction on the fixed-point iterations $(\pi_\beta)_{\beta\in\mathsf{On}}$ and $(\sigma_\beta)_{\beta\in\mathsf{On}}$ for φ and $\neg\varphi$ and show that π_β and σ_β are complementary.
  - For $\beta = 0$, we have $\overline{h(\top)} = \overline{1} = 0 = h(\bot)$.
  - For $\beta{+}1$, we proceed as in the proof of proposition 5.14. Set $\ell' = \ell[R/\pi_\beta, \neg R/\sigma_\beta]$ so we can write $\pi_{\beta+1} = [\![\vartheta]\!]^\bullet_{\ell'}$ and $\sigma_{\beta+1} = [\![\neg\vartheta]\!]^\bullet_{\ell'}$. Note that ℓ' is complementary as π_β, σ_β are complementary. The claim thus holds by induction on ϑ.
  - For limit ordinals, we use that h is continuous:

$$\overline{h(\pi_\lambda)} = \overline{h\big(\bigsqcap\{\pi_\beta \mid \beta < \lambda\}\big)} = \overline{\bigsqcap\{h(\pi_\beta) \mid \beta < \lambda\}} = \overline{h(\pi_{\beta^*})}$$

    Analogously, $h(\sigma_\lambda) = h(\sigma_{\beta^*})$. The ordinal $\beta^* < \lambda$ must exist as $\mathsf{PosBool}[X, \bar{X}]$ is finite. By induction, $\overline{h(\pi_{\beta^*})} = h(\sigma_{\beta^*})$ and the claim follows.

  The claim for φ follows from the inductive claim (for sufficiently large β). □

We can apply this result to $\mathbb{W}[X, \bar{X}]$, $\mathbb{S}^\infty[X, \bar{X}]$ and $\mathsf{PosBool}[X, \bar{X}]$ via the canonical homomorphisms to $\mathsf{PosBool}[X, \bar{X}]$, in particular for model-compatible interpretations.

Corollary 6.63. *Let ℓ be a model-compatible interpretation in $\mathbb{W}[X, \bar{X}]$, $\mathbb{S}^\infty[X, \bar{X}]$ or $\mathsf{PosBool}[X, \bar{X}]$. Then for every formula φ, $[\![\varphi]\!]_\ell \cdot [\![\neg\varphi]\!]_\ell = 0$ and $[\![\varphi]\!]_\ell + [\![\neg\varphi]\!]_\ell \neq 0$.*

Proof. Consider complementarity via the canonical homomorphisms to $\mathsf{PosBool}[X, \bar{X}]$. Then ℓ is complementary and thus $[\![\varphi]\!]_\ell$ and $[\![\neg\varphi]\!]_\ell$ are complementary as well.

Let $P = [\![\varphi]\!]_\ell$ and $Q = [\![\neg\varphi]\!]_\ell$. We know that $\overline{h(P)} = h(Q)$ and thus $0 = h(P) \cdot h(Q) = h(PQ)$. As h is one of the canonical homomorphisms, this implies $PQ = 0$ as claimed.

Assume that $P + Q = 0$. Then $P = 0$ and $Q = 0$ which leads to a contradiction, as $\overline{h(0)} = 1 \neq 0 = h(0)$. We must thus have $P + Q \neq 0$ as claimed. □

Note that the duality is in fact much stronger than this corollary. For example, assume that we start with a model-compatible interpretation in $\mathbb{S}^\infty[X, \bar{X}]$ and obtain $[\![\varphi]\!]_\ell = x\bar{y}^2 + z^\infty$. We then already know that $[\![\neg\varphi]\!]_\ell = \bar{x}^a\bar{z}^b + y^c\bar{z}^d$ for some exponents $a, b, c, d \in \mathbb{N}^\infty$.

Applications in Reverse Analysis

So far, we have mostly shown that dual polynomials have similar properties as the original polynomial semirings. We now turn our attention to applications for provenance analysis. Let us start with a motivating example.

Example 6.64. Consider the following formula which states that on all paths from u, the proposition P holds for every node.

$$\varphi(u) = \big[\mathsf{gfp}\, R\, x.\ Px \wedge \forall y\big(\underbrace{\neg Exy \vee Ry}_{Exy \to Ry}\big)\big](u)$$

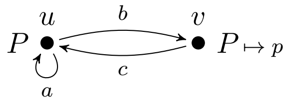

We interpret this formula over the graph shown on the right using the model-compatible interpretation ℓ indicated in the picture, e.g. $\ell(Pu) = 1$, $\ell(Pv) = p$, $\ell(Euu) = a$ and $\ell(\neg Euv) = \bar{b}$. That is, we fix that P holds at u, but leave open whether it also holds at v and whether any of the edges exist. This interpretation results in the following fixed-point iteration in $\mathsf{PosBool}[X, \overline{X}]$:

$$\begin{array}{ll} \pi(u): & 1 \\ \pi(v): & 1 \end{array} \;\mapsto\; \begin{array}{l} 1 \\ p \end{array} \;\mapsto\; \begin{array}{l} \bar{b} + p \\ p \end{array} \;\mapsto\; \begin{array}{l} \bar{b} + p \\ p \end{array}$$

We obtain the overall value $[\![\varphi(u)]\!]_\ell = \bar{b} + p$ which tells us that the formula holds in all models where P holds at v and further in all models where edge b is missing. Note that the variables a and c do not occur in the resulting polynomial, as their presence (or absence) is not relevant for the truth of φ. In this sense, model-compatible interpretations allow us to reason about several models at once. In particular, we see that $\varphi(u)$ does not hold in all models where edge b is present and P does not hold at v.

Using the more expressive semiring $\mathbb{S}^\infty[X, \overline{X}]$ reveals more information about the proofs of $\varphi(u)$ in these models. The iteration then continues in the following way:

$$\begin{array}{ll} \pi(u): & \bar{b} + p \\ \pi(v): & p \end{array} \;\mapsto\; \begin{array}{l} \bar{a}\bar{b} + \bar{a}p + \bar{b}^2 + \bar{b}p + p^2 \\ \bar{b}p + \bar{c}p + p^2 \end{array} \;\mapsto\; \begin{array}{l} \bar{a}\bar{b}+\bar{a}\bar{c}p+\bar{a}p^2+\bar{b}^3+\bar{b}^2p+\bar{b}p^2+\bar{c}p^3+p^4 \\ \bar{a}\bar{b}p+\bar{a}p^2+\bar{b}^2p+\bar{b}p^2+\bar{c}p+p^3 \end{array} \quad \cdots$$

We see that the size of the polynomials quickly increases, so it is not obvious what the fixed point is. Careful examination of all possible monomial chains leads to the following result. In the next chapter, we show how the same result can be obtained without being overly careful.

$$[\![\varphi(u)]\!]_\ell = \bar{a}\bar{b} + \bar{a}\bar{c}p + \bar{b}^\infty + p^\infty$$

First note that this polynomial simplifies to the earlier result $\bar{b} + p$ by dropping the exponents. Due to the universal quantification, witnesses for the truth of $\varphi(u)$ are not as simple as in earlier examples and we again refer to the forthcoming chapter for a concise explanation. For now, we note that one way to satisfy $\varphi(u)$ is to ensure that all paths are finite (and P holds on these paths), which corresponds to the monomials $\bar{a}\bar{b}$ and $\bar{a}\bar{c}p$. If infinite paths are possible, we either have to make sure that P holds at v, leading to p^∞, or that v is unreachable, leading to $\bar{b}^\infty$. In both cases, the literals corresponding to p and $\bar{b}$ have to be used infinitely often due to the infinite paths.

Following [GT17a], we can make the relation between interpretations and models precise. Let $\ell : \mathsf{Lit}_A \to \mathbb{P}[X, \overline{X}]$ be a model-compatible interpretation. Then a model $\mathfrak{A}$ over the same alphabet A is said to be **compatible** with ℓ if $\mathfrak{A} \models L$ for all literals with $\ell(L) = 1$. We denote the set of models compatible with ℓ by Mod_ℓ. A *variable assignment* $f : X \cup \overline{X} \to \mathbb{B}$ with $f(x) = 1$ and $f(\bar{x}) = 0$, or vice versa, allows us to switch from a model-compatible interpretation to a model-defining one.

For the converse, the intuition is that the interpretation $\ell_{\mathfrak{A}}$ induced by a model $\mathfrak{A} \in \mathsf{Mod}_\ell$ is a specialization of ℓ in the sense that $\ell_{\mathfrak{A}}$ results from ℓ by a similar variable assignment. For this to be possible, we require ℓ to be **non-overlapping**. That is, no two different literals are mapped to the same variable x (if the same variable is used for two literals L and L', then a model with $\mathfrak{A} \models L$ but $\mathfrak{A} \not\models L'$ cannot be defined by applying a variable assignment to ℓ). Combining these observations with the universal properties, we can formulate the following adaption of a result[9] in [GT17a] for LFP.

Proposition 6.65. *Let ℓ be a non-overlapping model-compatible S-interpretation where S is one of $\mathbb{W}[X, \overline{X}]$, $\mathbb{S}^\infty[X, \overline{X}]$ or $\mathsf{PosBool}[X, \overline{X}]$. Then a formula φ is Mod_ℓ-satisfiable if, and only if, $[\![\varphi]\!]_\ell \neq 0$ (and φ is Mod_ℓ-valid if, and only if, $[\![\neg\varphi]\!]_\ell = 0$).*

Proof. We assume w.l.o.g. that all variables in X appear in ℓ. First assume that there is $\mathfrak{A} \in \mathsf{Mod}_\ell$ with $\mathfrak{A} \models \varphi$. Consider the $\mathbb{B}$-interpretation $\ell_{\mathfrak{A}}$ induced by $\mathfrak{A}$. We now define a variable assignment $f : X \cup \overline{X} \to \mathbb{B}$ to turn ℓ into $\ell_{\mathfrak{A}}$. For $x \in X$, let L be the unique literal with $\ell(L) = x$, $\ell(\neg L) = \bar{x}$ and set $f(x) = \ell_{\mathfrak{A}}(L)$ and $f(\bar{x}) = \ell_{\mathfrak{A}}(\neg L)$.

Note that f is well-defined since ℓ is non-overlapping. Moreover, f uniquely induces a cpo-semiring homomorphism $h : S \to \mathbb{B}$ due to the previous universality results. By definition of f, we have $\ell_{\mathfrak{A}} = h \circ \ell$ and thus obtain:

$$\mathfrak{A} \models \varphi \quad\Longleftrightarrow\quad 0 \neq [\![\varphi]\!]_{\ell_{\mathfrak{A}}} = [\![\varphi]\!]_{h \circ \ell} = h\big([\![\varphi]\!]_\ell\big) \quad\Longrightarrow\quad [\![\varphi]\!]_\ell \neq 0$$

For the other direction, we assume $[\![\varphi]\!]_\ell \neq 0$ and can thus fix a monomial $m \in [\![\varphi]\!]_\ell$. We again define a variable assignment $f : X \cup \overline{X} \to \mathbb{B}$ such that for each $x \in X$:

- if x appears in m, then $f(x) = 1$ and $f(\bar{x}) = 0$,
- if $\bar{x}$ appears in m, then $f(x) = 0$ and $f(\bar{x}) = 1$,
- otherwise, $f(x) = 1$ and $f(\bar{x}) = 0$ (this is an arbitrary choice).

Note that f is well-defined as m is non-conflicting. Then f induces a cpo-semiring homomorphism $h : S \to \mathbb{B}$ with $h(m) = 1$ and $h \circ \ell$ is model-defining. Let $\mathfrak{A} = \mathfrak{A}_{h \circ \ell}$ be

[9] In [GT17a], the interpretation is not explicitly assumed to be non-overlapping. However, the result formulated there is only correct for non-overlapping interpretations.

the corresponding model. Clearly, $\mathfrak{A} \in \mathsf{Mod}_\ell$ and the claim follows:

$$1 = h(m) \le h(\llbracket \varphi \rrbracket_\ell) = \llbracket \varphi \rrbracket_{h \circ \ell} \quad \Longrightarrow \quad \mathfrak{A} \models \varphi$$

The statement about Mod_ℓ-validity follows from the statement on satisfiability. □

The above proof justifies our explanation in the preceding example: Each monomial m in $\llbracket \varphi \rrbracket_\ell$ corresponds to one (or more) models compatible with ℓ.

As a conclusion of this chapter, we can say that absorptive polynomials $\mathbb{S}^\infty[X]$ and their dual-indeterminate version $\mathbb{S}^\infty[X, \bar{X}]$ are the most interesting semirings for provenance analysis of fixed-point logic. While more general semirings, such as formal power series $\mathbb{N}^\infty\llbracket X \rrbracket$, are available for positive LFP, greatest fixed points require chain-positive semirings and we further need absorption to rule out monomials not corresponding to actual proofs, as seen for the semiring $\mathbb{W}[X]$. The desirable properties of $\mathbb{S}^\infty[X]$, most notably chain-positivity and universality, come at the price of restricting provenance information to what we described as shortest proofs, corresponding to the absorption-maximal monomials. We have, however, seen in several examples that this is often sufficient and that for some applications, we can even work with the much simpler semiring $\mathsf{PosBool}[X]$ which does not track multiplicities (that is, exponents) of variables. From an algebraic point of view, there is no hard requirement for absorption, but we have motivated this property both by the problematic example in $\mathbb{W}[X]$ and by symmetry arguments.

What remains is a precise understanding of the provenance information $\llbracket \varphi \rrbracket_\ell$ we compute in $\mathbb{S}^\infty[X]$ or $\mathsf{PosBool}[X]$. So far, we have always argued in terms of proofs on an intuitive level, but we have seen in example 6.64 that this does not work well in all cases. To improve upon this situation, the following chapter considers winning strategies in model checking games as a well-defined notion of proofs.

7 Understanding Provenance via Games

This chapter presents an alternative view on semiring semantics via model checking games. The goal is a characterization of $[\![\varphi]\!]_\ell$ in terms of winning strategies in the game for φ, in analogy to the proof-tree characterization for $\mathbb{N}[X]$-provenance of FO in [GT17a]. This does not only lead to a better understanding of semiring semantics, but further shows what kind of information semiring provenance can reveal about a formula.

The focus is on the semiring $\mathbb{S}^\infty[X]$ which exhibits a tight connection between the computation of $[\![\varphi]\!]_\ell$ and (truncations of) winning strategies in the associated model checking game. We start by introducing the required notions of parity games, model checking games and strategies. Before we proceed with the two main results, the puzzle lemma and the characterization theorem, we take the opportunity to use parity games as an exemplary application of semiring provenance in $\mathsf{PosBool}[X, \bar{X}]$. At the end of the chapter, we lift our results to absorptive continuous semirings via the universality of $\mathbb{S}^\infty[X]$ and sketch ideas to generalize the characterization to non-absorptive semirings.

7.1 Preliminaries: Parity Games

Towards model checking games, we start by briefly introducing parity games. More background can be found in [GKL+07, chapter 3].

Definition 7.1. A *parity game* $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ is a two player game played on a directed graph (V, E) with $V = V_0 \,\dot\cup\, V_1$ and a node labeling $\Omega : V \to \mathbb{N}$.

For $v \in V$, we denote the set of successor positions by $vE = \{w \mid (v, w) \in E\}$. Positions v with $vE = \emptyset$ are called *terminal* positions. Player 0 moves at the positions in V_0 and player 1 at V_1, possible moves are described by E.

We call $\Omega(v)$ the *priority* of the position $v \in V$. The maximal priority assigned by Ω must be finite and is called the *index* of the game $\mathcal{G}$.

The semantics of a game $\mathcal{G}$ are as follows: If a player cannot move, they lose (this happens at terminal positions). If a play does not end in a terminal position, then it is infinite and the winner is determined by the parity of the smallest priority that occurs infinitely often in the play. Priorities which only occur finitely often are irrelevant.

Definition 7.2. A *play* in $\mathcal{G}$ from v_0 is a sequence $\rho = v_0 v_1 v_2 \ldots$ of positions with $(v_i, v_{i+1}) \in E$ (for all i) which is either infinite or ends in a terminal position.

The play ρ is *winning* for player 0 if either,

(1) $\rho = v_0 v_1 \dots v_k$ and $v_k \in V_1$ (so player 1 cannot move), or

(2) ρ is infinite and the *smallest* number occurring infinitely often in the sequence $\Omega(v_0)\Omega(v_1)\Omega(v_2)\dots$ is *even* (we call this the *parity condition*).

Otherwise, ρ is winning for player 1.

Following [GT19], we represent strategies via the tree unraveling of a game.

Definition 7.3. Let $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ be a parity game. The *tree unraveling* of $\mathcal{G}$ from v_0 is the tree $\mathcal{T}(\mathcal{G}, v_0) = (V^\#, V_0^\#, V_1^\#, E^\#)$ where

- $V^\#$ is the set of all finite paths $\pi = v_0 v_1 \dots v_k$ (for some $k < \omega$) through $\mathcal{G}$,
- $V_\sigma^\# = \{\pi v \in V^\# \mid v \in V_\sigma\}$ is the set of paths ending in a node of player σ,
- $E^\# = \{(\pi v, \pi v v') \mid \pi v \in V^\# \text{ and } (v, v') \in E\}$.

For $\pi, \pi' \in V^\#$, we write $\pi \sqsubseteq \pi'$ if π is a prefix of π'. For a node πv, we call $V(\pi v) = v$ the *position* of πv. For $\pi = v_0 \dots v_k$, we write $|\pi| = k + 1$ for the length of π. Each play $\rho = v_0 v_1 v_2 \dots$ in $\mathcal{G}$ induces a unique path $(v_0), (v_0 v_1), (v_0 v_1 v_2), \dots$ through $\mathcal{T}(\mathcal{G}, v_0)$.

Strategies can then be defined as subtrees of the tree unraveling, which allows for a more visual way to reason about strategies. The model checking games we consider are always finite, so the tree unraveling and all strategies are finitely branching.

Definition 7.4. A *strategy* $\mathcal{S}$ of player $\sigma \in \{0, 1\}$ from v_0 in $\mathcal{G}$ is a subtree of $\mathcal{T}(\mathcal{G}, v_0)$ of the form $\mathcal{S} = (W, F)$ with $W \subseteq V^\#$ and $F \subseteq (W \times W) \cap E^\#$ that satisfies the following conditions. Let $\widehat{V}_\sigma^\#$ be the set $V_\sigma^\#$ without terminal nodes (i.e., leaves).

(1) W is closed under predecessors: if $\pi v \in W$, then also $\pi \in W$,

(2) player σ makes unique choices: if $\pi \in W \cap \widehat{V}_\sigma^\#$, then $|\pi F| = 1$,

(3) all choices of the opponent are considered: if $\pi \in W \cap V_{1-\sigma}^\#$, then $\pi F = \pi E^\#$.

Equivalently, we sometimes view a strategy as a function $\mathcal{S} : W \cap \widehat{V}_\sigma^\# \to V$ with $\mathcal{S}(\pi v) \in vE$ encoding the choices of player σ in (2).

A play π is *consistent* with $\mathcal{S}$ if the corresponding path in $\mathcal{T}(\mathcal{G}, v_0)$ is contained in $\mathcal{S}$. The strategy $\mathcal{S}$ is *winning* if all plays consistent with $\mathcal{S}$ are winning (for player σ).

Two kinds of strategies are of particular interest to us: A *positional* strategy requires the player to make a unique choice for each position, independent of the history, and it is well known that parity games are *positionally determined* (see e.g. [GKL+07] for a proof). We also consider *weakly positional* strategies in which each play must be positional but the player can make different decisions across different plays.

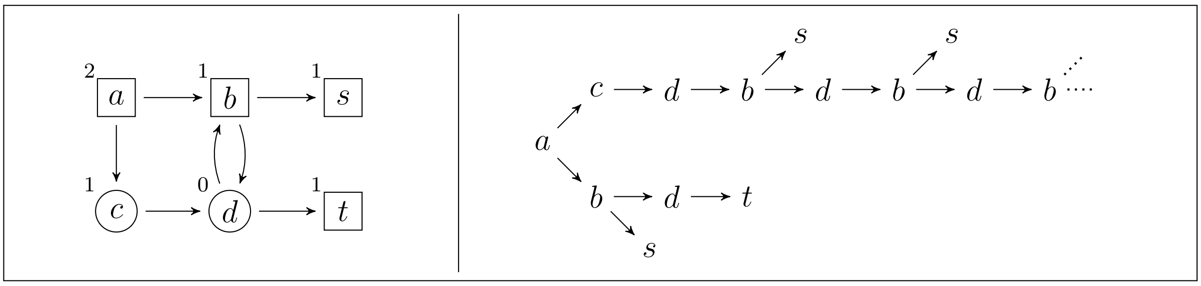

Figure 6: *An example of a small parity game and a depiction of a winning strategy for player* 0 *from position* a*. The strategy is weakly positional, but not positional (at position* d*).*

Definition 7.5. Let $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ be a parity game, $\sigma \in \{0,1\}$ and let $\widehat{V}_\sigma$ be the set V_σ without terminal positions. A strategy $\mathcal{S} = (W, F)$ for player σ in $\mathcal{G}$ is

(1) *positional*, if $\mathcal{S}(\pi v) = \mathcal{S}(\pi' v)$ for all $v \in \widehat{V}_\sigma$ and $\pi, \pi' \in W$,

(2) *weakly positional*, if $\mathcal{S}(\pi v) = \mathcal{S}(\pi' v)$ for all $v \in \widehat{V}_\sigma$ and $\pi, \pi' \in W$ with $\pi \sqsubseteq \pi'$.

An example of a parity game and a winning strategy for player 0 is shown in figure 6. We always depict V_0 by circular and V_1 by $\boxed{\text{rectangular}}$ nodes. The small numbers indicate the priorities of each node. For the strategy, we denote a node πv simply by its position v (note that π is already determined by the path from the root to the node πv).

7.2 Excursion: Strategy Computation with Semirings

Before we look into model checking games, let us consider an application of semiring provenance to parity games. This involves many of the topics we considered so far, most importantly homomorphisms, dual polynomials and absorption.

For this excursion, we fix a parity game $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ with index d (that is, $\Omega : V \to \{0, \ldots, d\}$). It is known that the winning region of player 0 (the nodes from which they have a winning strategy) is definable in LFP by the following construction due to Walukiewicz [Wal02]. In the modal μ-calculus (which can be embedded into LFP and is more concise in this case, see e.g. [GKL+07]) the formula can be written as follows, where we interpret P_i as the set of all positions with priority i.

$$\varphi = \sigma X_d \ldots \mu X_1.\, \nu X_0.\, \Big(\big(V_0 \to \Diamond\big(\bigwedge_{i=0}^{d} P_i \to X_i\big)\big) \;\wedge\; \big(V_1 \to \Box\big(\bigwedge_{i=0}^{d} P_i \to X_i\big)\big) \Big)$$

Here, σ is either ν or μ, depending on whether d is even or odd. This formula induces an LFP formula $\varphi(u)$ defining the winning region. We have not introduced the modal μ-calculus, but the exact definition of $\varphi(u)$ is mostly irrelevant for our application. Due

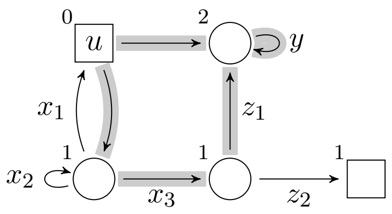

$$[\![\varphi(u)]\!]_\ell \;=\; x_1y \;+\; x_3yz_1 \;+\; x_3yz_2$$

Figure 7: *A parity game with edge labeling ℓ (left) and the corresponding provenance information (right). The highlighted monomial corresponds to the highlighted positional winning strategy.*

to the positional determinacy, we know that $\varphi(u)$ holds if, and only if, player 0 has a positional winning strategy from the position u.

Our goal is to compute all positional winning strategies for player 0 in $\mathcal{G}$. To this end, we work in $\mathsf{PosBool}[X, \bar{X}]$ and define the model-compatible interpretation ℓ as follows. We view $\mathcal{G}$ as a structure over the signature $\{V_0, V_1, P_0, \dots, P_d, E\}$ as used in the formula above. The idea is that we fix the priorities and all edges of the opponent and track edges of player 0 by variables.

- We set $X = \{x_{vw} \mid (v, w) \in E,\, v \in V_0\}$.
- For $(v, w) \in E$ with $v \in V_0$, we set $\ell(Evw) = x_{vw}$ and $\ell(\neg Evw) = \bar{x}_{vw}$.
- All other literals are interpreted by $\{0, 1\}$ according to $\mathcal{G}$.

An example (with simplified variable names) is shown in figure 7. One can see that the monomials of $[\![\varphi]\!]_\ell$ correspond to the positional winning strategies and we now want to prove this observation in general. The following lemma is the only consideration on the formula $\varphi(u)$ we need for this purpose.

Lemma 7.6. *The dual variables $\bar{x}_{vw}$ do not occur in $[\![\varphi(u)]\!]_\ell$.*

Proof. Negative literals of the form $\neg Exy$ only occur in the subformula

$$V_1 \to \Box\Big(\bigwedge_{i=0}^{d} P_i \to X_i\Big) \qquad \text{or, in LFP,} \qquad \neg V_1x \,\vee\, \forall y\Big(\neg Exy \vee \bigwedge_{i=0}^{d} \neg P_iy \vee X_iy\Big)$$

We only introduce variables $\bar{x}_{vw}$ for $v \in V_0$. But for $v \in V_0$, we have $v \notin V_1$ and thus

$$[\![\neg V_1\, v \,\vee \dots]\!]_\ell = [\![\neg V_1 v]\!]_\ell + [\![\dots]\!]_\ell = 1 + [\![\dots]\!]_\ell = 1$$

That is, the value of the subformula involving the literals $\neg Exy$ is absorbed. $\square$

The monomial corresponding to a strategy can be defined directly: Given a positional strategy $\mathcal{S}$, the *strategy monomial* $m_{\mathcal{S}}$ is defined as

$$m_{\mathcal{S}} = \{x_{vw} \mid v \in V_0 \text{ and } (v, w) \text{ appears in } \mathcal{S} \text{ (i.e., } \mathcal{S}(\pi v) = \pi v w \text{ for some } \pi)\}.$$

We first make a general observation about strategy monomials of positional strategies.

Lemma 7.7. *Let $\mathcal{S}_1$ and $\mathcal{S}_2$ be two different positional winning strategies from u. Then $m_{\mathcal{S}_1}$ and $m_{\mathcal{S}_2}$ are incomparable (none absorbs the other).*

Proof. Let $m_1 = m_{\mathcal{S}_1}$ and $m_2 = m_{\mathcal{S}_2}$ and assume that $m_1 < m_2$, i.e., $m_2 \subset m_1$. Consider the subgame $\mathcal{G}_2$ which results from $\mathcal{G}$ by removing all edges $(v, w) \in E$ with $v \in V_0$ for which $x_{vw} \notin m_2$. Further consider the component of u in $\mathcal{G}_2$. For each node $v \in V_0$ in this component, there is a unique w such that $x_{vw} \in m_2$ because $\mathcal{S}_2$ is a positional strategy.

Let $\mathcal{G}_1$ be the subgame defined in the same way for m_1. Because of $m_2 \subset m_1$, all edges of $\mathcal{G}_2$ are also contained in $\mathcal{G}_1$. Since $\mathcal{S}_1$ is positional, this means that $\mathcal{G}_1$ and $\mathcal{G}_2$ coincide on the component of u.

Due to $m_2 \subset m_1$, there must further be a variable $x_{vw} \in m_1 \setminus m_2$. The corresponding edge (v, w) of $\mathcal{G}_1$ is not part of $\mathcal{G}_2$ and can thus not lie in the component of u. But then v is not reachable in $\mathcal{G}_1$ and can thus not occur in $\mathcal{S}_1$, which is a contradiction. □

We further need two main steps that relate positional strategies and monomials in $[\![\varphi(u)]\!]_\ell$. The first one is that winning strategies induce monomials of $[\![\varphi(u)]\!]_\ell$.

Proposition 7.8. *If $\mathcal{S}$ is a positional winning strategy for player 0 from v in $\mathcal{G}$, then $m_{\mathcal{S}} \le [\![\varphi(u)]\!]_\ell$. That is, there is a monomial $m \in [\![\varphi(u)]\!]_\ell$ with $m_{\mathcal{S}} \le m$.*

Proof. Consider the subgame $\mathcal{G}'$ which results from $\mathcal{G}$ by removing all edges $(v, w) \in E$ with $v \in V_0$ which are not used in $\mathcal{S}$ (so $x_{vw} \notin m_{\mathcal{S}}$). Then $\mathcal{G}'$ is compatible with ℓ, so there is an assignment $f : X \cup \bar{X} \to \mathbb{B}$ which lifts to a homomorphism $h : \mathsf{PosBool}[X, \bar{X}] \to \mathbb{B}$ such that $h \circ \ell = \ell_{\mathcal{G}'}$.

Since $\mathcal{S}$ is a winning strategy in $\mathcal{G}'$, we know that $\mathcal{G}' \models \varphi(u)$ and thus $h([\![\varphi(u)]\!]_\ell) = [\![\varphi(u)]\!]_{h \circ \ell} = \top$. Then there must be a monomial $m \in [\![\varphi(u)]\!]_\ell$ with $h(m) = \top$. This means that for all variables $x_{vw} \in m$, we must have $h(x_{vw}) = \top$. But by definition of $\mathcal{G}'$ and h, this only holds for variables $x_{vw} \in m_{\mathcal{S}}$. Hence $m \subseteq m_{\mathcal{S}}$ and thus $m_{\mathcal{S}} \le m$. □

Conversely, monomials of $[\![\varphi(u)]\!]_\ell$ induce positional winning strategies from u by applying the positional determinacy of parity games.

Proposition 7.9. *Let $m \in [\![\varphi(u)]\!]_\ell$. Then there is a positional winning strategy $\mathcal{S}$ for player 0 from u with $m \leq m_\mathcal{S}$.*

Proof. Consider the homomorphism $h : \mathsf{PosBool}[X, \bar{X}] \to \mathbb{B}$ induced by m, i.e., $h(x_{vw}) = \top$ if $x_{vw} \in m$ and $h(x_{vw}) = \bot$ otherwise. Then $h(m) = \top$ and thus $[\![\varphi(u)]\!]_{h \circ \ell} = \top$.

Note that $h \circ \ell$ is model-defining. The induced model $\mathcal{G}'$ is the subgame which results from $\mathcal{G}$ by removing all edges (v, w) for which $x_{vw} \notin m$. Because $[\![\varphi(u)]\!]_{h \circ \ell} = \top$, we have $\mathcal{G}' \models \varphi(u)$. Hence u is winning for player 0, so by positional determinacy there is a positional winning strategy $\mathcal{S}$ from u in $\mathcal{G}'$. Because $\mathcal{S}$ can only use the edges of $\mathcal{G}'$, it follows that $m_\mathcal{S} \subseteq m$ and thus $m \leq m_\mathcal{S}$. □

Combining these observations indeed shows that the monomials in $[\![\varphi(u)]\!]_\ell$ precisely correspond to the positional winning strategies from u in $\mathcal{G}$, as seen in figure 7. While there are certainly other ways to compute all winning strategies, this illustrates that semiring provenance can provide useful information and that $\mathsf{PosBool}[X, \bar{X}]$ is already sufficient if we only need to know whether tracked literals are used in a proof (or a positional strategy) or not. We further see that the universality results are of great importance as they connect model-compatible interpretations (such as ℓ above) and actual models (such as the subgames $\mathcal{G}'$ used in the proofs).

Corollary 7.10. *The monomials in $[\![\varphi(u)]\!]_\ell$ are in one-to-one correspondence with the positional winning strategies for player 0 from position u in the game $\mathcal{G}$:*

$$[\![\varphi(u)]\!]_\ell = \{m_\mathcal{S} \mid \mathcal{S} \textit{ is a positional winning strategy from } u\}$$

Proof. Let $\mathcal{S}$ be a positional winning strategy. Then there is $m' \in [\![\varphi(u)]\!]_\ell$ with $m_\mathcal{S} \leq m'$. Considering m', there is a positional winning strategy $\mathcal{S}'$ with $m' \leq m_{\mathcal{S}'}$. Lemma 7.7 implies that $\mathcal{S} = \mathcal{S}'$ and hence $m_\mathcal{S} = m' \in [\![\varphi(u)]\!]_\ell$.

Conversely, let $m \in [\![\varphi(u)]\!]_\ell$. Then there is a positional winning strategy $\mathcal{S}$ with $m \leq m_\mathcal{S}$. We have already shown that $m_\mathcal{S} \in [\![\varphi(u)]\!]_\ell$. But this implies $m = m_\mathcal{S}$, as otherwise m would be absorbed by $m_\mathcal{S}$. □

In this excursion, we have seen how to compute all positional winning strategies via semiring provenance, which in particular allows us to count the number of these strategies. In [GT19], it is shown that counting positional winning strategies (of acyclic reachability games) using K-valuations is impossible. These valuations map nodes to semiring values and are induced by annotations of both terminal positions and edges. K-valuations are preserved by *counting bisimulations* that respect these annotations, which means that a game and its tree unraveling yield the same K-valuation. We instead use S-interpretations to track only the edges (of one player) and we use different variables for every edge. For the tree unraveling (of an acyclic reachability game), this means that we use *different variables* for all edges in the tree, even if they correspond to the *same edge* in the game graph. For this reason, their argument (based on bisimulations that respect the annotation of edges) does not apply to our setting.

7.3 Model Checking Games

Coming back from our excursion, we now want to turn the tables: Instead of using semiring provenance to compute strategies, we want to use strategies to characterize semiring semantics $[\![\varphi]\!]_\ell$. To achieve this, we consider strategies in model checking games. In LFP, these are parity games (we again refer to [GKL+07] for more details), as the parity condition can be used to encode the semantics of least and greatest fixed points (and especially their alternation).

We first define the classical model checking game for a model $\mathfrak{A}$ (with universe A) and a formula φ. In the context of semirings semantics, we have a formula φ and a universe A – but no model. We thus define the notion of *generic* model[10] checking games. Intuitively, these are games in which the interpretation of literals is not yet fixed. Instead, we fix it afterwards (in the next section) by applying a semiring interpretation ℓ to *strategies* of the generic model checking game.

For technical reasons, we assume the formula φ to be *well-named*. That is, all fixed-point subformulae use different relation symbols and these symbols do not occur outside of the scope of the respective subformula. Recall that we write $[\mathsf{fp}\, R\,\mathbf{x}.\ \vartheta](\mathbf{y})$ for either $[\mathsf{lfp}\, R\,\mathbf{x}.\ \vartheta](\mathbf{y})$ or $[\mathsf{gfp}\, R\,\mathbf{x}.\ \vartheta](\mathbf{y})$.

Definition 7.11. Let φ be a well-named LFP formula in negation normal form over the finite relational signature τ. Let $\mathfrak{A}$ be a τ-model with universe A. The *(classical) model checking game* for $\mathfrak{A}$ and φ is the parity game $\mathcal{G}(\mathfrak{A},\varphi) = (V, V_0, V_1, E, \Omega)$ defined as follows. We call player 0 the *Verifier* and player 1 the *Falsifier*.

The positions are all subformulae ψ in which all free variables are instantiated by elements of the universe A, so $V = \{\psi(\mathbf{a}) \mid \psi(\mathbf{x}) \text{ is a subformula of } \varphi\}$.

We define the the sets V_0, V_1 and E for each position $\psi(\mathbf{a})$ depending on ψ:

- If $\psi = R\mathbf{x}$ belonging to a fixed-point formula $[\mathsf{fp}\, R\,\mathbf{x}.\ \vartheta](\mathbf{y})$, then $R\mathbf{a} \in V_0$ and the only outgoing edge is $R\mathbf{a} \to \vartheta(\mathbf{a})$.
- If $\psi \in \mathsf{Lit}_A$ (not belonging to a fixed-point formula), we consult the model: If $\mathfrak{A} \models \psi(\mathbf{a})$, then $\psi(\mathbf{a}) \in V_1$ (Verifier wins), otherwise $\psi(\mathbf{a}) \in V_0$.

[10] Truth be told, these are actually generic checking games, as they are independent of any models.

- If $\psi = \psi_1 \vee \psi_2$, then $\psi(\mathbf{a}) \in V_0$ with edges $\psi(\mathbf{a}) \to \psi_1(\mathbf{a})$ and $\psi(\mathbf{a}) \to \psi_2(\mathbf{a})$.
  If $\psi = \psi_1 \wedge \psi_2$, then $\psi(\mathbf{a}) \in V_1$ with edges $\psi(\mathbf{a}) \to \psi_1(\mathbf{a})$ and $\psi(\mathbf{a}) \to \psi_2(\mathbf{a})$.
- If $\psi = \exists y\, \vartheta(y)$, then $\psi(\mathbf{a}) \in V_0$ with edges $\psi(\mathbf{a}) \to \vartheta(\mathbf{a}, b)$ for each $b \in A$.
  If $\psi = \forall y\, \vartheta(y)$, then $\psi(\mathbf{a}) \in V_1$ with edges $\psi(\mathbf{a}) \to \vartheta(\mathbf{a}, b)$ for each $b \in A$.
- If $\psi = [\mathsf{fp}\, R\, \mathbf{x}.\, \vartheta](\mathbf{y})$, then $\psi(\mathbf{a}) \in V_0$. For $\psi(\mathbf{a}) = [\mathsf{fp}\, R\, \mathbf{x}.\, \vartheta](\mathbf{b})$, we add the single edge $\psi(\mathbf{a}) \to \vartheta(\mathbf{b})$.

The priorities are defined for nodes $R\mathbf{a}$ belonging to fixed-point formulae such that,

(1) if $R\mathbf{a}$ belongs to an lfp-formula, then $\Omega(R\mathbf{a})$ is odd,

(2) if $R\mathbf{a}$ belongs to a gfp-formula, then $\Omega(R\mathbf{a})$ is even,

(3) if $R\mathbf{a}$ belongs to $\psi_R = [\mathsf{fp}\, R\, \mathbf{x}.\, \vartheta](\mathbf{y})$ and $P\mathbf{b}$ belongs to $\psi_P = [\mathsf{fp}\, P\, \mathbf{x}.\, \vartheta'](\mathbf{y})$ where ψ_P is a subformula of ϑ, then $\Omega(R\mathbf{a}) \leq \Omega(P\mathbf{b})$.

That is, outermost fixed-point formulae get the smallest priorities. We want the priorities of all other positions to be irrelevant. If d is the highest priority used above, we thus label all other positions with $d + 1$.

Winning strategies for the game $\mathcal{G}(\mathfrak{A}, \varphi)$ constitute proofs of $\mathfrak{A} \models \varphi$ which leads to the following central property (see [GKL+07] for a proof).

Theorem 7.12. *Let φ and $\mathfrak{A}$ be as in the above definition. Then $\mathfrak{A} \models \varphi(\mathbf{a})$ if, and only if, Verifier has a winning strategy in $\mathcal{G}(\mathfrak{A}, \varphi)$ from position $\varphi(\mathbf{a})$.*

In the context of semiring semantics, we do not have a model $\mathfrak{A}$ but instead only fix the universe A and an S-interpretation ℓ. We separate the two and first define *generic* model checking games in which all terminal positions are won by Verifier. We later re-introduce ℓ when we define the *outcome* of (winning) strategies.

Definition 7.13. Let φ be a well-named LFP formula in negation normal form over the finite relational signature τ and let A be a finite universe. The *(generic) model checking game* for φ (over A and τ) is the parity game $\mathcal{G}_A(\varphi)$ defined exactly as $\mathcal{G}(\mathfrak{A}, \varphi)$ except for its terminal positions (other positions, edges and priorities are as in $\mathcal{G}(\mathfrak{A}, \varphi)$):

- If $\psi \in \mathsf{Lit}_A$ (not belonging to a fixed-point formula), then $\psi(\mathbf{a}) \in V_1$
  (we consider all terminal positions to be winning for Verifier)

We write $\mathcal{G}_A(\varphi(\mathbf{a}))$ for the game $\mathcal{G}_A(\varphi)$ with distinguished initial position $\varphi(\mathbf{a})$ and where we only keep positions reachable from $\varphi(\mathbf{a})$. Strategies for $\mathcal{G}_A(\varphi(\mathbf{a}))$ are always strategies from the initial position.

$$\varphi(x) = \Big[\mathsf{gfp}\, X\, x.\ \big[\mathsf{lfp}\, Y\, x.\ \exists y\big(Exy \wedge ((Xy \wedge Py) \vee Yy)\big)\big](x)\Big](x)$$

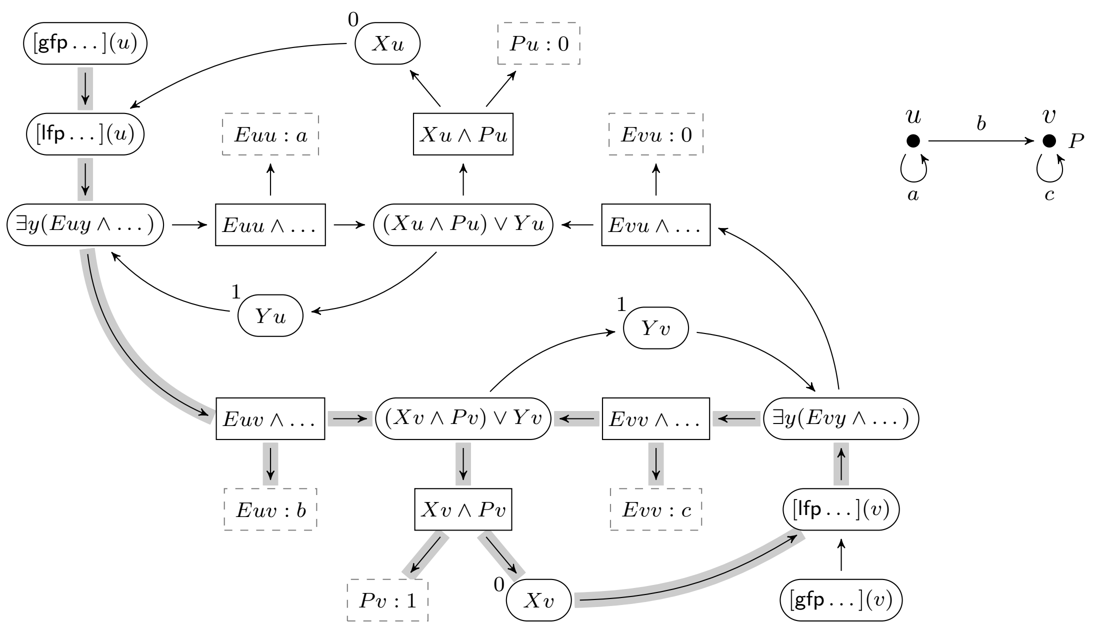

Figure 8: *Example of a model checking game $\mathcal{G}_A(\varphi)$ for the given formula φ over $A = \{u, v\}$. Terminal positions (dashed border) include the value assigned by ℓ, if ℓ is defined as illustrated on the right. Only the relevant priorities are shown. The highlighted edges indicate a winning strategy with non-zero outcome.*

We call such a model checking game *generic* as it is the same for all models with universe A (and signature τ). An example together with an S-interpretation for the terminal positions is shown in figure 8. Our goal is to use generic model checking games together with interpretations ℓ to describe the result of a provenance computation $[\![\varphi]\!]_\ell$ by means of the values of literals appearing in winning strategies for the game $\mathcal{G}(\varphi)$.

Before we proceed, we need further notation. Recall that we always assume a finite universe A. We write $\mathcal{G}(\varphi)$ for the game $\mathcal{G}_A(\varphi)$ if A is clear from the context. When discussing strategies, we always take the perspective of Verifier, so winning means winning for Verifier. The set of all winning strategies of $\mathcal{G}(\varphi(\mathbf{a}))$, which are the winning strategies from position $\varphi(\mathbf{a})$ in $\mathcal{G}(\varphi)$, is denoted by $\mathcal{W}_{\varphi(\mathbf{a})}$. All strategies we consider in the following are winning strategies which we usually denote by the letter $\mathcal{S}$. Given $\mathcal{S} = (W, F)$, we often write $\pi \in \mathcal{S}$ instead of $\pi \in W$.

The following observations are used frequently in proofs and are clear from the definition of strategies and the winning condition.

Lemma 7.14. *Consider $\mathcal{G}(\varphi) = (V, V_0, V_1, E, \Omega)$ for some formula φ.*

(1) *Let $\rho = \rho_0\, \rho'$ be a play in $\mathcal{G}(\varphi)$ where ρ_0 is a finite prefix and ρ' is not empty. Then ρ is winning if, and only if, ρ' is winning.*

(2) *Let $\mathcal{S}$ be a (winning) strategy in $\mathcal{G}(\varphi)$ from some position $v \in V$ and let $\pi \in \mathcal{S}$. Then the subtree of $\mathcal{S}$ rooted at π (i.e., the subgraph induced by all nodes $\pi' \in \mathcal{S}$ with $\pi \sqsubseteq \pi'$) is a (winning) strategy from $V(\pi)$.*

We are especially concerned with literals appearing in strategies. Recall that Lit_A is the finite set of literals (over A and τ). For a strategy $\mathcal{S}$ and a literal $L \in \mathsf{Lit}_A$, we write $|\mathcal{S}|_L$ for the number occurrences of L in leaves of $\mathcal{S}$, or, formally:

Definition 7.15. Let φ be a formula, $\mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}$ and $L \in \mathsf{Lit}_A$. Moreover, let

$$W_L = \{\pi L \in \mathcal{S} \mid L \text{ is a terminal position in } \mathcal{G}(\varphi)\}$$

That is, W_L is the set of leaves of $\mathcal{S}$ with position L. We then define

$$|\mathcal{S}|_L = \begin{cases} |W_L|, & \text{if } W_L \text{ is finite} \\ \infty, & \text{otherwise} \end{cases}$$

Note that we always have $|\mathcal{S}|_{R\mathbf{a}} = 0$ if R belongs to a fixed-point formula $[\mathsf{fp}\, R\mathbf{x}.\, \vartheta](\mathbf{y})$, as the positions $R\mathbf{a}$ have outgoing edges (to $\vartheta(\mathbf{a})$) in the model checking game and are thus not terminal (and therefore no leaves of $\mathcal{S}$).

The characterization we want to achieve is of the form

$$[\![\varphi(\mathbf{a})]\!]_\ell \;=\; \sum_{\mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}} [\![\mathcal{S}]\!]_\ell$$

where the *outcome* $[\![\mathcal{S}]\!]_\ell$ of the strategy $\mathcal{S}$ should intuitively be the product of all literals (or rather, their values under ℓ) appearing in $\mathcal{S}$. Note that this is not yet well-defined, as both the number of winning strategies and the number of literals in $\mathcal{S}$ can be infinite. We consider this characterization for the important class of absorptive continuous semirings where it can be phrased in a reasonable way (mostly due to the infinitary power a^∞ and because $\top = 1$), but we believe that similar statements are possible also for other semirings such as $\mathbb{N}^\infty$ and $\mathbb{W}[X]$.

Recall that we can define $a^\infty = \bigsqcap_n a^n$ in all absorptive semirings S. Given a strategy $\mathcal{S}$, we define its *outcome* $[\![\mathcal{S}]\!]_\ell$ under an S-interpretation ℓ as the product of all literals appearing in leaves of $\mathcal{S}$. The infinitary power is needed in case of $|\mathcal{S}|_L = \infty$.

Definition 7.16. Let $\mathcal{S}$ be a strategy in the game $\mathcal{G}(\varphi)$ for some formula φ and let ℓ be an S-interpretation for an absorptive semiring S. The *outcome* of $\mathcal{S}$ is the product

$$[\![\mathcal{S}]\!]_\ell = \prod_{L \in \mathsf{Lit}_A} \ell(L)^{|\mathcal{S}|_L}$$

If ℓ is model-defining, we can identify winning strategies in $\mathcal{G}(\varphi)$ with non-zero outcome with winning strategies in the classical model checking game $\mathcal{G}(\mathfrak{A}_\ell, \varphi)$. We must additionally impose a positivity requirement on the semiring to ensure that the outcome can only be zero if $\mathcal{S}$ contains a zero-valued literal.

Proposition 7.17. *Let S be a positive, chain-positive and absorptive semiring, let ℓ be a model-defining S-interpretation and φ a formula. Then the winning strategies $\mathcal{S}$ in the game $\mathcal{G}(\varphi(\mathbf{a}))$ with $[\![\mathcal{S}]\!]_\ell \neq 0$ are precisely the winning strategies from position $\varphi(\mathbf{a})$ in the classical model checking game $\mathcal{G}(\mathfrak{A}_\ell, \varphi)$.*

Proof. The games $\mathcal{G}(\varphi(\mathbf{a}))$ and $\mathcal{G}(\mathfrak{A}_\ell, \varphi)$ are equivalent except for the literals. In $\mathcal{G}(\varphi(\mathbf{a}))$, all literals are considered to be winning (for Verifier), whereas in $\mathcal{G}(\mathfrak{A}_\ell, \varphi)$, only the literals L with $\mathfrak{A}_\ell \models L$ are winning. As the model $\mathfrak{A}_\ell$ is defined by ℓ, literals with $\mathfrak{A}_\ell \not\models L$ are annotated by $\ell(L) = 0$. Hence strategies visiting L have the outcome 0.

For the formal argument, let $\mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}$. Due to the positivity and chain-positivity of S, we have $\ell(L)^{|\mathcal{S}|_L} = 0$ if, and only if, $\ell(L) = 0$ and $|\mathcal{S}|_L > 0$. By definition of the outcome, this implies that $[\![\mathcal{S}]\!]_\ell \neq 0$ if, and only if, $|\mathcal{S}|_L = 0$ for all literals L with $\ell(L) = 0$, i.e., with $\mathfrak{A}_\ell \not\models L$. It follows that $\mathcal{S}$ is a winning strategy in $\mathcal{G}(\mathfrak{A}_\ell, \varphi)$ if, and only if, $[\![\mathcal{S}]\!]_\ell \neq 0$. □

This observation makes clear that the characterization of $[\![\varphi(\mathbf{a})]\!]_\ell$ is all about winning strategies in the classical sense and can thus be seen as a description of semiring semantics in terms of proofs of the formula $\varphi(\mathbf{a})$. In order to establish this result for all absorptive continuous semirings, we show that it holds for absorptive polynomials $\mathbb{S}^\infty[X]$ and make use of their universality.

7.4 Characterization for Absorptive Polynomials

In the semiring $\mathbb{S}^\infty[X]$ of absorptive polynomials, we can express the sum over all strategies as supremum (recall that $\mathbb{S}^\infty[X]$ is idempotent and a lattice semiring). The characterization we want to prove can then be formulated as follows. The rest of this section is devoted to the proof and consequences of this result.

Theorem 7.18 (strategy characterization)**.** *Let ℓ be an $\mathbb{S}^\infty[X]$-interpretation and $\varphi(\mathbf{x})$ a formula. Then*

$$[\![\varphi(\mathbf{a})]\!]_\ell = \bigsqcup\{[\![\mathcal{S}]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\}$$

Absorption on Strategies

Given two strategies $\mathcal{S}_1$ and $\mathcal{S}_2$ in a game $\mathcal{G}(\varphi(\mathbf{a}))$ and an $\mathbb{S}^\infty[X]$-interpretation ℓ, we say that $\mathcal{S}_1 \leq \mathcal{S}_2$ (that is, $\mathcal{S}_2$ *absorbs* $\mathcal{S}_1$) if $[\![\mathcal{S}_1]\!]_\ell \leq [\![\mathcal{S}_2]\!]_\ell$. When we consider the supremum $\bigsqcup\{[\![\mathcal{S}]\!]_\ell \mid \mathcal{S} \in \mathcal{G}(\varphi(\mathbf{a}))\}$, only the absorption-maximal strategies are relevant. An interesting question is what kind of strategies these are. Thinking of examples such as graph reachability, it is clear that absorption-maximal strategies must not correspond to paths with unnecessary repetitions. This suggests that absorption-maximal strategies are related to positional strategies. While this is almost true, consider this example:

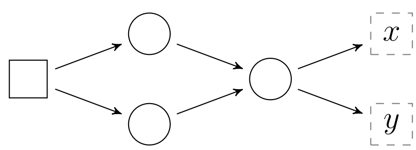

This game is a simple reachability game (without cycles), yet the only two positional strategies have the outcomes x^2 and y^2. So although the outcome xy is absorption-maximal, the corresponding strategy is not positional. It is, however, *weakly positional*, as it only makes different decisions in different plays. This example is also considered in [GT19] where it is claimed that the restriction to weakly positional strategies does not reduce strategic power in reachability games with cycles. We show that the same holds for the larger class of parity games.

This observation is applicable not only to model checking games, but to finite parity games with labeled terminal positions in general. Given a finite parity game $\mathcal{G} = (V, V_0, V_1, E, \Omega)$, let $T = \{v \in V \mid vE = \emptyset\}$ be the set of terminal positions (of either player) and consider a labeling $\ell : T \to \mathbb{S}^\infty[X]$. We can adapt the definition of outcomes accordingly:

$$[\![\mathcal{S}]\!]_\ell = \prod_{t \in T} \ell(t)^{|\mathcal{S}|_t}$$

where $\mathcal{S}$ is a strategy in $\mathcal{G}$ and $|\mathcal{S}|_t$ is the number of occurrences (possibly ∞) of the terminal position t. If $\mathcal{G} = \mathcal{G}(\varphi)$, then this definition coincides with the original one.

Proposition 7.19. *Let $\mathcal{G}$, T and ℓ be defined as above and let $\mathcal{S}$ be a any winning strategy (for player 0) from position $v \in V$. Then there is a weakly positional winning strategy $\mathcal{S}'$ (for player 0) from v such that $\mathcal{S} \leq \mathcal{S}'$.*

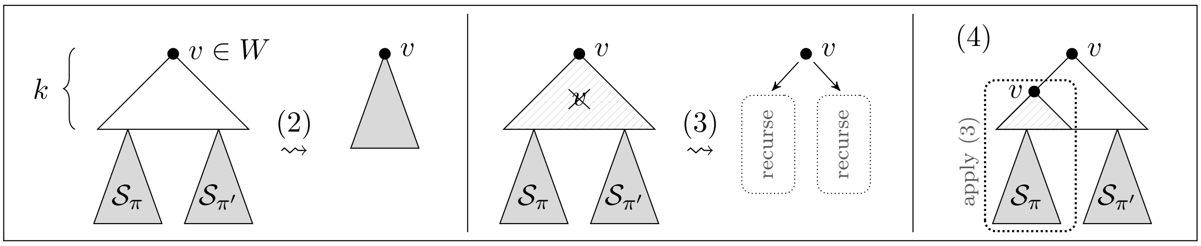

Figure 9: *Illustration of the recursive construction in the proof of proposition 7.19. Grey subtrees are positional winning strategies in $\mathcal{G}'$, striped parts do not contain v.*

Proof. We show how to construct a weakly positional strategy $\mathcal{S}'$ such that $|\mathcal{S}|_t \geq |\mathcal{S}'|_t$ for all $t \in T$, since this implies $\mathcal{S} \leq \mathcal{S}'$. In particular, $\mathcal{S}$ is winning and thus visits no terminals $t \in T \cap V_0$, so the same must hold for $\mathcal{S}'$.

Let $T_\infty = \{t \in T \mid |\mathcal{S}|_t = \infty\}$ be the set of terminals appearing infinitely often in $\mathcal{S}$. These impose no restriction, as we always have $|\mathcal{S}|_t \geq |\mathcal{S}'|_t$ for $t \in T_\infty$. We are thus more concerned with the literals $T_{\text{fin}} = T \setminus T_\infty$ which we call *problematic*.

First consider the parity game $\mathcal{G}' = (V, V_0', V_1', E, \Omega)$ with $V_0' = (V_0 \setminus T) \cup T_{\text{fin}}$ and $V_1' = (V_1 \setminus T) \cup T_\infty$. That is, we consider all terminals in T_∞ to be winning and all problematic terminals to be losing (from the perspective of player 0). Let $W \subseteq V$ be the set of all positions from which there is a positional winning strategy in $\mathcal{G}'$.

Let further $n = \sum_{t \in T_{\text{fin}}} |\mathcal{S}|_t$. Then n is finite (note that T is finite), so there is a depth k such that no more occurrences of T_{fin} happen below depth k (that is, for all $\pi \in \mathcal{S}$ with $|\pi| \geq k$ we have $V(\pi) \notin T_{\text{fin}}$). Consider any node $\pi \in \mathcal{S}$ with $|\pi| = k$. Then the subtree of $\mathcal{S}$ rooted at π is a winning strategy in $\mathcal{G}'$ from $V(\pi)$ (because it avoids T_{fin}). Due to positional determinacy, there is a positional winning strategy $\mathcal{S}_\pi$ from $V(\pi)$ in $\mathcal{G}'$.

Let $\mathcal{T}$ result from $\mathcal{S}$ by replacing all $\pi \in \mathcal{S}$ with $|\pi| = k$ by the these strategies $\mathcal{S}_\pi$. We describe a recursive process to transform $\mathcal{T}$ into a weakly positional strategy $\mathcal{T}'$.

(1) Let $v \in V$ be the position of the root of $\mathcal{T}$.

(2) If $v \in W$, then there is a positional winning strategy $\mathcal{T}'$ from v in $\mathcal{G}'$. By definition of $\mathcal{G}'$, we have $|\mathcal{S}|_t \geq |\mathcal{T}'|_t$ for all $t \in T$, so we stop and return $\mathcal{T}'$.

(3) If v does not occur in $\mathcal{T}$ below the root, then $\mathcal{T}$ is already weakly positional with respect to position v. Let $\mathcal{T}_1, \ldots, \mathcal{T}_n$ be the subtrees rooted in the children of $\mathcal{T}$'s root. We obtain weakly positional versions $\mathcal{T}_1', \ldots, \mathcal{T}_n'$ by recursion. We then replace each subtree $\mathcal{T}_i$ by $\mathcal{T}_i'$ and return the result.

(4) Otherwise, v does occur somewhere else in $\mathcal{T}$. Note that v cannot occur in the strategies $\mathcal{S}_\pi$, as we would then have $v \in W$. Hence all occurrences of v occur in a finite prefix of $\mathcal{T}$. There must thus be a node $\pi \in \mathcal{T}$ with $V(\pi) = v$ such that the

subtree $\mathcal{T}'$ rooted at π contains no further occurrences of v. We then proceed with $\mathcal{T}'$ as in step (3) and return the result.

Let $\mathcal{S}'$ be the overall result of applying these steps to $\mathcal{T}$. Note that the process terminates as we always recurse on subtrees and step (2) applies when we reach one of the subtrees $\mathcal{S}_\pi$. Moreover, $\mathcal{S}'$ is weakly positional and winning: Each sufficiently long play eventually enters one of the subtrees $\mathcal{S}_\pi$. So each play has the form $\rho = \rho_0\,\rho'$ where ρ' is the (possibly empty) play through $\mathcal{S}_\pi$. Steps (3) and (4) ensure that each position in ρ_0 occurs only once in the entire play ρ. The remaining play ρ' follows the positional strategy $\mathcal{S}_\pi$. Hence ρ is positional and $\mathcal{S}'$ is thus weakly positional. If ρ' is empty, then the play ends in a terminal. Assuming $|\mathcal{S}|_t \geq |\mathcal{S}'|_t$ for all $t \in T$, this terminal position must be winning. If ρ' is nonempty, then ρ' (and thus ρ) is winning due to $\mathcal{S}_\pi$.

It remains to show that $|\mathcal{S}|_t \geq |\mathcal{S}'|_t$ for all $t \in T$. First note that by definition of $\mathcal{G}'$, the strategies $\mathcal{S}_\pi$ only visit literals in T_∞. Hence $|\mathcal{S}|_t \geq |\mathcal{T}|_t$ for all $t \in T$. The recursive construction only removes parts of $\mathcal{T}$ or replaces subtrees with positional winning strategies for $\mathcal{G}'$ which do not visit T_{fin}. Hence $\mathcal{S}'$ contains fewer (or equally many) problematic literals than $\mathcal{T}$ and thus $|\mathcal{S}|_t \geq |\mathcal{T}|_t \geq |\mathcal{S}'|_t$ for all $t \in T$. □

Towards the Characterization: Strategy Truncations

The main idea to prove the characterization is to relate the fixed-point iteration with prefixes of the strategy $\mathcal{S}$. That is, we cut off the tree $\mathcal{S}$ after n occurrences of the relation symbol R belonging to the fixed-point formula. The outcome of the resulting subtree of $\mathcal{S}$ then corresponds to the n-th step of the fixed-point iteration.

Definition 7.20. Let $\mathcal{S} = (W, F)$ be a strategy in $\mathcal{G}(\varphi) = (V, V_0, V_1, E, \Omega)$ and let R be a relation symbol of arity r. For $\pi = v_0 v_1 \dots v_k \in W$, we define

$$|\pi|_R = \big|\{i \mid v_i = R\mathbf{a} \text{ for some } \mathbf{a} \in A^r\}\big|$$

The (R, n)-*truncation* of $\mathcal{S}$ is the tree (W', F') defined as follows. Its nodes are finite sequences over $V \cup \{✂\}$, where ✂ is a special symbol which marks the nodes at which we cut off subtrees of $\mathcal{S}$. For $n \geq 1$, we define

- $W' = \{\pi \in W \mid |\pi|_R < n\} \;\cup\; \{\pi ✂ \mid \pi\, R\mathbf{a} \in W,\ |\pi|_R = n-1\}$,
- $F' = F \cap (W' \times W') \;\cup\; \{(\pi, \pi ✂) \mid \pi ✂ \in W'\}$.

For $n = 0$, we instead set $W' = \{✂\}$ and $F' = \emptyset$. If R is clear from the context, we write $\mathcal{S}|_n$ for the (R, n)-truncation of $\mathcal{S}$.

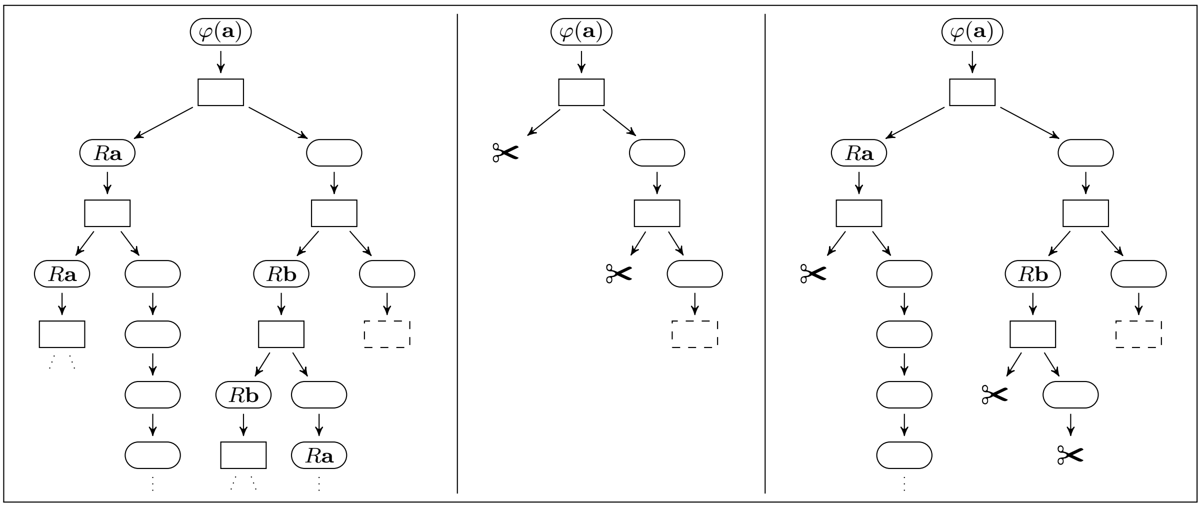

Figure 10: *A visualization of a strategy $\mathcal{S}$ and its (R, n)-truncations for $n = 1, 2$.*

The idea of strategy truncations is roughly comparable to the unfolding of parity games in [GKL+07, chapter 3]. However, we need more information than just winning or non-winning, so we have to work with strategies instead of modifying the game.

See figure 10 for an example. We lift the definition of the outcome $[\![\mathcal{S}]\!]_\ell$ to truncations $[\![\mathcal{S}|_n]\!]_\ell$ by treating ✂ as an additional literal. That is,

$$[\![\mathcal{S}|_n]\!]_\ell = \ell(\text{✂})^{|\mathcal{S}|_{\text{✂}}} \cdot \prod_{L \in \mathsf{Lit}_A} \ell(L)^{|\mathcal{S}|_L}$$

As a first observation, we note that for a fixed strategy $\mathcal{S}$, the outcomes of its truncations converge towards the outcome of $\mathcal{S}$.

Lemma 7.21. *Let ℓ be an $\mathbb{S}^\infty[X]$-interpretation.*

(1) *Let $\varphi = [\mathsf{lfp}\, R\mathbf{x}.\, \vartheta](\mathbf{y})$ and let $\mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}$. If we extend ℓ by $\ell(\text{✂}) = 0$, then*

$$\bigsqcup_{n<\omega} [\![\mathcal{S}|_n]\!]_\ell = [\![\mathcal{S}]\!]_\ell$$

(2) *Let $\varphi = [\mathsf{gfp}\, R\mathbf{x}.\, \vartheta](\mathbf{y})$ and let $\mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}$. If we extend ℓ by $\ell(\text{✂}) = 1$, then*

$$\bigsqcap_{n<\omega} [\![\mathcal{S}|_n]\!]_\ell = [\![\mathcal{S}]\!]_\ell$$

Proof. For (1), note that $[\![\mathcal{S}|_n]\!]_\ell = 0$ whenever $\mathcal{S}$ has a path with at least n R-nodes (so $\mathcal{S}|_n$ contains a ✂-node). We show that there is a k such that all paths of $\mathcal{S}$ have less than k R-nodes. Assume towards a contradiction that this is not the case. Then consider the subgraph of $\mathcal{S}$ consisting of all nodes from which a path to an R-node exists. This subgraph must then have paths of arbitrary length and by Kőnigs lemma (note that $\mathcal{S}$ is finitely branching), it must have an infinite path. This is a contradiction, as this infinite path would contain an infinite number of R-nodes and would thus be losing. Hence $[\![\mathcal{S}|_n]\!]_\ell = [\![\mathcal{S}]\!]_\ell$ for all $n \geq k$ and the claim follows.

For (2), we first note that the truncations $[\![\mathcal{S}|_n]\!]_\ell$ indeed form a chain. The reason is that $\ell(\text{✂}) = 1$ is the greatest element, so replacing subtrees of $\mathcal{S}$ by ✂ leads to a larger outcome. Using the splitting lemma, the infimum can be written as follows (we can ignore the value $\ell(\text{✂})$ appearing in the outcome, as 1 is also the neutral element).

$$\bigsqcap_{n<\omega} [\![\mathcal{S}|_n]\!]_\ell = \prod_{L \in \mathsf{Lit}_A} \bigsqcap_{n<\omega} \ell(L)^{|\mathcal{S}|_n|_L} = \prod_{L \in \mathsf{Lit}_A} \ell(L)^{c_L}, \quad \text{where } c_L = \bigsqcup_{n<\omega} \big|\mathcal{S}|_n\big|_L$$

The main observation is that each node of $\mathcal{S}$ is eventually contained in $\mathcal{S}|_n$ (for sufficiently large n). Consider a literal L. If $|\mathcal{S}|_L$ is finite, then for sufficiently large n, we have $\big|\mathcal{S}|_n\big|_L = |\mathcal{S}|_L$ and thus $c_L = |\mathcal{S}|_L$. If $|\mathcal{S}|_L = \infty$, then for each k there is a sufficiently large n such that $\big|\mathcal{S}|_n\big|_L \geq k$ and thus $c_L = \infty$, which closes the proof. ☐

The Puzzle Lemma

As in the universality proof for $\mathbb{S}^\infty[X]$, greatest fixed points are more challenging than least fixed points. To overcome this obstacle, we make use of the characterization of infima from theorem 6.27 by means of monomial chains. In the current setting, the monomials are outcomes of strategy truncations and the following lemma plays a central role in the proof of the characterization theorem.

Lemma 7.22 (puzzle lemma). *Let $\varphi = [\mathsf{gfp}\, R\, \mathbf{x}.\ \vartheta](\mathbf{y})$, let r be the arity of R and let $\mathbf{a} \in A^r$. Let ℓ be an $\mathbb{S}^\infty[X]$-interpretation extended by $\ell(\text{✂}) = 1$. Let further $(\mathcal{S}_i)_{i<\omega}$ be a family of strategies in $\mathcal{W}_{\varphi(\mathbf{a})}$ such that $([\![\mathcal{S}_i|_i]\!]_\ell)_{i<\omega}$ is a descending chain. Then there is a winning strategy $\mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}$ with $[\![\mathcal{S}]\!]_\ell \geq \bigsqcap_i [\![\mathcal{S}_i|_i]\!]_\ell$.*

For an intuition why this result is not obvious, we consider the following example. The key problem is that the strategies $\mathcal{S}_i$ can all be different. In particular, It can happen that for every i, the outcome of the truncation $\mathcal{S}_i|_i$ is larger than the outcome of the full strategy $\mathcal{S}_i$. The insight of the lemma is that we can always use one of the truncations $\mathcal{S}_i|_i$ (for sufficiently large i) to construct a strategy $\mathcal{S}$ with the desired property. This construction has to be done carefully to ensure that the resulting strategy $\mathcal{S}$ is winning.

Example 7.23. Consider the following setting:

$$\varphi_{\text{infpath}}(u) = [\mathsf{gfp}\, R\, x.\ \exists y(Exy \wedge Ry)](u)$$

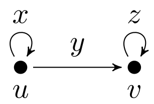

Let $\mathcal{S}_i$ be the strategy corresponding to the infinite path that cycles $i-1$ times via x, then uses edge y and finally cycles via z. The i-truncation then cuts off $\mathcal{S}_i$ after taking the edge y and we obtain the outcomes

$$[\![\mathcal{S}_i]\!]_\ell = x^i y z^\infty \qquad \text{and} \qquad \bigsqcap_{i<\omega} [\![\mathcal{S}_i|_i]\!]_\ell = \bigsqcap_{i<\omega} x^i y = x^\infty y.$$

We see that the infimum only contains the variables x and y, although there is no winning strategy with this outcome. Instead, we obtain $\mathcal{S}$ by repeating the cycling part of any truncation $\mathcal{S}_i|_i$ (without the problematic literal y). This results in the strategy $\mathcal{S}$ with outcome $[\![\mathcal{S}]\!]_\ell = x^\infty$ that corresponds to the path always cycling via x. This path is not consistent with any of the strategies $\mathcal{S}_i$. In general, we have to make sure that the additional plays in $\mathcal{S}$ (which result from the repetition of $\mathcal{S}_i|_i$) are always winning.

As a first step, we can apply the splitting lemma to the infimum and obtain:

$$\bigsqcap_{i<\omega} [\![\mathcal{S}_i|_i]\!]_\ell = \prod_{L \in \mathsf{Lit}_A} \ell(L)^{n_L} \quad \text{with} \quad n_L = \bigsqcup_{i<\omega} \big|\mathcal{S}_i|_i\big|_L$$

Literals with $n_L = \infty$ (such as the edge x in the example) can appear arbitrarily often in $\mathcal{S}$, so they do not impose any restrictions. If $n_L < \infty$, then we must have $|S|_L \leq n_L$ to guarantee that the outcome of $\mathcal{S}$ is larger than the infimum. We therefore call literals L with $n_L < \infty$ (such as the edge y) *problematic*. The outline of the proof is as follows:

- We decompose the trees $\mathcal{S}_i|_i$ into *layers* based on the appearance of R-nodes.
- We choose a sufficiently large i such that there is one such layer in $\mathcal{S}_i|_i$ which does not contain any problematic literals at all.
- We construct $\mathcal{S}$ by first following $\mathcal{S}_i|_i$ and then repeating this layer ad infinitum. For the construction, we collect several subtrees (which we call *puzzle pieces*) from this layer which we can then join together to form the repetition.
- The form of the puzzle pieces ensures that $\mathcal{S}$ is winning. In particular, we only join the pieces at R-nodes. Paths through infinitely many pieces are thus guaranteed to satisfy the parity condition.

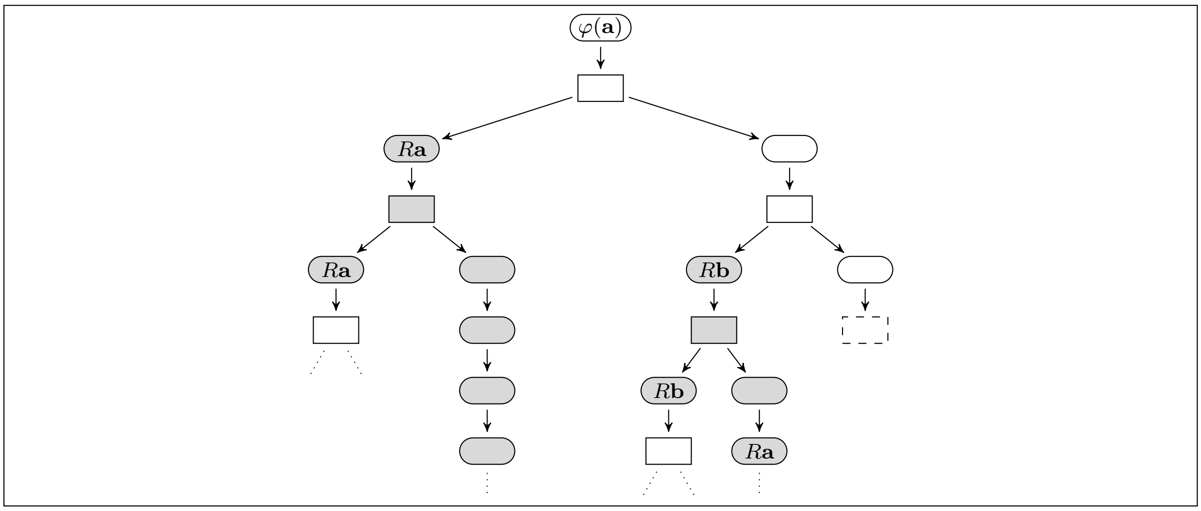

Figure 11: *A visualization of a strategy. The gray nodes form the first layer (for $k = 1$). The two trees in this layer are puzzle pieces, the left one has an infinite winning path.*

Decomposition into layers

Fix an i and let $\mathcal{S}_i|_i = (W, F)$. We call each node $\pi \in W$ with $V(\pi) = R\mathbf{a}$ (for any $\mathbf{a} \in A^r$) an R-node. If an R-node happens to be a leaf, we speak of an R-leaf.

For each $n \geq 0$, we define the sets

$$W_{\leq n} = \{\pi \in W \mid |\pi|_R \leq n\}, \quad W^+_{\leq n} = W_{\leq n} \cup \{\pi v \in W \mid \pi \in W_{\leq n},\, v \in V\}$$

We sort the nodes $\pi \in W$ into layers based on the number of R-nodes on the path to π. For now, think of a layer as a forest in which all roots and most of the leaves are R-nodes. The R-leaves of one layer are the root nodes of the next layer, apart from this layers do not overlap. The constant k controls the thickness of the layer (the maximal number of R-nodes that can occur on paths through the layer).

For any $j \geq 1$, the *j-th layer* is the subgraph of $\mathcal{S}_i|_i$ induced by the node set

$$W_j = W^+_{\leq j \cdot k} \setminus W_{\leq (j-1) \cdot k} \quad \text{where} \quad k = |A|^r + 2.$$

Note that each tree in a layer is a strategy (i.e., satisfies conditions (1)-(3) of definition 7.4) except for its leaves. See figure 11 for a visualization.

Avoiding problematic literals

Let $n = \sum\{n_L \mid L \in \mathsf{Lit}_A,\, n_L < \infty\}$ be the sum of the problematic n_L, which is an upper bound on the number of problematic literals appearing in any truncation $\mathcal{S}_i|_i$. Note that n is always finite. We now choose any i such that:

$$i \;\geq\; (n+1) \cdot k = (n+1) \cdot (|A|^r + 2)$$

From now on, we only work with $\mathcal{S}_i|_i = (W, F)$. Consider the layers $W_1, \ldots, W_{n+1}$ of $\mathcal{S}_i|_i$. First assume that there is a j such that $W_j = \emptyset$. By definition of the layers, we thus have $|\pi|_R \leq (j-1) \cdot k < i$ for all $\pi \in \mathcal{S}_i|_i$. But this means that each path in $\mathcal{S}_i|_i$ has less than i R-nodes. By definition of the truncation, this means that $\mathcal{S}_i|_i = \mathcal{S}_i$. In this case we can simply set $\mathcal{S} = \mathcal{S}_i$ and are done.

Otherwise, there are $n+1$ nonempty layers and at most n occurrences of problematic literals. Hence there must be a layer j such that W_j does not contain any problematic literals. In the following, we concentrate only on this layer W_j.

Collecting puzzle pieces

We want to build the strategy $\mathcal{S}$ from the prefix of $\mathcal{S}_i|_i$ up to layer W_j and then continue by always repeating the layer W_j. Because W_j does not contain any problematic literals, this eventually yields $[\![\mathcal{S}]\!]_\ell \geq [\![\mathcal{S}_i|_i]\!]_\ell$ as required.

Let T be one of the components in W_j, so T is a tree. We call a path in T *winning* if it is infinite or ends in a terminal position, so it corresponds to a play consistent with $\mathcal{S}_i$. Paths ending in R-leaves of T (which could be continued by leaving the layer W_j) are not considered to be winning.

Definition 7.24. A *puzzle piece* $P = (W', F')$ is a subtree of W_j such that

(a) The root of P is an R-node,

(b) For each inner node $\pi \in P$, we have $\pi F' = \pi F$ (P contains all successors),

(c) Each maximal path through P is either winning or ends in an R-node.

A puzzle piece P with root π *matches* a node $\pi' \in W_j$ if $V(\pi) = V(\pi')$. A *complete puzzle* is a set of puzzle pieces such that for each piece in the set and all R-leaves π of this piece, the set contains a puzzle piece that matches π.

First observe that T itself is a puzzle piece: Each maximal path through T which does not end in an R-node must visit less than k R-nodes. If we append this path to the unique path from the root of $\mathcal{S}_i|_i$ to the root of T, then the resulting path contains less than i R-nodes. Hence the path is not truncated in $\mathcal{S}_i|_i$, so it is also a maximal path of $\mathcal{S}_i$ and thus winning. However, a single piece does not make a complete puzzle. Instead, we collect smaller pieces from T by the following algorithm:

(1) Initialize $L = \{\hat{\pi}\}$ where $\hat{\pi}$ is the root of T (which is an R-node).

(2) Pick a node $\pi \in L$ and remove it from L (if L is empty, terminate).

(3) If we have already found a puzzle piece matching π, go to step (2).

(4) Let P be the subgraph of T induced by the following set of nodes. Then P is a puzzle piece matching π and we add it to our set of pieces.

$$W' := \{\pi' \in T \mid \pi \sqsubseteq \pi' \text{ and there is no } R\text{-node } \pi'' \text{ with } \pi \sqsubset \pi'' \sqsubset \pi')\}$$

(5) For each $\mathbf{a} \in A^r$: If P has a leaf π' with $V(\pi') = R\mathbf{a}$, add one such leaf π' to L.

(6) Go back to step (2).

If the definition of P in step (4) is correct, then this algorithm clearly terminates after finding at most $|A|^r$ puzzle pieces and the resulting set of pieces is a complete puzzle. For step (4), recall the definition of W_j. For the root of T, we have $|\hat{\pi}|_R = (j-1)k+1$ and W_j contains in particular the nodes π with $(j-1)k < |\pi|_R \leq jk$.

Assume that in (2), we picked a node π with $|\pi|_R = n$ (for some n). By definition of W', the piece P only contains nodes π' with $|\pi'|_R \leq n+1$. In particular, the leaves that we add to L in step (5) all satisfy $|\pi'|_R \leq n+1$. We start with $|\hat{\pi}|_R = (j-1)k+1$ and perform at most $k-2$ iterations, hence we always have $|\pi|_R < jk$ for all $\pi \in L$.

This guarantees that P is always a puzzle piece in step (4): Inner nodes of P cannot be R-nodes and hence P always contains all successors of inner nodes, so (b) is satisfied. For (c), assume towards a contradiction that there is a maximal path through P which does not end in an R-node and is not winning. This path is also a path in $\mathcal{S}_i$ and because all infinite paths of $\mathcal{S}_i$ are winning, the path must be finite. Because all terminal positions in $\mathcal{S}_i$ are winning, the path must end in a leaf of T which is not a leaf of $\mathcal{S}_i$. But such leaves of T must be R-nodes by definition of the layers, which is a contradiction. Hence (c) holds as well and P is a puzzle piece.

We proceed in the same way for all other components of W_j and obtain a complete puzzle for each component. The overall result is the union of all these puzzles, which is again a complete puzzle. An illustration of such a puzzle (as individual pieces and in assembled form) is shown in figure 12; the next step is to perform the assembly.

Completing the puzzle

We now have a complete puzzle with a matching piece for all root nodes of W_j (these are precisely the R-leaves of the preceding layer W_{j-1}). All that remains is to join the pieces together to form the strategy $\mathcal{S}$. Note that puzzle pieces can contain infinite paths or even infinitely many R-leaves. We therefore construct $\mathcal{S}$ recursively layer by layer.

- S_0 is the subgraph induced by $W^+_{\leq(j-1)k}$, i.e., the prefix of $\mathcal{S}_i|_i$ up to layer W_j. By definition of the layers, all leaves of S_0 are either leaves of $\mathcal{S}_i$ or R-leaves of W_{j-1}.

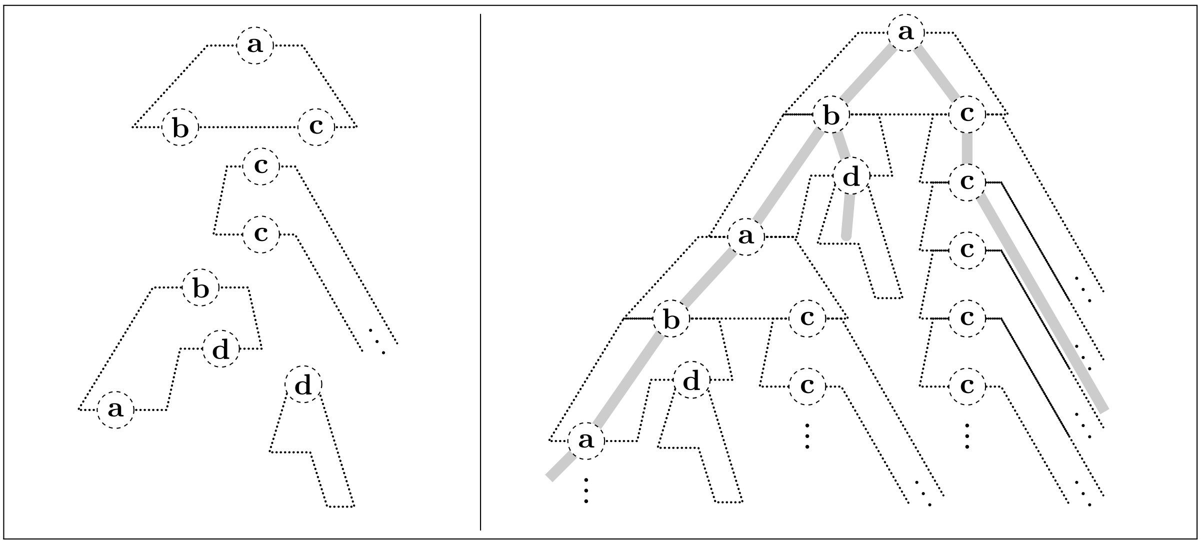

Figure 12: *A schematic illustration of the pieces in a complete puzzle and their infinite repetition. We abbreviate R-nodes R**a** (at which we join pieces) by just **a**. The gray lines indicate three paths: One through infinitely many pieces, a finite one and an infinite one that only visits finitely many pieces (from left to right). All three are winning by construction.*

- Given S_n, we construct S_{n+1} as follows. Recall that for $\pi = v_0 \dots v_l \in S_n$, we write $|\pi| = l$ for its length (which equals the depth of π in S_n). Consider the set

$$X = \{\pi \in S_n \mid \pi \text{ is an } R\text{-leaf of } S_n \text{ with } |\pi| = n\}$$

  Because S_n is finitely branching (as we construct it from subtrees of $\mathcal{S}_i|_i$), this set is finite. Moreover, each $\pi \in X$ is either the R-leaf of a puzzle piece or, initially, the root of one component of W_j. In both cases, the complete puzzle contains a piece matching π. The tree S_{n+1} results from S_n by replacing all leaves $\pi \in X$ with the unique puzzle piece matching π. Then S_{n+1} has no more R-leaves at depth n (note that the puzzle pieces we collected always consist of at least two nodes).

  When we replace $\pi \in X$ by a piece P, we rename the nodes of P accordingly, to match our definition of strategies (if $\hat{\pi}$ is the root of P, we rename each node $\hat{\pi}\pi' \in P$ to $\pi\pi'$ when adding it to S_{n+1}).

- We define $\mathcal{S} = \bigcup_{n<\omega} S_n$, so $\mathcal{S}$ contains no more R-leaves.

Then $\mathcal{S}$ is a strategy: Each node $\pi \in \mathcal{S}$ corresponds to a node $\pi' \in \mathcal{S}_i|_i$ (either π' is an inner node of a puzzle piece, or $\pi' \in S_0$) and π has the same successors as π'. Moreover, the outcome is $[\![\mathcal{S}]\!]_\ell \geq [\![\mathcal{S}_i|_i]\!]_\ell$ as desired, because the repetition of puzzle pieces does not contain any problematic literals. Lastly, $\mathcal{S}$ is a winning strategy: Consider a play consistent with $\mathcal{S}$ and the corresponding maximal path through $\mathcal{S}$. If the path is finite, it ends in a leaf of $\mathcal{S}$ which corresponds to a leaf of $\mathcal{S}_i$ and is therefore winning. If the path visits infinitely many puzzle pieces (whose root nodes are R-nodes), then it visits infinitely many R-nodes and is thus winning by the parity condition (note that R belongs to the outermost fixed-point formula in φ). If the path is infinite and stays in S_0, then it corresponds to an infinite path of $\mathcal{S}_i|_i$ and is thus winning. Otherwise, the path is infinite, leaves S_0 at some point and visits only finitely many puzzle pieces. This means that it must from some point on stay in one piece, so it is winning by definition of puzzle pieces. We have therefore completed the puzzle (lemma). □

The Characterization

We are now ready to prove the characterization:

$$[\![\varphi(\mathbf{a})]\!]_\ell = \bigsqcup\{[\![\mathcal{S}]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\}$$

The interesting part is the proof for fixed-point formulae $\varphi = [\mathsf{fp}\, R\, \mathbf{x}.\, \vartheta](\mathbf{y})$. A strategy $\mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}$ may then look as in the picture below. We write $L/\dots$ to denote either a literal or an infinite path (without occurrences of R-nodes).

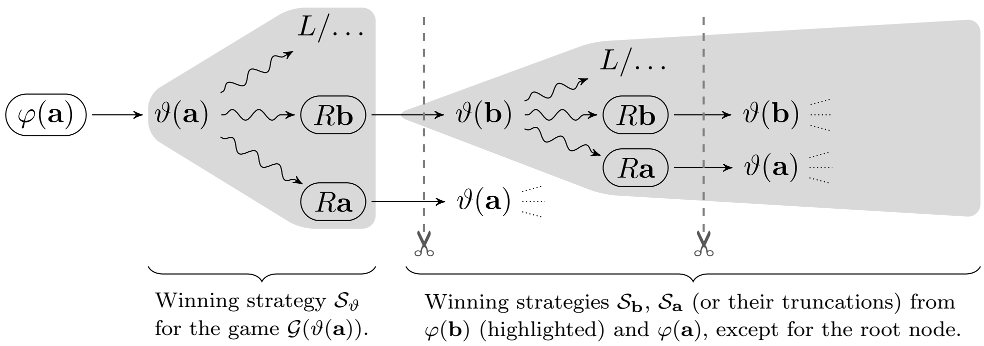

The strategy $\mathcal{S}$ must first move to $\vartheta(\mathbf{a})$ and thus contains a winning strategy from $\vartheta(\mathbf{a})$ in $\mathcal{G}(\varphi)$. If we only consider the strategy from $\vartheta(\mathbf{a})$ up to the first occurrence of an R-node, as indicated above, we obtain a winning strategy for the game $\mathcal{G}(\vartheta(\mathbf{a}))$. Note that in $\mathcal{G}(\vartheta(\mathbf{a}))$, R-nodes are terminals and are thus winning.

In $\mathcal{G}(\varphi(\mathbf{a}))$, these R-nodes are not terminals. Hence $\mathcal{S}$ must further contain substrategies for these R-nodes ($R\mathbf{b}$ and $R\mathbf{a}$ above). Because the positions $\varphi(\mathbf{b})$ and $R\mathbf{b}$ must both be followed by $\vartheta(\mathbf{b})$, we can view the substrategy from $R\mathbf{b}$ as a winning strategy $\mathcal{S}_\mathbf{b}$ for $\mathcal{G}(\varphi(\mathbf{b}))$ (as highlighted above).

We thus see that each winning strategy $\mathcal{S}$ in $\mathcal{G}(\varphi(\mathbf{a}))$ can be decomposed into a prefix $\mathcal{S}_\vartheta$ which we can identify with a winning strategy for $\mathcal{G}(\vartheta(\mathbf{a}))$ and, for all R-leaves of $\mathcal{S}_\vartheta$, substrategies which are winning strategies in $\mathcal{G}(\varphi)$. Conversely, every winning strategy $\mathcal{S}_\vartheta$ for $\mathcal{G}(\vartheta(\mathbf{a}))$ can be combined with winning strategies from $\varphi(\mathbf{a}), \varphi(\mathbf{b}), \dots$ for all the R-leaves of $\mathcal{S}_\vartheta$ to form a winning strategy in $\mathcal{G}(\varphi(\mathbf{a}))$.

If we build the strategy by starting with $\mathcal{S}_\vartheta$ but then appending the truncations $\mathcal{S}_\mathbf{a}|_n$ and $\mathcal{S}_\mathbf{b}|_n$ instead of $\mathcal{S}_\mathbf{a}, \mathcal{S}_\mathbf{b}$ (as indicated by the dashed lines in the picture), then the result is the $n+1$-truncation $\mathcal{S}|_{n+1}$ of a winning strategy $\mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}$, because $\mathcal{S}_\vartheta$ contains at most one R-node on each path. We exploit this observation in an inductive proof that relates the n-truncations of winning strategies with the n-th step of the fixed-point iteration.

Proof of theorem 7.18. Induction on the negation normal form of $\varphi(\mathbf{a})$:

- $\varphi(\mathbf{a}) = L \in \mathsf{Lit}$: Then $\mathcal{G}(\varphi(\mathbf{a}))$ consists only of a terminal position (which is winning). There is only one (trivial) strategy with $[\![\mathcal{S}]\!]_\ell = \ell(L) = [\![\varphi(\mathbf{a})]\!]_\ell$.

- $\varphi(\mathbf{a}) = \varphi_1(\mathbf{a}) \lor \varphi_2(\mathbf{a})$. The game $\mathcal{G}(\varphi(\mathbf{a}))$ is shown on the right. Each strategy $\mathcal{S}$ for $\mathcal{G}(\varphi(\mathbf{a}))$ makes a unique choice at $\varphi(\mathbf{a})$ and thus either consists of a strategy $\mathcal{S}_1$ for $\mathcal{G}(\varphi_1(\mathbf{a}))$ or a strategy $\mathcal{S}_2$ for $\mathcal{G}(\varphi_2(\mathbf{a}))$, but not both. Conversely, a strategy $\mathcal{S}_i$ for $\mathcal{G}(\varphi_i(\mathbf{a}))$ lifts to a strategy $\mathcal{S}$ from $\varphi(\mathbf{a})$ (for $i \in \{0,1\}$). We thus have:

$$\begin{tikzcd}[column sep=tiny] & \varphi(\mathbf{a}) \arrow[dl] \arrow[dr] & \\ \mathcal{G}(\varphi_1(\mathbf{a})) & & \mathcal{G}(\varphi_2(\mathbf{a})) \end{tikzcd}$$

$$\begin{aligned} \bigsqcup\{[\![\mathcal{S}]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\} &= \bigsqcup\{[\![\mathcal{S}_i]\!]_\ell \mid \mathcal{S}_i \in \mathcal{W}_{\varphi_i(\mathbf{a})},\ i \in \{0,1\}\} \\ &= \bigsqcup\{[\![\mathcal{S}_1]\!]_\ell \mid \mathcal{S}_1 \in \mathcal{W}_{\varphi_1(\mathbf{a})}\} \sqcup \bigsqcup\{[\![\mathcal{S}_2]\!]_\ell \mid \mathcal{S}_2 \in \mathcal{W}_{\varphi_2(\mathbf{a})}\} \\ &\overset{\text{IH}}{=} [\![\varphi_1(\mathbf{a})]\!]_\ell \sqcup [\![\varphi_2(\mathbf{a})]\!]_\ell = [\![\varphi_1(\mathbf{a})]\!]_\ell + [\![\varphi_2(\mathbf{a})]\!]_\ell = [\![\varphi(\mathbf{a})]\!]_\ell \end{aligned}$$

- $\varphi(\mathbf{a}) = \varphi_1(\mathbf{a}) \land \varphi_2(\mathbf{a})$. The reasoning is similar: Each strategy $\mathcal{S}$ for $\mathcal{G}(\varphi(\mathbf{a}))$ consists of both a strategy $\mathcal{S}_1$ for $\mathcal{G}(\varphi_1(\mathbf{a}))$ and a strategy $\mathcal{S}_2$ for $\mathcal{G}(\varphi_2(\mathbf{a}))$. The converse direction (all $\mathcal{S}_1$ and $\mathcal{S}_2$ together induce a strategy $\mathcal{S}$) holds as well.

$$\begin{tikzcd}[column sep=tiny] & \boxed{\varphi(\mathbf{a})} \arrow[dl] \arrow[dr] & \\ \mathcal{G}(\varphi_1(\mathbf{a})) & & \mathcal{G}(\varphi_2(\mathbf{a})) \end{tikzcd}$$

  If $\mathcal{S}$ consists of the two strategies $\mathcal{S}_1$ and $\mathcal{S}_2$, then we further have $[\![\mathcal{S}]\!]_\ell = [\![\mathcal{S}_1]\!]_\ell \cdot [\![\mathcal{S}_2]\!]_\ell$ by definition of the outcome and lemma 6.41 (1). The claim follows by induction and continuity of $\mathbb{S}^\infty[X]$ (for suprema of arbitrary sets, by corollary 4.24):

$$\begin{aligned} [\![\varphi_1(\mathbf{a})]\!]_\ell \cdot [\![\varphi_1(\mathbf{a})]\!]_\ell &\overset{\text{IH}}{=} \bigsqcup\{[\![\mathcal{S}_1]\!]_\ell \mid \mathcal{S}_1 \in \mathcal{W}_{\varphi_1(\mathbf{a})}\} \cdot \bigsqcup\{[\![\mathcal{S}_2]\!]_\ell \mid \mathcal{S}_2 \in \mathcal{W}_{\varphi_2(\mathbf{a})}\} \\ &= \bigsqcup\{[\![\mathcal{S}_1]\!]_\ell \cdot [\![\mathcal{S}_2]\!]_\ell \mid \mathcal{S}_i \in \mathcal{W}_{\varphi_i(\mathbf{a})} \text{ for } i \in \{1,2\}\} \\ &= \bigsqcup\{[\![\mathcal{S}]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\} \end{aligned}$$

The cases for $\exists$ and $\forall$ follow by the same arguments (except that we have $|A|$ instead of 2 child nodes). For fixed-point formulae, we use the decomposition into $\mathcal{S}_\vartheta$ and $\mathcal{S}_\mathbf{a}, \mathcal{S}_\mathbf{b}$ as motivated above. Consider the fixed-point iteration $(\pi_\beta)_{\beta \in \mathsf{On}}$ for $\varphi = [\mathsf{fp}\, R\,\mathbf{x}.\ \vartheta](\mathbf{y})$, where R is of arity r.

We first show by induction that the following holds for all $n < \omega$ and all $\mathbf{a} \in A^r$, where we set $\ell(✂) = 0$ for $\varphi = [\mathsf{lfp}\, R\mathbf{x}.\, \vartheta](\mathbf{y})$ and $\ell(✂) = 1$ for $\varphi = [\mathsf{gfp}\, R\mathbf{x}.\, \vartheta](\mathbf{y})$:

$$\pi_n(\mathbf{a}) = \bigsqcup\{[\![\mathcal{S}|_n]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\}$$

- For $n = 0$, we trivially have $\pi_0 = 0$ and $[\![\mathcal{S}|_0]\!]_\ell = 0$ for least fixed points and $\pi_0 = 1$ and $[\![\mathcal{S}|_0]\!]_\ell = 1$ for greatest fixed points.
- For the induction step $n \to n+1$, we first show $\pi_{n+1}(\mathbf{a}) \leq \bigsqcup\{[\![\mathcal{S}|_{n+1}]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\}$. To simplify notation, we set $\mathsf{Lit}_A^* = \mathsf{Lit}_A \setminus \{R\mathbf{a} \mid \mathbf{a} \in A^r\}$. By the induction hypothesis on ϑ and on π_n, we can write π_{n+1} as follows:

$$\begin{aligned}\pi_{n+1}(\mathbf{a}) = [\![\vartheta(\mathbf{a})]\!]_{\ell[R/\pi_n]} &= \bigsqcup\left\{[\![\mathcal{S}_\vartheta]\!]_{\ell[R/\pi_n]} \;\middle|\; \mathcal{S}_\vartheta \in \mathcal{W}_{\vartheta(\mathbf{a})}\right\} \\ &= \bigsqcup\left\{\prod_{L \in \mathsf{Lit}_A^*} \ell(L)^{n_L} \cdot \prod_{\mathbf{b} \in A^r} \underbrace{\left(\bigsqcup\{[\![\mathcal{S}_\mathbf{b}|_n]\!]_\ell \mid \mathcal{S}_\mathbf{b} \in \mathcal{W}_{\varphi(\mathbf{b})}\}\right)^{\mathbf{b}_L}}_{(*)} \;\middle|\; \mathcal{S}_\vartheta \in \mathcal{W}_{\vartheta(\mathbf{a})}\right\}\end{aligned}$$

  where we set $n_L = |\mathcal{S}_\vartheta|_L$ and $\mathbf{b}_L = |\mathcal{S}_\vartheta|_{R\mathbf{b}}$, depending on $\mathcal{S}_\vartheta$. Let us first fix a strategy $\mathcal{S}_\vartheta$ and $\mathbf{b} \in A^r$ and consider the term $(*)$. Recall that absorptive polynomials are finite and that for a set S, $\bigsqcup S = \mathsf{maximals}(\bigcup S)$. We can thus write

$$(*) = \left([\![\mathcal{S}_\mathbf{b}^1|_n]\!]_\ell + \cdots + [\![\mathcal{S}_\mathbf{b}^k|_n]\!]_\ell\right)^{\mathbf{b}_L}$$

  for some $\mathcal{S}_\mathbf{b}^1, \dots, \mathcal{S}_\mathbf{b}^k \in \mathcal{W}_{\varphi(\mathbf{b})}$. If $\mathbf{b}_L = \infty$, then by lemma 6.41:

$$(*) = [\![\mathcal{S}_\mathbf{b}^1|_n]\!]_\ell^{\,\infty} + \cdots + [\![\mathcal{S}_\mathbf{b}^k|_n]\!]_\ell^{\,\infty}$$

  Otherwise, $\mathbf{b}_L = l < \infty$. Then each monomial of $(*)$ is of the form

$$[\![\mathcal{S}_\mathbf{b}^{i_1}|_n]\!]_\ell \cdots [\![\mathcal{S}_\mathbf{b}^{i_l}|_n]\!]_\ell, \quad \text{where } i_1, \dots, i_l \in \{1, \dots, k\}$$

  To prove that $\pi_{n+1}(\mathbf{a}) \leq \bigsqcup\{[\![\mathcal{S}|_{n+1}]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\}$, we show that each monomial of $\pi_{n+1}(\mathbf{a})$ is absorbed by the right-hand side. Given the above considerations, we know that these monomials are of the form

$$m = \prod_{L \in \mathsf{Lit}_A^*} \ell(L)^{n'_L}, \quad n'_L = n_L + \sum_{\substack{\mathbf{b} \in A^r \\ \mathbf{b}_L = \infty}} \big|\mathcal{S}_\mathbf{b}|_n\big|_L \cdot \infty + \sum_{\substack{\mathbf{b} \in A^r \\ \mathbf{b}_L = l < \infty}} \big|\mathcal{S}_\mathbf{b}^1|_n\big|_L \cdots \big|\mathcal{S}_\mathbf{b}^l|_n\big|_L$$

  for some $\mathcal{S}_\vartheta \in \mathcal{W}_{\vartheta(\mathbf{a})}$ (which defines n_L) and some $\mathcal{S}_\mathbf{b}$, $\mathcal{S}_\mathbf{b}^i \in \mathcal{W}_{\varphi(\mathbf{b})}$ (for all $\mathbf{b}$ and i).

  Now consider the strategy which starts with $\mathcal{S}_\vartheta$ and for each $\mathbf{b} \in A^r$, we replace all $R\mathbf{b}$-leaves of $\mathcal{S}_\vartheta$ by either $\mathcal{S}_\mathbf{b}$ (if there are infinitely many such leaves) or by the strategies $\mathcal{S}_\mathbf{b}^1, \dots, \mathcal{S}_\mathbf{b}^l$ if there are $l < \infty$ such leaves (it does not matter in which order these strategies are assigned to the leaves). The result is a strategy $\mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}$.

We further see that $\mathcal{S}|_{n+1}$ results from $\mathcal{S}_\vartheta$ in the same way if we replace leaves by the truncations $\mathcal{S}_\mathbf{b}|_n$ instead of $\mathcal{S}_\mathbf{b}$ (and $\mathcal{S}^i_\mathbf{b}|_n$ instead of $\mathcal{S}^i_\mathbf{b}$), because $\mathcal{S}_\vartheta$ contains R-nodes only as leaves. By this construction, we see that $|\mathcal{S}|_L = n'_L$ for each $L \in \mathsf{Lit}^*_A$ and hence $m = [\![\mathcal{S}|_{n+1}]\!]_\ell$ and thus $m \leq \bigsqcup\{[\![\mathcal{S}|_{n+1}]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\}$ as claimed.

- For the other direction, we fix a strategy $\mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}$ and show that $[\![\mathcal{S}|_{n+1}]\!]_\ell \leq \pi_{n+1}(\mathbf{a})$. In the case that $\ell(\text{✂}) = 0$ and ✂ appears in $\mathcal{S}|_{n+1}$, we have $[\![\mathcal{S}|_{n+1}]\!]_\ell = 0$ and there is nothing to show. If $\ell(\text{✂}) = 1$, the appearance of ✂ does not affect the outcome $[\![\mathcal{S}|_{n+1}]\!]_\ell$. We again decompose $\mathcal{S}$ into a prefix $\mathcal{S}_\vartheta$ corresponding to a winning strategy in $\mathcal{G}(\vartheta(\mathbf{a}))$ and substrategies from all R-leaves of $\mathcal{S}_\vartheta$.

  We first consider any $\mathbf{b} \in A^r$ for which $|\mathcal{S}_\vartheta|_{R\mathbf{b}} = \infty$ such that we have infinitely many such substrategies from $R\mathbf{b}$-leaves. We make the following claim: There is a strategy $\mathcal{S}'_\mathbf{b} \in \mathcal{W}_{\varphi(\mathbf{b})}$ such that the strategy $\mathcal{S}'$ which is like $\mathcal{S}$ but uses $\mathcal{S}'_\mathbf{b}$ for all of the infinitely many $R\mathbf{b}$-leaves of $\mathcal{S}_\vartheta$ satisfies $[\![\mathcal{S}|_{n+1}]\!]_\ell \leq [\![\mathcal{S}'|_{n+1}]\!]_\ell$.

  Proof: Let $(\mathcal{S}^i_\mathbf{b})_{i<\omega}$ be the family of all substrategies $\mathcal{S}$ uses from $R\mathbf{b}$-leaves of $\mathcal{S}_\vartheta$. As in the proof of the puzzle lemma, we call a literal L problematic if $\big|\mathcal{S}|_{n+1}\big|_L < \infty$. Let L be a problematic literal. Then there can only be finitely many i with $\big|\mathcal{S}^i_\mathbf{b}|_n\big|_L > 0$ (otherwise L would occur infinitely often in $\mathcal{S}|_{n+1}$). As ω is infinite while the number of literals is finite, there is an $i < \omega$ such that $\mathcal{S}^i_\mathbf{b}|_n$ contains no problematic literals at all. We then set $\mathcal{S}'_\mathbf{b} = \mathcal{S}^i_\mathbf{b}$ and the claim follows.

  Due to this claim and because A^r is finite, we obtain a strategy $\mathcal{S}'$ with $[\![\mathcal{S}|_{n+1}]\!]_\ell \leq [\![\mathcal{S}'|_{n+1}]\!]_\ell$ such that for all $\mathbf{b} \in A^r$ with $|\mathcal{S}_\vartheta|_{R\mathbf{b}} = \infty$, $\mathcal{S}'$ uses the same strategy from all $R\mathbf{b}$-leaves of $\mathcal{S}_\vartheta$. From the other direction, we know that

$$\pi_{n+1}(\mathbf{a}) \geq \prod_{L \in \mathsf{Lit}^*_A} \ell(L)^{n_L} \cdot \prod_{\mathbf{b} \in A^r} \Big(\underbrace{\bigsqcup\{[\![\mathcal{S}_\mathbf{b}|_n]\!]_\ell \mid \mathcal{S}_\mathbf{b} \in \mathcal{W}_{\varphi(\mathbf{b})}\}}_{=: P_\mathbf{b}}\Big)^{\mathbf{b}_L}$$

  where $n_L = |\mathcal{S}_\vartheta|_L$ and $\mathbf{b}_L = |\mathcal{S}_\vartheta|_{R\mathbf{b}}$. Consider the strategies $\mathcal{S}'$ uses from the R-leaves of $\mathcal{S}_\vartheta$: For $\mathbf{b} \in A^r$ with $|\mathcal{S}_\vartheta|_{R\mathbf{b}} = \infty$, let $\mathcal{S}_\mathbf{b}$ be the strategy that $\mathcal{S}'$ uses from all $R\mathbf{b}$-leaves. For $\mathbf{b}$ with $|\mathcal{S}_\vartheta|_{R\mathbf{b}} = l < \infty$, let $\mathcal{S}^1_\mathbf{b}, \dots, \mathcal{S}^l_\mathbf{b}$ be the strategies $\mathcal{S}'$ uses from the $R\mathbf{b}$-leaves of $\mathcal{S}_\vartheta$. Let further

$$n'_L = n_L + \sum_{\substack{\mathbf{b} \in A^r \\ \mathbf{b}_L = \infty}} \big|\mathcal{S}_\mathbf{b}|_n\big|_L \cdot \infty + \sum_{\substack{\mathbf{b} \in A^r \\ \mathbf{b}_L = l < \infty}} \big|\mathcal{S}^1_\mathbf{b}|_n\big|_L \cdots \big|\mathcal{S}^l_\mathbf{b}|_n\big|_L$$

  We can apply commutativity and lemma 6.41 to conclude

$$\begin{aligned}[\![\mathcal{S}'|_{n+1}]\!]_\ell = \prod_{L \in \mathsf{Lit}^*_A} \ell(L)^{n'_L} &= \prod_{L \in \mathsf{Lit}^*_A} \ell(L)^{n_L} \cdot \prod_{\substack{\mathbf{b} \in A^r \\ \mathbf{b}_L = \infty}} [\![\mathcal{S}_\mathbf{b}|_n]\!]_\ell{}^\infty \cdot \prod_{\substack{\mathbf{b} \in A^r \\ \mathbf{b}_L = l < \infty}} [\![\mathcal{S}^1_\mathbf{b}|_n]\!]_\ell \cdots [\![\mathcal{S}^l_\mathbf{b}|_n]\!]_\ell \\ &\leq \prod_{L \in \mathsf{Lit}^*_A} \ell(L)^{n_L} \cdot \prod_{\substack{\mathbf{b} \in A^r \\ \mathbf{b}_L = \infty}} P_\mathbf{b}^\infty \cdot \prod_{\substack{\mathbf{b} \in A^r \\ \mathbf{b}_L = l < \infty}} P_\mathbf{b}^l \leq \pi_{n+1}(\mathbf{a})\end{aligned}$$

This proves the inductive claim. We now show that for least and greatest fixed points,

$$\pi_\omega(\mathbf{a}) = \bigsqcup\{[\![\mathcal{S}]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\}$$

- For $\varphi = [\mathsf{lfp}\, R\, \mathbf{x}.\, \vartheta](\mathbf{y})$, this follows via lemma 7.21 by swapping suprema:

$$\begin{aligned}\pi_\omega(\mathbf{a}) = \bigsqcup_{n<\omega} \pi_n(\mathbf{a}) &= \bigsqcup_{n<\omega} \bigsqcup\{[\![\mathcal{S}|_n]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\} \\ &= \bigsqcup\Big\{\bigsqcup_{n<\omega} [\![\mathcal{S}|_n]\!]_\ell \;\Big|\; \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\Big\} = \bigsqcup\{[\![\mathcal{S}]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\}\end{aligned}$$

- For $\varphi = [\mathsf{gfp}\, R\, \mathbf{x}.\, \vartheta](\mathbf{y})$, the proof is more difficult and requires the puzzle lemma 7.22. We first note that one direction is trivial (using lemma 7.21 in the last step):

$$\bigsqcap_{n<\omega} \bigsqcup\{[\![\mathcal{S}|_n]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\} \geq \bigsqcup\Big\{\bigsqcap_{n<\omega} [\![\mathcal{S}|_n]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\Big\} = \bigsqcup\{[\![\mathcal{S}]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\}$$

  For the other direction, we use the characterization of infima from theorem 6.27:

$$\pi_\omega(\mathbf{a}) = \bigsqcap_{n<\omega} \underbrace{\bigsqcup\{[\![\mathcal{S}|_n]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\}}_{=:P_n} = \bigsqcup\Big\{\bigsqcap_{n<\omega} \mathfrak{m}_n \;\Big|\; \mathfrak{m} \in \mathfrak{M}\Big\}$$

  where $\mathfrak{M}$ is the set of monomial chains through the polynomials P_n. That is, $\mathfrak{m}_n = [\![\mathcal{S}_n|_n]\!]_\ell$ for some strategy $\mathcal{S}_n \in \mathcal{W}_{\varphi(\mathbf{a})}$ (for each $n < \omega$). Consider one monomial chain $\mathfrak{m} \in \mathfrak{M}$. By the puzzle lemma, there is a strategy $\mathcal{S}_\mathfrak{m} \in \mathcal{W}_{\varphi(\mathbf{a})}$ such that $[\![\mathcal{S}_\mathfrak{m}]\!]_\ell \geq \bigsqcap_n \mathfrak{m}_n$. This implies the nontrivial direction:

$$\pi_\omega(\mathbf{a}) = \bigsqcup\Big\{\bigsqcap_{n<\omega} \mathfrak{m}_n \;\Big|\; \mathfrak{m} \in \mathfrak{M}\Big\} \leq \bigsqcup\{[\![\mathcal{S}_\mathfrak{m}]\!]_\ell \mid \mathfrak{m} \in \mathfrak{M}\} \leq \bigsqcup\{[\![\mathcal{S}]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\}$$

By corollary 6.44, we have $[\![\varphi(\mathbf{a})]\!]_\ell = \pi_\omega(\mathbf{a})$, finally closing the proof. □

Consequences

This result generalizes to all absorptive continuous semirings due to the universal property of $\mathbb{S}^\infty[X]$. We first observe that outcomes of strategies are preserved (using lemma 6.41):

Lemma 7.25. *Let S and T be absorptive continuous semirings, let $h : S \to T$ be a cpo-semiring homomorphism. Let further ℓ be an S-interpretation, φ a formula and $\mathcal{S}$ a strategy in $\mathcal{G}(\varphi)$. Then*

$$h([\![\mathcal{S}]\!]_\ell) = [\![\mathcal{S}]\!]_{h\circ\ell}$$

Corollary 7.26. *Let T be an absorptive continuous semiring, let ℓ be a T-interpretation and $\varphi(\mathbf{x})$ a formula. Then*

$$[\![\varphi(\mathbf{a})]\!]_\ell = \bigsqcup\{[\![\mathcal{S}]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\}$$

Proof. Consider the most general $\mathbb{S}^\infty[X]$-interpretation ℓ^* with $\ell^*(L) = x_L$ and $X = \{x_L \mid L \in \mathsf{Lit}_A\}$. By the universality theorem, there is a cpo-semiring homomorphism $h : \mathbb{S}^\infty[X] \to T$ with $\ell = h \circ \ell^*$. Then,

$$\begin{aligned}[\![\varphi(\mathbf{a})]\!]_\ell = h([\![\varphi(\mathbf{a})]\!]_{\ell^*}) &= h\big(\bigsqcup\{[\![\mathcal{S}]\!]_{\ell^*} \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\}\big) \\ &\overset{(*)}{=} \bigsqcup\{h([\![\mathcal{S}]\!]_{\ell^*}) \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\} = \bigsqcup\{[\![\mathcal{S}]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\}\end{aligned}$$

where $(*)$ holds by proposition 6.33. $\square$

When working with *model-defining* interpretations, we can consider strategies in the *classical* model checking game $\mathcal{G}(\mathfrak{A}_\ell, \varphi)$ instead of the generic game (see proposition 7.17). Our considerations about absorption on strategies further show that we can restrict the supremum to only consider *weakly positional* strategies (see proposition 7.19).

Corollary 7.27. *Let T be an absorptive continuous semiring, let ℓ be a model-defining T-interpretation and $\varphi(\mathbf{x})$ a formula. Then*

$$[\![\varphi(\mathbf{a})]\!]_\ell = \bigsqcup\{[\![\mathcal{S}]\!]_\ell \mid \mathcal{S} \textit{ is a weakly positional winning strategy in } \mathcal{G}\big(\mathfrak{A}_\ell, \varphi(\mathbf{a})\big)\}$$

Remark: The finiteness of absorptive polynomials implies that for every formula $\varphi(\mathbf{a})$, there is only a finite number of absorption-maximal strategies. This also holds for the absorptive continuous semiring T, as homomorphisms respect absorption.

The last corollary we mention considers the semiring $\mathsf{PosBool}[X]$ in which we can further restrict the strategies to be positional instead of weakly positional. The reason is that $\mathsf{PosBool}[X]$ is multiplicatively idempotent. For the outcome, it is thus only relevant which literals occur in $\mathcal{S}$, but not how often. The positional determinacy of parity games then yields an analogue of proposition 7.19 that, given a strategy $\mathcal{S}$, guarantees the existence of a positional strategy $\mathcal{S}' \geq \mathcal{S}$. As in the excursion in section 7.2, we can thus use $\mathsf{PosBool}[X]$ if we are interested in information about positional strategies.

Corollary 7.28. *let ℓ be a $\mathsf{PosBool}[X]$-interpretation (or a $\mathsf{PosBool}[X, \overline{X}]$-interpretation) and $\varphi(\mathbf{x})$ a formula. Then*

$$[\![\varphi(\mathbf{a})]\!]_\ell = \bigsqcup\{[\![\mathcal{S}]\!]_\ell \mid \mathcal{S} \textit{ is a positional winning strategy in } \mathcal{G}\big(\varphi(\mathbf{a})\big)\}$$

To close this section, let us recall the reverse analysis example from the last chapter where we used dual polynomials to reason about several models at once. Recall that dual polynomials $\mathbb{S}^\infty[X, \overline{X}]$ form an absorptive continuous semiring and are thus covered by the above corollaries.

Example 7.29. Recall the setting of example 6.64:

$$\varphi(u) = \big[\mathsf{gfp}\, R\, x.\ Px \wedge \forall y\big(\underbrace{\neg Exy \vee Ry}_{Exy \to Ry}\big)\big](u)$$

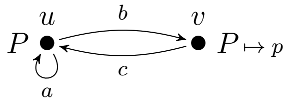

We have seen the results $[\![\varphi(u)]\!]_\ell = \bar{b} + p$ in $\mathsf{PosBool}[X, \overline{X}]$ and (perhaps less obvious) $[\![\varphi(u)]\!]_\ell = \bar{a}\bar{b} + \bar{a}\bar{c}p + \bar{b}^\infty + p^\infty$ in $\mathbb{S}^\infty[X, \overline{X}]$. The strategy characterization provides a more intuitive approach to reason about the evaluation of $\varphi(u)$ by considering the following model checking game. We shorten the node labels for space reasons replace terminals by their semiring value (see figure 8 for a more detailed example).

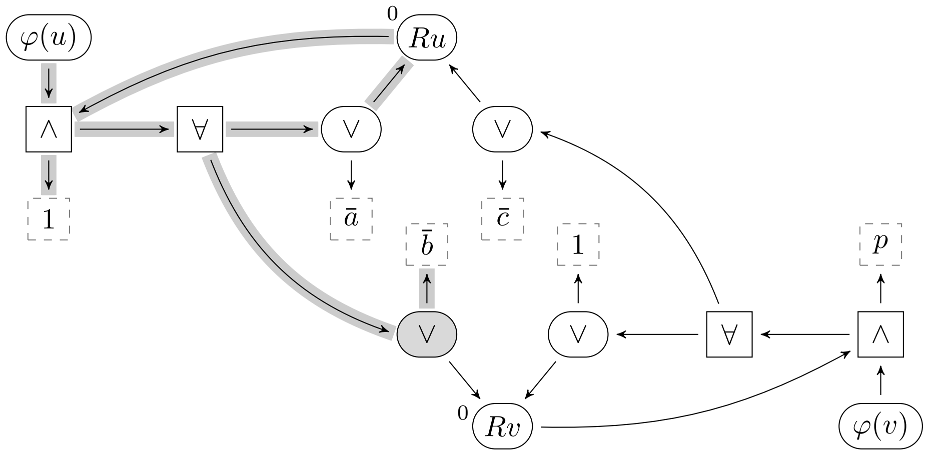

The only choices for Verifier occur in the nodes labeled $\vee$. We can thus compute the outcome of all positional strategies by considering all combinations of these choices. As we are only interested in absorption-maximal strategies and 1 is the greatest element, we can w.l.o.g. assume that Verifier always moves to 1 from the lower right $\vee$-node. The remaining 8 combinations then lead to strategies with the following outcomes:

$$\underbrace{\overbrace{\bar{a}\bar{b},\quad \bar{a}\bar{b},}^{\text{moving to } \bar{b}}\quad \bar{a}\bar{c}p,\quad \bar{a}^\infty p^\infty}_{\text{moving to } \bar{a}},\quad \overbrace{\bar{b}^\infty,\quad \bar{b}^\infty,}^{\text{moving to } \bar{b}}\quad \bar{c}^\infty p^\infty,\quad p^\infty$$

The strategy with outcome $\bar{b}^\infty$ is highlighted in the picture above. In $\mathsf{PosBool}[X, \overline{X}]$, these outcomes simplify to just $\bar{b} + p$ (by dropping exponents and due to absorption). For

$\mathbb{S}^\infty[X, \bar{X}]$, we must additionally consider weakly positional strategies. As an example, consider the highlighted node and the strategy for $\bar{b}^\infty$. As this node is reached via the $\forall$-node belonging to Falsifier, Verifier is allowed to make different decisions whenever the highlighted node is visited. However, it is easy to see that for the game above, doing so would only lead to smaller outcomes (such as $\bar{b}^\infty\bar{c}^5$ instead of $\bar{b}^\infty$). We thus obtain the following result from the strategy characterization:

$$[\![\varphi(u)]\!]_\ell = \bar{a}\bar{b} + \bar{a}\bar{c}p + \bar{b}^\infty + p^\infty$$

Although the semiring $\mathsf{PosBool}[X, \bar{X}]$ leads to a more intuitive result for this example (and is sufficient when we just want to know in which models $\varphi(u)$ holds), we see that the information computed in $\mathbb{S}^\infty[X]$ can indeed be explained by viewing winning strategies as witnesses for the truth of $\varphi(u)$.

7.5 Thoughts Beyond Absorption

The proof of the strategy characterization theorem relied on the structure of absorptive polynomials $\mathbb{S}^\infty[X]$ for some arguments. Nevertheless, we conjecture that the connection to model checking games can be extended also to non-absorptive semirings such as $\mathbb{W}[X]$ and $\mathbb{N}^\infty$. Without absorption, we have to adapt the definition of the outcome as follows:

$$[\![\mathcal{S}]\!]_\ell = \begin{cases} \prod_{L \in \mathsf{Lit}_A} \ell(L)^{|\mathcal{S}|_L}, & \text{if } \mathcal{S} \text{ is finite} \\ \top \cdot \prod_{L \in \mathsf{Lit}_A} \ell(L)^{|\mathcal{S}|_L}, & \text{if } \mathcal{S} \text{ is infinite} \end{cases}$$

First note that this coincides with our earlier definition in $\mathbb{S}^\infty[X]$ (and in all absorptive continuous semirings), as there $\top = 1$. For the rationale behind this generalization, recall that we defined the outcome $[\![\mathcal{S}]\!]_\ell$ as product of all literals occurring in $\mathcal{S}$. Instead, we can equivalently define it as the product over all plays, where we associate each finite play with the terminal it ends in (to be precise, with the value ℓ assigns to this terminal). By considering only the literals in $\mathcal{S}$, we ignore infinite plays, so we may associate them with the value 1. Here we instead associate infinite plays with the value $\top$.

Infinite winning strategies can only arise from greatest fixed points, so this definition essentially distinguishes between least and greatest fixed points and introduces an asymmetry if $\top \neq 1$. This is in line with our observations for the semiring $\mathbb{N}^\infty$, in which this asymmetry often leads to the value $\top = \infty$ for gfp-formulae.

The definition of $[\![\mathcal{S}]\!]_\ell$ requires the existence of an infinitary power a^∞. Note that this is only required when $\mathcal{S}$ is infinite. We make the following observations.

Lemma 7.30. *Let S be a cpo semiring. Then*

(1) $\top \cdot \top = \top$,

(2) *for all* $a, b \in S$: $(a + ab) \cdot \top = a \cdot \top$ *(absorption relative to* $\top$*)*,

(3) *for all* $a \in S$: $(a^n \cdot \top)_{n<\omega}$ *is a descending chain.*

Proof. We trivially have $\top^2 \leq \top$. For the other direction, note that $1 \leq \top$ and hence monotonicity of multiplication implies $\top = 1 \cdot \top \leq \top \cdot \top$. For (2), we use a similar argument: We trivially have $(a + ab)\top = a\top + ab\top \geq a\top$ and for the other direction, $(1 + b)\top \leq \top$ implies $(a + ab)\top \leq a\top$ by monotonicity of multiplication. For (3), we trivially have $a\top \leq \top$ and thus $a^2\top \leq a\top$. The claim follows by induction on n. $\square$

Similar to the generalization of $[\![\mathcal{S}]\!]_\ell$, we can view the chain $a^n\top$ as a generalization of the chain a^n we have in absorptive semirings. Following this analogy, we define:

$$a^\infty := \bigsqcap_{n<\omega} a^n\top$$

With these considerations, we formulate the following conjecture which states that the strategy characterization applies to all continuous semirings.

Conjecture. *Let $\varphi(\mathbf{x})$ be a formula, S a continuous semiring and ℓ an S-interpretation. Then,*

$$[\![\varphi(\mathbf{a})]\!]_\ell = \sum\{[\![\mathcal{S}]\!]_\ell \mid \mathcal{S} \in \mathcal{W}_{\varphi(\mathbf{a})}\}$$

where we define the summation of a countable subset $A = \{a_i \mid i < \omega\} \subseteq S$ by the supremum over all partial summations:

$$\sum A = \bigsqcup_{n<\omega} a_1 + \cdots + a_n$$

Note that, although the number of strategies can be uncountable, there are only countably many different outcomes due to the finite number of literals, so the sum is well-defined.

The motivation behind this conjecture stems on the one hand from the well-known equivalence between model checking games and logic which should also apply to semirings semantics. On the other hand, it is motivated by an analysis of our proof of the characterization theorem. We have used induction to relate the fixed-point iteration with truncations of strategies. For these truncations, we now set $\ell(\text{✂}) = \top$ instead of $\ell(\text{✂}) = 1$, following our previous adaptions. It is then easy to see that lemma 7.21 still applies (in particular, the sequence of truncations $[\![\mathcal{S}|_n]\!]_\ell$ still forms a chain). An important observation was the decomposition of winning strategies $\mathcal{S}$ into a prefix $\mathcal{S}_\vartheta$

and several strategies for the R-leaves of $\mathcal{S}_\vartheta$, which is independent of the semiring and still applies. The second key step was the puzzle lemma. Here, the main idea of repeating a literal-free layer ad infinitum is independent of the semiring as well. Our reasoning on problematic literals was based on absorption, but we can instead use absorption relative to $\top$ as shown above.

The step that relied the most on $\mathbb{S}^\infty[X]$ is the actual computation in the induction step of the final proof, where we reasoned via monomials. For the limit ordinal step π_ω, we have further used the characterization of infima via monomial chains. In specific semirings, we can find alternative arguments: For $\mathbb{W}[X]$, we can similarly reason via monomials and the proof for π_ω can exploit the finiteness of this semiring. For $\mathbb{N}^\infty$, we can use the well-foundedness for π_ω and can otherwise rely on the simple structure and case distinctions between finite and infinite number of outcomes. In both semirings, the computation in the induction step can be simplified by observing that $a^2\top = a\top$ for all elements a, which allows us to simplify the outcome of infinite strategies $\mathcal{S}$ by omitting the exponents $|\mathcal{S}|_L$. Note that the refined strategy characterization is in line with our examples in chapter 6: In $\mathbb{N}^\infty$, we have $\top = \infty$, so whenever there is a winning strategy which permits an infinite play, the formula evaluates to ∞. This explains our examples with infinite paths. In $\mathbb{W}[X]$, we obtain more information, but due to the multiplication with $\top$ we also obtain monomials not corresponding to proofs. We have seen this in example 6.13 where we got the monomial abc as part of the polynomial $ab\top$.

The arguments for these semirings do not work in general (for example, $x\top \neq x^2\top$ in $\mathbb{B}[\![X]\!]$). Instead, a convenient way to establish the conjecture would be to work with a universal semiring more general than $\mathbb{S}^\infty[X]$ that captures all continuous semirings. It is not clear whether such a semiring exists, but we want to suggest an interesting candidate. Consider the extension of formal power series $\mathbb{N}^\infty[\![X]\!]$ and $\mathbb{B}[\![X]\!]$ by permitting the exponent ∞, in the same way that $\mathbb{S}^\infty[X]$ extends $\mathbb{S}[X]$. That is, we allow monomials such as x^∞ and x^3y^∞ to occur in power series. In contrast to $\mathbb{S}^\infty[X]$, this does not solve the problem of chain-positivity, as the same counterexample is still possible. However, fixed-point iterations induced by gfp-formulae always start at $\top$ which now includes x^∞, so instead of $\bigsqcap_n x^n\top = 0$ (our counterexample for chain-positivity in $\mathbb{B}[\![X]\!]$) we now have $\bigsqcap_n x^n\top = x^\infty\top$. So while the extension of formal power series is not chain-positive in general, the intuition is that chain-positivity might still hold for chains resulting from formulae. Similarly, homomorphisms induced by variable assignments are not continuous, but could still suffice to preserve fixed-point iterations of formulae.

An important detail is that the conjecture is related to our earlier stated open question whether fixed-point iterations terminate at ω in continuous semirings (see section 5.5). If our inductive proof using strategy truncations can be used to establish the conjecture, then this implies a positive answer to the open question. The reason is that at step ω, the supremum (or infimum) over the outcomes of all n-truncations of a strategy yields the outcome of the complete strategy (due to lemma 7.21). This means that π_ω can be described by the outcomes of all winning strategies and is thus already the overall result of the fixed-point iteration.

Even if the ideas sketched above turn out to be feasible and formal power series can thus capture computations in all continuous semirings, our statement that $\mathbb{S}^{\infty}[X]$ is the sensible choice for provenance analysis of LFP still stands. The reason is that, as we have seen for $\mathbb{W}[X]$, the refined outcome with $\top \neq 1$ can lead to monomials not corresponding to strategies. We therefore do not view the proposed extension of power series as an alternative to $\mathbb{S}^{\infty}[X]$, but rather as a potential way to reason about semiring semantics (via model checking games) beyond absorption.

8 Conclusion

Summary

In this thesis, we have defined and analyzed semiring semantics for the fixed-point logic LFP as a means to perform provenance analysis. While semiring provenance for positive LFP has already been considered in [GT19], we have also studied greatest fixed points which turn out to be challenging in several aspects. Greatest fixed points have already been studied for the logic CTL in [Mrk18] and we have expanded on these results both by considering a more expressive logic and by permitting a more general notion of semirings.

The first main result is that semiring semantics for LFP are well-defined for all *cpo semirings*, in particular for continuous semirings which include a wide range of natural examples and applications. These notions are inspired by an order-theoretic perspective which allows a uniform treatment of both least and greatest fixed points. We have argued that cpo semirings are possibly the most general class of semirings suitable for LFP, with an alternative being fully ω-continuous semirings (for alternation-free LFP or in the context of absorption). To provide a sound foundation, we have discussed the relation between the concepts of cpo, continuous and lattice semirings as well as algebraic properties such as idempotence and absorption with many examples, including counterexamples that answer a question on absorptive lattice semirings from [Mrk18].

Our analysis of the resulting semantics was motivated by questions from a provenance point of view, most notably the compatibility with standard semantics and the interplay with negation. To overcome challenges arising from greatest fixed points, we have introduced the concept of *chain-positivity* to guarantee that semiring semantics preserve truth. This justifies our view of these semantics as a *generalization of standard semantics by multiple truth values* and includes the concept of duality of least and greatest fixed points. This view is also evident in the truth projection homomorphism $\dagger_S$ which connects semiring and standard semantics. In general, homomorphisms have played a central role in many arguments and we have seen that *cpo-semiring homomorphisms* are required to preserve semiring semantics for LFP. Another difficulty of greatest fixed points becomes apparent for polynomial semirings which provide the most general (and thus most interesting) provenance information. While formal power series $\mathbb{N}^\infty[\![X]\!]$ can be used for positive LFP, the interpretation of greatest fixed points is more involved and we have argued by symmetry that we need *absorptive* semirings to obtain reasonable information. In particular, we have analyzed (generalized) absorptive polynomials $\mathbb{S}^\infty[X]$ in detail and have shown them to be the *most general absorptive continuous semiring* (in terms of cpo-semiring homomorphisms), thus taking the place of $\mathbb{N}^\infty[\![X]\!]$ as the semiring of choice for provenance analysis of full LFP.

To formalize the meaning of *reasonable provenance information*, we have shown that provenance computations in absorptive continuous semirings can be characterized by *winning strategies* in model checking games. As a consequence, these semirings provide information about absorption-maximal winning strategies and we have seen that these are related to (weakly) positional strategies or, in other words, to *shortest proofs* of formulae. This leads to the overall conclusion that, under the assumptions of chain-positivity and absorption, the idea of semiring provenance can successfully be applied to full LFP.

Future Work

An important follow-up question is how semiring semantics can actually be computed, given a formula φ and an interpretation ℓ. We can certainly compute information in the finite semiring $\mathsf{PosBool}[X]$, but the most important semiring $\mathbb{S}^\infty[X]$ is infinite and thus requires further work. Instead of the fixed-point iteration, model checking games may provide a more accessible approach to compute $\mathbb{S}^\infty[X]$-provenance.

Motivated by observations for positive and alternation-free LFP, we have further posed the open question (in section 5.5) whether the fixed-point iteration always stops at step ω in continuous semirings (which would imply that we can work with fully ω-continuous semirings). As the continuity of update operators F_ℓ^ϑ is not guaranteed for alternating formulae, we have instead used the universality of $\mathbb{S}^\infty[X]$ to answer this question for absorptive semirings. The general case remains open and is closely related to our conjecture that the strategy characterization applies to all continuous semirings (in section 7.5). We have suggested an extension of formal power series by the exponent ∞ as one possible way to resolve both issues.

In chapter 7, we have used the outcomes of winning strategies in parity games in order to understand semiring semantics. A related question is how one can define semiring provenance for parity games in the first place, similar to the provenance analysis of reachability games proposed in [GT19]. Our excursion on the computation of positional winning strategies illustrates that one can perform provenance analysis of parity games through the evaluation of logical formulae, but there might also be a more direct approach.

Leaving the context of provenance analysis, semiring semantics are also interesting in their own right. For example, it is known that on finite structures, every LFP formula is equivalent to a formula in positive LFP (e.g., [GKL+07]). This does not hold for $\mathbb{S}^\infty[X]$-semantics, where lfp-iterations are always finite. Starting from an interpretation that maps literals to variables, lfp-formulae can thus never evaluate to monomials of the form x^∞. This illustrates that it might be worthwhile to study the properties of semiring semantics for certain (classes of) semirings in future work.

References

[BKWC01] Peter Buneman, Sanjeev Khanna, and Tan Wang-Chiew. Why and where: A characterization of data provenance. In *International Conference on Database Theory*, pages 316–330. Springer, 2001.

[BMR97] Stefano Bistarelli, Ugo Montanari, and Francesca Rossi. Semiring-based constraint satisfaction and optimization. *Journal of the ACM (JACM)*, 44(2):201–236, 1997.

[DG19] Katrin Dannert and Erich 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*. Springer, 2019.

[DK09] Manfred Droste and Werner Kuich. Semirings and formal power series. In *Handbook of Weighted Automata*, pages 3–28. Springer, 2009.

[DMRT14] Daniel Deutch, Tova Milo, Sudeepa Roy, and Val Tannen. Circuits for datalog provenance. In *Proceedings of the 17th International Conference on Database Theory*, pages 201–212, 2014.

[DP02] Brian A Davey and Hilary A Priestley. *Introduction to lattices and order*. Cambridge university press, 2002.

[FFSS11] Diego Figueira, Santiago Figueira, Sylvain Schmitz, and Philippe Schnoebelen. Ackermannian and primitive-recursive bounds with dickson's lemma. In *2011 IEEE 26th Annual Symposium on Logic in Computer Science*, pages 269–278. IEEE, 2011.

[GKL+07] Erich Grädel, Phokion G Kolaitis, Leonid Libkin, Maarten Marx, Joel Spencer, Moshe Y Vardi, Yde Venema, and Scott Weinstein. *Finite Model Theory and Its Applications*. Springer, 2007.

[GKT07] Todd J Green, Grigoris Karvounarakis, and Val Tannen. Provenance semirings. In *Proceedings of the twenty-sixth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, pages 31–40. ACM, 2007.

[Gol99] Jonathan S Golan. *Semirings and their Applications*. Springer, 1999.

[Gre11] Todd J Green. Containment of conjunctive queries on annotated relations. *Theory of Computing Systems*, 49(2):429–459, 2011.

[GT17a] Erich Grädel and Val Tannen. Semiring provenance for first-order model checking. *arXiv preprint arXiv:1712.01980*, 2017.

[GT17b] Todd J Green and Val Tannen. The semiring framework for database provenance. In *Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 93–99. ACM, 2017.

[GT19] Erich Grädel and Val Tannen. Provenance analysis for logic and games, submitted for publication, 2019.

[Huw18] Lukas Huwald. Provenance analysis for logics with team semantics, Master's thesis, RWTH Aachen, 2018.

[Mar76] George Markowsky. Chain-complete posets and directed sets with applications. *Algebra universalis*, 6(1):53–68, 1976.

[Mrk18] Lovro Mrkonjić. Provenance analysis for temporal and dynamic logics, Bachelor's thesis, RWTH Aachen, 2018.

[Ran52] George N Raney. Completely distributive complete lattices. *Proceedings of the American Mathematical Society*, 3(5):677–680, 1952.

[Wal02] Igor Walukiewicz. Monadic second-order logic on tree-like structures. *Theoretical computer science*, 275(1-2):311–346, 2002.