Logic, Semirings, and Fixed Points

Von der Fakultät für Mathematik, Informatik und Naturwissenschaften der
RWTH Aachen University zur Erlangung des akademischen Grades
eines Doktors der Naturwissenschaften genehmigte Dissertation

vorgelegt von

Matthias Ferdinand Naaf, M.Sc.
aus Aachen

Berichter: Universitätsprofessor Dr. Erich Grädel
Professor Dr. Javier Esparza

Tag der mündlichen Prüfung: 30. August 2024

Diese Dissertation ist auf den Internetseiten der Universitätsbibliothek online verfügbar.

Abstract

Semiring provenance is a successful approach in database theory, developed over the last two decades, that provides a unifying framework for various forms of provenance computations. This is achieved by annotating the database with semiring values and then defining a semiring semantics for query languages that combines these values using the semiring operations, ultimately producing a semiring value for each answer of the query. By using different semirings, this approach can cover a wide range of applications: the Boolean semiring yields standard semantics, the natural semiring can be used for counting and bag semantics, semirings on real numbers can model costs or probabilities, and semirings of polynomials can be used to track which facts affect the evaluation of a query.

This approach was originally confined to query languages without negation, but a decade later this limitation was lifted and a semiring semantics for first-order logic was proposed. This provides an interesting generalisation of the classical Boolean semantics to multiple truth values, combining logic and algebra in a unique way. Indeed, algebraic properties, homomorphisms, and universal semirings play an important role in the study of logics under this new semantics. This becomes even more important for the semiring semantics of fixed-point logic (LFP), where additional completeness and continuity assumptions (called *full continuity*) on the semiring are needed to guarantee that both least and greatest fixed points exist and are well behaved. To ensure that the semantics provides meaningful information (from a provenance perspective), it is further assumed that the semirings are *absorptive*. Of particular interest is thus the universal fully-continuous and absorptive semiring $\mathbb{S}^\infty(X)$, both for provenance applications and as an important tool for understanding the semantics.

In this thesis, we advance the study of semiring semantics for fixed-point logic in several directions. An algorithmic aspect concerns the computation of fixed points arising in the evaluation of LFP-formulae. Several techniques have been developed to compute *least* fixed points in semirings, with the generalisation of Newton's method to semirings emerging as the most general technique, but it was not known how to compute *greatest* fixed points. We address this problem by providing a closed-form solution for greatest fixed points of polynomial systems in absorptive, fully-continuous semirings. This is a general result about semirings, but also has implications for the expressive power of LFP compared to infinitary logic under semiring semantics. Further contributions on the algebraic side are generalisations of the universal semiring $\mathbb{S}^\infty(X)$ by coefficients and by infinitely many indeterminates, as well as extensions of semirings by infinitary operations that lay the foundation for semiring semantics on infinite structures and for infinitary logic.

In the direction of semiring provenance, we show in a case study on Büchi games how semiring semantics for LFP can be used for provenance analysis of infinite games. The resulting provenance values provide detailed information about winning strategies and can be used, for instance, to find minimal repairs of the game arena. Lastly, we take the perspective of model theory and study how the classical 0-1 laws about the asymptotic behaviour of logic on random structures can be generalised to semiring semantics of first-order and infinitary logic (and hence LFP) for several different semirings.

Zusammenfassung

Provenienzanalyse in Halbringen ist ein erfolgreicher Ansatz aus der Datenbanktheorie der letzten zwei Jahrzehnte, der mehrere frühere Ansätze zur Provenienzanalyse vereinheitlicht. Dazu werden die Einträge einer Datenbank zunächst mit Halbringwerten beschriftet und anschließend wird eine Halbringsemantik für Datenbankabfragen definiert, die diese Werte mithilfe der Halbringoperationen verrechnet und so einen Halbringwert für jede Antwort der Abfrage liefert. Mittels verschiedener Halbringe kann ein breites Spektrum an Anwendungen abgedeckt werden: der Boolesche Halbring ergibt die konventionelle Semantik, natürliche Zahlen bilden Multimengen ab und reellwertige Halbringe Kosten oder Wahrscheinlichkeiten. Mit Polynomhalbringen lässt sich nachverfolgen, welche Einträge die Auswertung einer Abfrage beeinflussen.

Dieser Ansatz war ursprünglich auf Abfragen ohne Negation beschränkt, was sich ein Jahrzehnt später durch die Einführung einer Halbringsemantik für die Prädikatenlogik geändert hat. Diese Semantik stellt eine interessante Verallgemeinerung der klassischen Booleschen Semantik auf mehrere Wahrheitswerte dar, bei der Logik und Algebra auf neuartige Weise kombiniert werden. Tatsächlich spielen algebraische Eigenschaften, Homomorphismen und universelle Halbringe eine wichtige Rolle für die Analyse dieser neuen Semantik. Das wird noch bedeutender für die Halbringsemantik von Fixpunktlogik (LFP), bei der zusätzliche Bedingungen an die Vollständigkeit und Stetigkeit (*vollständige Stetigkeit*) des Halbrings nötig sind, um Existenz und Verhalten von kleinsten wie größten Fixpunkten sicherzustellen. Damit die Semantik aussagekräftige Informationen liefert (im Sinne der Provenienzanalyse), wird zudem angenommen, dass der Halbring *absorptiv* ist. Damit ist der universelle vollständig stetige und absorptive Halbring $\mathbb{S}^\infty(X)$ von besonderem Interesse, sowohl für die Provenienzanalyse als auch für ein besseres Verständnis der Semantik.

In dieser Arbeit werden Beiträge zur Halbringsemantik für Fixpunktlogik in mehreren Bereichen vorgestellt. Eine algorithmische Fragestellung ist die Berechnung der Fixpunkte, die bei der Auswertung von LFP auftauchen. Es wurden bereits mehrere Methoden entwickelt, um *kleinste* Fixpunkte in Halbringen zu berechnen, wobei sich die Verallgemeinerung des Newton-Verfahrens auf Halbringe als die allgemeinste solche Methode herausgestellt hat. Bisher war jedoch nicht bekannt, wie *größte* Fixpunkte berechnet werden können. Für dieses Problem wird hier eine geschlossene Formel zur Berechnung größter Fixpunkte in absorptiven, vollständig stetigen Halbringen vorgestellt. Das ist zunächst ein allgemeines Ergebnis über Halbringe, hat aber auch Konsequenzen für die Ausdrucksstärke von LFP im Vergleich zu infinitärer Logik in der Halbringsemantik. Weitere Beiträge zu algebraischen Grundlagen sind die Erweiterung von $\mathbb{S}^\infty(X)$ um unendlich viele Unbestimmte und Koeffizienten, sowie die Einführung von unendlichen Operationen in Halbringen, die für Halbringsemantik in unendlichen Strukturen und für infinitäre Logik benötigt werden.

Im Bereich der Provenienzanalyse wird anhand einer Fallstudie für Büchi-Spiele gezeigt, wie sich die Halbringsemantik von LFP für die Provenienzanalyse von unendlichen Spielen nutzen lässt. Die resultierenden Halbringwerte liefern detaillierte Informationen über Gewinnstrategien und können beispielsweise zum Finden minimaler Reparaturen am Spielgraphen verwendet werden. Der letzte Bereich ist die Modelltheorie. Hier wird untersucht, wie sich die bekannten 0-1-Gesetze über das asymptotische Verhalten von Logik auf Zufallsstrukturen für die Halbringsemantik der Prädikaten- und infinitären Logik (und damit auch LFP) in verschiedenen Halbringen formulieren lassen.

Contents

Chapter 1

Introduction

Let us begin in an unconventional way by playing a game. It takes place in the game arena shown below, with the following rules. Starting at the initial position v, in each step one of the two players, let us call them Alice and Bob, chooses a successor position that can be reached by following an edge. It is Alice's turn to choose at the positions drawn as circles, and Bob's turn at the rectangular ones. The resulting sequence of positions is called a *play*, and since every position has at least one outgoing edge, this play is of infinite length. We then say that Alice wins such an infinite play if infinitely many target positions are visited, highlighted in red (this is called the *Büchi winning condition*).

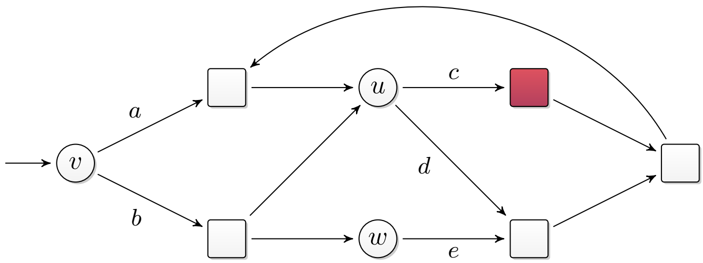

Can Alice win this game? That is, does she have a strategy to select successors in her turns, so that no matter what Bob does, she wins the resulting infinite play? It is not difficult to see that Alice can indeed win. Let us take a closer look at her possible winning strategies. One option is to start by taking edge a at v, and then always choose c at u. This ensures that the play loops in a cycle containing the target position, so that Alice wins. Notice that this strategy makes use of edge a only once, but edge c is taken infinitely often. There are more winning strategies; in particular, Alice can also win without using edge a, by first taking edge b. She must then be prepared to continue from either u or w, depending on Bob's choice. She may select the responses c and e, and when the play again reaches u she always picks c, ending in the same loop as before. Her choices in these two strategies can be summarised as follows:

Strategy 1: $a\colon 1$, $c\colon \infty$

Strategy 2: $b\colon 1$, $e\colon 1$, $c\colon \infty$

These are not the only options. In fact, there are infinitely many winning strategies, as Alice can alternate between choosing c and d at u in any pattern, as long as c is taken infinitely often. However, all of these strategies take each edge at least as often as either Strategy 1 or Strategy 2 (not distinguishing how edges that appear infinitely often are used), so we may say that these are the two simplest strategies.

This game is, of course, a toy example. What we are really interested in is a general approach that not only determines the winner of such a game, but also provides an explanation that tells us *why* or *how* the game is won – in this case, this explanation could be a description of Alice's winning strategies. This approach should be general in the sense that it can also provide explanations for games with different winning conditions, and perhaps even for similar problems such as the evaluation of logical formulae or database queries. The idea is that we start with a formal description of the question at hand, which provides a yes/no answer (in our case, whether Alice wins the game), and then use this description to generate explanations. In the case of Büchi games, it is well-known (e.g., [CGLP15, Wal02]) that the statement "Alice wins from v" can be expressed as a formula in the fixed-point logic LFP:

$$\mathsf{win}_A(v) \coloneqq \big[\,\mathbf{gfp}\, Yy.\ [\mathbf{lfp}\, Zz.\ (Fz \wedge \varphi_{\mathsf{move}}(Y,z)) \vee (\neg Fz \wedge \varphi_{\mathsf{move}}(Z,z)](y)\big](v),$$

where

$$\varphi_{\mathsf{move}}(X,v) \coloneqq (V_A v \wedge \exists w(Evw \wedge Xw)) \vee (V_B v \wedge \forall w(Evw \rightarrow Xw)).$$

This formula may appear complicated at first (and this impression may not be completely inaccurate), but what is important here is that $\mathsf{win}_A(v)$ is a formal description, uniform for all Büchi games, that Alice wins from position v in the game that is represented by the sets V_A and V_B of Alice's and Bob's positions, the binary edge relation E, and the set F of target positions. If we evaluate this formula for the above game, the answer will be: yes, Alice wins from position v. The important insight is that this formula is more than just a yes/no answer – it is a description of *how* this answer is computed. If we can somehow track the evaluation of $\mathsf{win}_A(v)$ on our game, we see which parts of the game are examined to arrive at the conclusion that Alice wins, and this information can constitute an explanation.

This insight is brought to life by the concept of semiring semantics. Roughly, the idea is that we annotate the parts of the game we are interested in, in our case the edges, by so-called *provenance tokens*, which are simply symbolic variables or indeterminates. We then change the semantics of the logic, in this case the fixed-point logic LFP, to use addition $+$ and multiplication $\cdot$ instead of the standard interpretations of the connectives $\vee$, $\wedge$, and the quantifiers $\exists$ and $\forall$. As a result, evaluating $\mathsf{win}_A(v)$ on the above game produces a polynomial expression, composed of the operations $+$, $\cdot$ and of the tokens we used as edge annotations. Formally, these polynomial expressions are elements of a *semiring*, which is an algebraic structure providing the two operations $+$ and $\cdot$. By choosing an appropriate semiring of polynomials, the same formula $\mathsf{win}_A(v)$ that under classical semantics merely tells us "yes, Alice wins", now evaluates to the polynomial (or rather, to some formal expression similar to a polynomial):

$$[\![\mathsf{win}_A(v)]\!] \;=\; ac^{\infty} \;+\; aec^{\infty}.$$

This is exactly the description of the simplest winning strategies we have seen earlier! In other words, we have turned a logical description of a decision problem ("does Alice win?") into an approach to generate explanations ("how does Alice win?") by modifying the semantics, so that we can evaluate formulae in more general algebraic structures which provide more information than just the two classical truth values.

This approach is called *semiring provenance*. It requires to first define a semiring semantics for the logic or query language of interest, and to then choose an appropriate semiring so that the evaluation of a, formerly Boolean, formula or query now produces the information one is interested in, from explanations up to costs or security clearances of the answers to a given query.

Semiring provenance and provenance semirings. The semiring framework for provenance analysis goes back to the seminal work of Green, Karvounarakis, and Tannen [GKT07]. In fact, their approach is similar, although developed independently, to a framework developed by Bistarelli, Montanari, and Rossi ten years earlier [BMR97] for constraint satisfaction problems.

In [GKT07], a semiring semantics is defined for database queries formulated in positive relational algebra (which covers the select ... from ... where statements of the common query language SQL) and datalog (a more expressive language with a least-fixed-point semantics). Starting from a database in which tuples are annotated with semiring values (called K-relations, for a commutative semiring K), these values are propagated through the evaluation of the query, resulting in a semiring value, such as $ac^\infty + aec^\infty$ above, for each answer tuple. This propagation follows the paradigm that *alternative use of information* is interpreted by semiring addition, such as unions and projections, whereas semiring multiplication is used for *joint use of information*, such as conjunctions and natural joins.

Many semirings can be used as annotations, making this a very general approach. First, it is a true generalisation of classical semantics, which is recovered by using the Boolean semiring $\mathbb{B} = (\{0,1\}, \vee, \wedge, 0, 1)$ of standard truth values (here using 0 for false and 1 for true). For a general commutative semiring $(K, +, \cdot, 0, 1)$, we always interpret 0 as *false*, and all nonzero values as *nuances of true*, which is to be understood as truth with additional information. For semirings over natural or real numbers, this information can model bag semantics, probabilities, or costs. By using different semirings of polynomials, called *provenance semirings* in [GKT07], one can further capture several forms of provenance computations for databases, such as why-provenance, incomplete databases, or minimal repairs, thus providing a unified framework for database provenance. This framework has been quite successful in database theory, with many extensions to other query languages and also practical implementations (see, e.g., [GT17b, Gla21] for surveys), and the original paper has since received a *test-of-time* award [LV17].

An important aspect of semiring provenance is the compatibility with semiring homomorphisms. Given a database annotated with values in a semiring K_1, a homomorphism $h\colon K_1 \to K_2$, and a positive relational algebra query, we obtain the same semiring value by either first evaluating the query and then applying h to the resulting semiring annotations of the answers, or by first applying h to the annotations of the input database and then using these to evaluate the query. This also applies to datalog, but requires a continuity assumption on h (to account for the least fixed points expressible in datalog).

Semiring homomorphisms thus provide a measure of how informative a provenance semiring is. For instance, the semiring $\mathbb{N}[X]$ of standard polynomials is the commutative semiring freely generated by the set X, and hence any assignment $h\colon X \to K$ of the indeterminates with values in a commutative semiring K extends to a homomorphism $h\colon \mathbb{N}[X] \to K$. Hence $\mathbb{N}[X]$ is the most informative provenance semiring for positive relational algebra. This is interesting not only for provenance analysis, but also means that in order to understand the evaluation of positive relational algebra in any commutative semiring, it suffices to understand the evaluation in $\mathbb{N}[X]$. Identifying such most informative semirings for various logics and query languages is thus an important step in the study of semiring semantics.

Towards the definition of semiring semantics for logics other than the ones in [GKT07], we thus use the following two properties as guiding principles of our semantics: first, standard semantics should be recovered by using the Boolean semiring; second, semiring semantics should be compatible with homomorphisms.

FO and the obstacle of negation. Semiring semantics has been developed with the goal of provenance analysis of databases, and has thus originally been confined to *finite* databases or structures (more on this later) and to query languages *without negation*. Indeed, neither positive relational algebra nor datalog include a negation operator. This is not a coincidence: semirings only have two operations $+$ and $\cdot$ that are used to interpret (joint or alternative) combinations of information, but without further assumptions semirings do not have a canonical negation or complement operation. This is an obstacle both for considering more expressive query languages on databases, such as stratified datalog, and also for extending semiring semantics to classical logics beyond the database setting.

The obvious candidate for the latter possibility is a semiring semantics for first-order logic (FO), which has been studied by Grädel and Tannen [GT17a, GT]. The idea for such a semantics follows the paradigm of [GKT07]: starting from a relational structure whose atoms are annotated by semiring values, such as the edge-annotated game arena above, semiring addition is used to interpret $\vee$ and $\exists$, while multiplication is used for $\wedge$ and $\forall$. More precisely, a quantifier $\exists x\, \varphi(x)$ is evaluated by first obtaining semiring values for each possible choice of x in $\varphi(x)$, and then taking the sum over these values (this is well-defined for finite structures), and $\forall$ is evaluated analogously by taking the product.

How can the obstacle of negation be overcome for first-order logic? The main idea is to define semiring semantics only for formulae in *negation normal form*, where negations have been pushed inside (using the classical dualities of $\vee$ and $\wedge$, and of $\exists$ and $\forall$) and appear only in front of atoms. The annotation of a relational structure, which is called *semiring interpretation* and usually denoted as π, then has to provide semiring values for all literals; that is, not only for atoms (such as Evw, referring to an edges in a graph), but also for their negations (such as $\neg Evw$, referring to the absence of an edge). Given such a semiring operation π and a formula φ in negation normal form, we can evaluate φ to its semiring value $\pi[\![\varphi]\!]$ using only the semiring operations $\cdot$ and $+$. Intuitively, the interpretation of negation is shifted from the semantics to the interpretation π, and to ensure a reasonable behaviour of negation we thus have to impose constraints on π. Following the intuition that 0 means false and positive semiring values mean truth with extra information, we usually require that for an atom, say Evw, only one of the annotations $\pi(Evw)$ and $\pi(\neg Evw)$ is a positive value, and the other is 0. Semiring interpretations with this property are called *model-defining*, as they induce the unique classical model satisfying exactly those literals annotated with positive values. With these preparations, we obtain a semiring semantics for first-order logic that meets both of our guiding principles: the evaluation in a model-defining

$\mathbb{B}$-interpretation is the same as in standard semantics, and we have compatibility with semiring homomorphisms in the sense that $h(\pi[\![\varphi]\!]) = (h \circ \pi)[\![\varphi]\!]$ for a K_1-interpretation π, a homomorphism $h\colon K_1 \to K_2$, and any formula φ.

The most informative semiring for first-order logic is, again, $\mathbb{N}[X]$. However, it may not be the right provenance semiring for FO. What we can do using $\mathbb{N}[X]$ is, for example, to annotate the edges of a graph with indeterminates in X, so that we learn on which edges the truth of a formula φ depends. But what if φ depends on the *absence* of an edge, using a negative literal $\neg Evw$? We could, of course, also annotate all negative literals with indeterminates, but such a $\mathbb{N}[X]$-interpretation π with $\pi(Evw) = x$ and $\pi(\neg Evw) = y$ is not model-defining and the polynomial $\pi[\![\varphi]\!]$ could contain nonsensical monomials such as xy, which would mean that the truth of φ depends simultaneously on the presence and absence of the edge $v \to w$, a contradiction. What we observe here is that $\mathbb{N}[X]$ does not properly capture the correspondence of positive and negative literals, so we again face an issue with negation. To overcome this second obstacle, Grädel and Tannen propose semirings of *dual-indeterminate polynomials* $\mathbb{N}[X, \overline{X}]$. These are constructed by associating with each indeterminate $x \in X$ a dual indeterminate $\overline{x} \in \overline{X}$, and then defining $\mathbb{N}[X, \overline{X}]$ as the quotient of $\mathbb{N}[X \cup \overline{X}]$ generated by the equations $x \cdot \overline{x} = 0$, effectively filtering out all nonsensical monomials. We would then use the annotation $\pi(Evw) = x$ and $\pi(\neg Evw) = \overline{x}$ so that the correspondence of the literals is reflected in the semiring. Such an interpretation π is not model-defining, but rather *model-compatible*: the truth of Evw is left unspecified, so π corresponds not to one, but to a set of classical structure. This enables applications such as minimal repairs, where one wants to compute how to modify a given structure so that a certain first-order property φ holds. In summary, $\mathbb{N}[X]$ is the most informative semiring from an algebraic point of view, while $\mathbb{N}[X, \overline{X}]$ is the more suitable provenance semiring from the perspective of provenance analysis. In this thesis, we usually take the algebraic point of view, as we often regard the most informative semirings as a tool for our proofs. We use the dual-indeterminate versions only for applications, such as the computation of repairs in Büchi games.

The semiring semantics for FO is the foundation for the semiring semantics of more expressive logics we consider in this thesis. We should note, however, that also first-order logic in semirings is far from being completely understood. While the Boolean semiring gives us back classical semantics, larger semirings with different algebraic properties lead to a very different semantics (see the paragraph on model theory below).

LFP and continuous semirings. An important extension of first-order logic, and the main focus of this thesis, is least fixed-point logic (LFP). It is well-known that, in classical semantics, FO cannot express properties such as transitive closure or connectedness of a graph, although these properties are simple from an algorithmic perspective. This issue is addressed in LFP by adding an iteration mechanism to the logic, in the form of fixed-point formulae $[\mathbf{lfp}\, R\boldsymbol{x}.\ \varphi(R,\boldsymbol{x})](\boldsymbol{y})$ and $[\mathbf{gfp}\, R\boldsymbol{x}.\ \varphi(R,\boldsymbol{x})](\boldsymbol{y})$. The semantics of an **lfp**-formula is defined by an iterative computation of the relation R: starting with $R = \emptyset$, the formula $\varphi(R,\boldsymbol{x})$ defines a relation (in $\boldsymbol{x}$) which becomes the new value of R, and this is repeated until a fixed-point is reached. The syntax of LFP requires φ to be positive in R, so that the evaluation of φ is monotone in R and the iteration is thus guaranteed to reach the *least* fixed point. The semantics of **gfp**-formulae can either be defined analogously, starting from the complete relation and reaching the *greatest* fixed point, or indirectly as negation of an **lfp**-formula (this is also the reason for the name LFP). In particular, having both **lfp**- and

gfp-formula allows us to bring any formula into negation normal form, via the well-known duality of least and greatest fixed points. We illustrate the evaluation of an LFP formula on the following example:

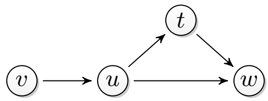

$$\psi(v,w) = [\mathbf{lfp}\, R\,x.\; x = v \lor \exists y(Eyx \land Ry)](w)$$

This formula computes the set of nodes reachable from v, by starting with v and then adding all nodes x with a predecessor reachable from v, and finally checks whether w is contained in the resulting set. The iterative computation of R on the given graph is as follows:

$$\emptyset \quad\mapsto\quad \{v\} \quad\mapsto\quad \{v,u\} \quad\mapsto\quad \{v,u,t,w\} \;\circlearrowleft$$

Since w is contained in this set, the graph satisfies $\psi(v,w)$.

The fixed-point logic LFP and its variants with inflationary and partial fixed points have been extensively studied in finite model theory and descriptive complexity (see, e.g., [GKL+07, EF95] for background). In particular, the iteration mechanism indeed brings the expressive power of LFP closer to the algorithmically feasible problems: it is well-known that LFP captures polynomial time on ordered finite structures. However, LFP is too weak to capture polynomial time on finite structures without a total order; one way to see this are the 0-1 laws that we also consider in this thesis.

How can a semiring semantics for LFP be defined? It turns out that extending the semantics of FO to LFP is rather straightforward. What is challenging are the algebraic and order-theoretic requirements that we have to impose on the semirings so that fixed points exists and provide reasonable information. The study of these requirements and the semiring semantics of LFP was the topic of my Master's thesis [Naa19] and the subsequent publication [DGNT21], inspired by earlier results on semiring provenance for positive LFP and safety games in [GT20]. To define the semantics, we recall the computation of R in the above example. A set over the universe V, such as $\{v,u\}$, can be equivalently viewed as an indicator function $g\colon V \to \mathbb{B}$, so that, e.g., $g(v) = 1$ and $g(t) = 0$. This readily generalises to functions $g\colon V \to K$ for any semiring K, representing annotated sets. For a simple example, consider the semiring interpretation π corresponding to the following graph with edges annotated by distances (or costs):

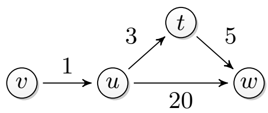

$$\psi(v,w) = [\mathbf{lfp}\, R\,x.\; x = v \lor \exists y(Eyx \land Ry)](w)$$

We view these distances as values in the tropical semiring $\mathbb{T} = (\mathbb{R}_{\geq 0} \cup \{\infty\}, \min, +, \infty, 0)$ over nonnegative real numbers, extended by an element ∞ with the obvious semantics. The iterative computation of R then has the following form, where we represent a function g as a tuple $(g(v), g(u), g(t), g(w))$. Each step corresponds to the evaluation of the inner formula $x = v \lor \exists y(Eyx \land Ry)$ as an FO-formula, with the current annotation of R. This yields:

$$\begin{pmatrix}\infty\\\infty\\\infty\\\infty\end{pmatrix} \quad\mapsto\quad \begin{pmatrix}0\\\infty\\\infty\\\infty\end{pmatrix} \quad\mapsto\quad \begin{pmatrix}0\\1\\\infty\\\infty\end{pmatrix} \quad\mapsto\quad \begin{pmatrix}0\\1\\4\\21\end{pmatrix} \quad\mapsto\quad \begin{pmatrix}0\\1\\4\\9\end{pmatrix} \;\circlearrowleft$$

Notice that we start with the function mapping all nodes to the zero-element of the semiring, which, counterintuitively, is ∞ in this case. The entry for w is 9, and we thus get $\pi[\![\psi(v,w)]\!] = 9$, which is indeed the shortest distance from v to w. A more formal definition of the semantics is given in Chapter 5, here we focus more on the algebraic side.

Which semirings are suitable for fixed-point logics? A first step is to equip semirings with an order, so that we can speak about *least* or *greatest* fixed points. This has already been considered in the original work of Green, Karvounarakis, and Tannen [GKT07] for datalog. Indeed, datalog can be seen as a fragment of LFP with only least fixed points, including for instance our reachability example. They propose to use the *natural order* induced by the addition operation of a given semiring. That is, $a \leq b$ if there is an element c with $a + c = b$. Only semirings where this defines an order are considered, which excludes rings (due to their negative elements). The advantage of using the natural order is that, by definition and distributivity, both semiring operations are monotone. For many semirings, the natural order coincides with the standard order, for instance on natural numbers $(\mathbb{N}, +, \cdot, 0, 1)$. In the tropical semiring $(\mathbb{R}^{\infty}_{\geq 0}, \min, +, \infty, 0)$, however, the natural order is the inverse of the standard order on real numbers, since addition is defined as minimum. In any case, the zero-element of the semiring is the least element of the natural order. A greatest element does not always exist, and hence additional completeness assumptions are needed. Indeed, the monotone function $n \mapsto n+1$ does not have a least fixed point in $\mathbb{N}$, so instead we have to consider its completion to $\mathbb{N}^{\infty} = \mathbb{N} \cup \{\infty\}$. The approach of [GKT07] is to work with *ω-continuous* semirings. These require that every ω-chain $a_0 < a_1 < a_2 < \dots$ of semiring elements has a supremum $\bigsqcup_{i<\omega} a_i$ (completeness) and that both semiring operations commute with these suprema (continuity). It follows that every operator built from the semiring operations, and hence every operator induced by a datalog program, satisfies the same continuity property, and thus has a well-defined least fixed-point. The notion of ω-continuous semirings has also been used in other contexts where least fixed points are considered, for example [Kui97, EKL10].

Semiring semantics for LFP needs stronger assumptions on the semirings, to guarantee the existence of both greatest and also nested fixed points. Indeed, the iterative computation of greatest fixed points forms a decreasing chain, and we must require infima for such chains. In addition, continuity of the semiring operations does not necessarily imply continuity of the operators induced by nested fixed points (with an alternation between an outer **lfp**- and an inner **gfp**-formula), so we must ensure that least and greatest fixed points exist for all monotone operators, not just for the continuous ones. The common way to achieve this is to ensure that the natural order is a complete lattice. Here we are a bit more general and only assume suprema and infima of chains instead of arbitrary sets, where a chain is any totally ordered set (of any cardinality). However, in many cases, and in particular for the absorptive semirings introduced below, this is actually equivalent to a complete lattice. In addition to this completeness assumption for chains (which would suffice for a well-defined semantics), we also assume *full continuity* of both semiring operations. This is similar to ω-continuity, but requires that $+$ and $\cdot$ commute with both suprema and infima of chains. Full continuity seems to be crucial to prove any interesting results about fixed points, and is satisfied by all relevant semirings with the chain-completeness property.

The semiring semantics of LFP is thus well-defined for all fully-continuous semirings, such as $\mathbb{N}^{\infty}$ or $\mathbb{T}$. It satisfies our two guiding principles: standard semantics is obtained for the semiring $\mathbb{B}$, and homomorphisms that are fully-continuous are compatible with our semantics. The status of a most informative provenance semiring is more interesting, and hints at a deeper problem. For datalog, this provenance semiring is given by the semiring of formal power series $\mathbb{N}^\infty[\![X]\!]$, which is both ω-continuous and fully continuous. However, the homomorphisms induced by assignments to the indeterminates are only ω-continuous and not fully continuous, and hence do not preserve greatest fixed points. This means that $\mathbb{N}^\infty[\![X]\!]$ is not the right provenance semiring for LFP, and it is not clear if such a semiring exists. A second observation is that the iterative computation of greatest fixed-points in $\mathbb{N}^\infty$ is in many cases trivial: it starts at the largest value ∞, and since both $+$ and $\cdot$ are increasing with respect to the natural order, this is often a fixed point. These issues are resolved in [Naa19, DGNT21] by restricting the semantics to *absorptive* semirings, defined by the equality $a + ab = a$. In these semirings, the neutral element 1 is the greatest element and multiplication is decreasing, which leads to a behaviour that is closer to the duality of $\vee$ and $\wedge$, and of **lfp**- and **gfp**-formulae known from standard semantics. The most informative semiring is then the semiring of generalised absorptive polynomials $\mathbb{S}^\infty(X)$, which is a quotient of standard polynomials $\mathbb{N}[X]$ by absorption, but extended by the exponent ∞. This is the free absorptive, fully-continuous semiring (in terms of fully-continuous homomorphisms), and examples show that it can provide useful provenance information for both least and greatest fixed points. We follow this approach and only consider absorptive, fully-continuous semirings for semiring provenance of LFP in this thesis, such as the tropical semiring that we used in our example.

Logic and games. There is a deep connection between logic and games in classical model theory: the semantics of FO or LFP can be understood through the associated model-checking games, which are acyclic games and parity games, respectively, and the power of first-order logic to distinguish two structures can be understood through Ehrenfeucht-Fraïssé games. We refer to [Grä11, EF95] for background on the classical theory. In the other direction, winning regions (and even winning strategies) in Büchi and parity games with bounded number of priorities can be defined in LFP [CGLP15], as we have seen for the introductory example.

This connection also applies to semiring semantics and is interesting in two ways, with the unifying theme of *sum-of-strategies* theorems. One way is to understand semiring semantics through games. The first result of this kind in [GT17a, GT] is in fact not about games, but about evaluation strategies in the form of proof trees for first-order logic. Given a semiring interpretation π, one assigns to each proof tree the semiring value that results from multiplying the annotations of all literals (according to π) occurring in the tree. This yields a sum-of-strategies theorem, saying that the the semiring value $\pi[\![\psi]\!]$ of a FO-sentence ψ is equal to the sum over the values of all proof trees. A similar sum-of-strategies theorem holds for LFP, by replacing proof trees with winning strategies in the associated model-checking game. We recall that the positions of the model-checking game for the evaluation of $\pi[\![\psi]\!]$ correspond to the instantiated subformulae of ψ, where the terminal positions are precisely the instantiated literals. The semiring value of a strategy can thus be defined in a similar way as for proof trees, by multiplying the annotations of all literals that can be reached by plays consistent with the strategy. The sum-of-strategies theorem for LFP then states that the value $\pi[\![\psi]\!]$ is equal to the sum over the values of all winning strategies in the model-checking game. This was first observed for positive LFP (with only least fixed points) in [GT20], and was later extended to full LFP in [Naa19, DGNT21] with more involved arguments to handle greatest fixed points and nesting of fixed-point formulae.

The other way is to define semiring provenance for games through semiring semantics of logics. In this setting, each position of a game receives a semiring value, with a similar goal as in our introductory example: to understand *how* the winning player wins from this position, by gathering information about the moves the player uses and the target positions they reach. For simple classes of games, this provenance information can be defined explicitly, by a recursive bottom-up approach for acyclic games, or as a least or greatest solution to a polynomial equation system for reachability and safety games [GT20]. For infinite games with more complex winning conditions, such as Büchi or parity games, an explicit definition becomes challenging, and it is more convenient to define the provenance information through a formula, such as $\mathsf{win}_A(v)$ for Büchi games. In any case, the goal is again to obtain a sum-of-strategies result, saying that the semiring value of a position v in a game corresponds to the values of the winning strategies from v, and thus offers meaningful provenance information.

Model theory. The field of semiring semantics originated in provenance analysis for databases, and so far we have mostly stayed true to this origin by looking at the semantics from a provenance perspective. A different perspective is that of (finite) model theory, where we forget about the original intention and instead study the properties of semiring semantics and its expressive power on its own. Besides the intrinsic interest, this may also lead to new insights about the proofs of common model-theoretic results: they all hold for semiring semantics in the Boolean semiring, but many fail for more general semirings, so one can try to pinpoint the algebraic properties of semirings needed for each individual proof.

We should note, however, that the expressive power of semiring semantics behaves quite differently from standard semantics. The reason lies in its provenance (no pun intended): the original goal was to turn a formula into an explanation, by starting with an annotated structure and tracking these annotations through the evaluation of the formula. It was not intended that semiring semantics could be used to express properties of annotated structures. This manifests itself in a key limitation of semiring semantics. Formulae evaluate to semiring values, but they cannot talk about these values, nor can they compare the values assumed by two subformulae. And indeed, a study of the distinguishing power of FO under semiring semantics in [GM21] paints a rather complex picture. Already for simple min-max semirings, or perhaps especially for those, there are examples of non-isomorphic interpretations that are indistinguishable in semiring semantics. And even for semirings where such counterexamples do not exist, such as the tropical semiring, one can show that the isomorphism class of a given interpretation can, in general, not be described by a single sentence, in contrast to the characteristic sentences in classical semantics.

This limitation should not prevent the study of model theory for semiring semantics, quite the contrary. The different behaviour and the inability of semirings to directly interpret negation often require new ideas for the generalisation of classical results. One example is the study of locality properties in [BGN23a], where the generalisation of Gaifman locality required a new construction of Gaifman normal forms without the use of negation. This not only lead to a proof for semiring semantics, but also to a strengthening of Gaifman's theorem for classical semantics. The model-theoretic perspective on semiring semantics has only recently gained attention; we survey first fundamental results in Section 7.1.

1.1 Contributions

As a summary of the introduction, we may say that semiring provenance combines logic and algebra to produce explanations of logically definable properties and queries. The contributions of this thesis cover both the algebraic and the logical side, as well as their combined application to infinite games.

Algebraic foundations. We begin with the algebraic side. For fixed-point logic, the required algebraic and order-theoretic properties of semirings have already been studied in [Naa19, DGNT21] and are summarised in Chapter 3, resulting in the two constraints of *full continuity* and *absorption*, with generalised absorptive polynomials $\mathbb{S}^\infty(X)$ as most informative semiring. Our first contribution is the analysis of the required properties for extending semiring semantics to infinitary logic and to infinite structures. In both cases, a notion of infinitary semiring operations is needed, either for the interpretation of infinitary conjunctions and disjunctions, or for the evaluation of quantifiers as sums or products over infinite universes. We thus consider extensions of semirings, called *infinitary semirings*, by the infinitary operations $\widehat{\sum}_{i\in I} a_i$ and $\widehat{\prod}_{i\in I} a_i$ that denote sums and products over families $(a_i)_{i\in I}$, where I is an arbitrary, possibly infinite, index set. Notice that we use families instead of just sets to account for non-idempotent operations. There are two philosophies to define such infinitary operations. One is the axiomatic one, where we characterise the operations through a list of axioms, and can then prove results about all infinitary semirings that satisfy these axioms. The other one is to identify a class of semirings that admits a natural explicit definition of infinitary operations, and then study the properties of these operations. We follow both philosophies, by providing an axiomatic definition of infinitary semirings and explicit definitions for the class of absorptive, fully-continuous semirings.

What are reasonable axioms for the new operations $\widehat{\sum}_{i\in I} a_i$ and $\widehat{\prod}_{i\in I} a_i$? There are certain obvious ones. If I is finite, then both operations should coincide with the finite semiring addition and multiplication. Moreover, we always assume $\widehat{\sum}$ and $\widehat{\prod}$ to be commutative (this also justifies our definition in terms of families $(a_i)_{i\in I}$ with an index *set* I, instead of using ordered sequences). Adapting the other properties of the semiring operations such as associativity and distributivity turns out to be more interesting, and we propose weak and strong versions of both. For associativity, we can either require that families can be split in two parts without changing the overall sum (or product), or we can consider arbitrary partitions into (possibly) infinitely many subfamilies. This latter form of strong associativity (or partition invariance) is quite powerful and implies, for instance, commutativity. However, it turns out to hold for all natural definitions of infinitary operations we consider, and we thus make it a requirement for infinitary semirings. This is different for distributivity. Here we require that finite multiplication, or infinitary multiplication in the strong version, distributes over infinite sums. Strong distributivity fails in important semirings, and hence only weak distributivity becomes a strict requirement. Further properties concern neutral and idempotent elements, infinite powers of an element, monotonicity, and a notion of compactness. We identify a suitable subset of these constraints that we propose as definition of an infinitary semiring, thus laying the algebraic foundations for a study of semiring provenance on infinite structures.

In the spirit of the second philosophy, we give natural definitions of $\widehat{\sum}$ and $\widehat{\prod}$ for the class of absorptive, fully-continuous semirings, by taking suprema or infima over all finite subfamilies.

These satisfy all requirements of infinitary semirings, and in fact all of the possible axioms we consider except for strong distributivity. The semiring $\mathbb{S}^\infty(X)$ then becomes the universal infinitary semiring that is both absorptive and fully-continuous, and, unlike the general case, it additionally satisfies strong distributivity. These infinitary operations are important for several reasons. First, many semirings for provenance applications are absorptive, such as the tropical semiring for costs, so these are interesting for provenance applications on infinite structures. On finite structures, they are the foundation for the semantics of infinitary logic and the embedding of LFP. Lastly, they can lead to more elegant arguments in proofs, in particular for the analysis of infinite games where we may need to sum or multiply over an infinite number of strategies or plays. This natural definition of infinitary operations does not work for larger classes of semirings, and it is not clear if a uniform definition is possible. Instead, we are interested in identifying the most informative infinitary semiring. Of course, the homomorphisms we consider should commute with the infinitary operations, so that they are compatible with semiring semantics on infinite structures. We prove that the semiring $\mathbb{N}^\infty_\infty[\![X]\!]$, the extension of formal power series $\mathbb{N}^\infty[\![X]\!]$ by the exponent ∞, is the freely generated infinitary semiring with strong distributivity. This leaves open whether a free infinitary semiring exists, but it does provide us with a tool to better understand all infinitary semirings that satisfy strong distributivity, and this can further be relaxed to require strong distributivity only for infinite products up to a certain cardinality.

These algebraic foundations for semiring provenance on infinite structures have been developed jointly with Sophie Brinke, Lovro Mrkonjić, and Erich Grädel. An account of this work, including also a study of semiring semantics for FO on infinite structures, will appear in [BGMN24]. The analysis of axioms for infinitary semirings (Section 3.4.1) was carried out in cooperation. I further contributed the results on absorptive semirings (Section 3.4.2 and Section 3.5.3) and on the freely generated semiring $\mathbb{N}^\infty_\infty[\![X]\!]$ (Section 3.4.3).

A second contribution to the algebraic foundations is the study of most informative or, more precisely, freely generated semirings. Given the importance of absorptive, fully-continuous semirings for fixed-point logics, for which the most informative semiring is $\mathbb{S}^\infty(X)$, we study two extensions of $\mathbb{S}^\infty(X)$. The first again aims at semiring provenance for infinite structures and extends $\mathbb{S}^\infty(X)$ by permitting infinitely many indeterminates. Since indeterminates are used as provenance tokens to track atoms of a given structure, we are naturally interested in an infinite set X to deal with infinite structures, but this leads to a number of complications. First, it requires a more general quotient construction of (generalised) absorptive polynomials, resulting in the semiring $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ for an infinite set $\boldsymbol{\mathcal{X}}$. This is only a technical issue; more severe is the different behaviour of the continuity properties and the associated definitions of infinitary operations, which we illustrate in a number of examples. Most importantly, $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ is no longer fully continuous: while addition remains fully continuous, and multiplication still commutes with suprema, the continuity of multiplication fails for infima. We show that one can nevertheless define a reasonable notion of infinitary operations, satisfying most, but not all, requirements of infinitary semirings. The emerging picture is that $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ is not well-suited for semiring provenance of LFP, but can still be interesting for simpler logics on infinite structures. We note that all of the issues disappear for the multiplicatively idempotent version of $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$, which is the free completely-distributive lattice $\mathrm{PosBool}(\!|\boldsymbol{\mathcal{X}}|\!)$. This semiring suffices for many provenance applications, such as the computation of minimal repairs, and may be an interesting candidate for semiring provenance in the infinite.

The second way we extend $\mathbb{S}^\infty(X)$ and its simpler variants PosBool(X) and $\mathbb{S}(X)$ is by adding coefficients from an absorptive, fully-continuous semiring K. One motivation for this is simply that standard polynomials $K[X]$ are not absorptive, even if K is, and depending on the situation it can be more convenient to work with an absorptive version of $K[X]$. A further motivation is that we can then combine abstract provenance information (i.e., monomials over provenance tokens) with values such as costs (coefficients in the tropical semiring), so that if there are different proofs of a formula, or different ways to derive an answer, we get information about the cost of each alternative. We develop a unified theory of absorptive polynomials with coefficients, of which the semirings $\mathbb{S}^\infty(X)$, PosBool(X), and $\mathbb{S}(X)$ are the special case with coefficients in $\mathbb{B}$. Our main results are two universal properties, one for finite and one for infinite coefficient semirings, with quite different arguments.

These two extensions of $\mathbb{S}^\infty(X)$ have not been published before. The semirings $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ and PosBool$(\!|\boldsymbol{\mathcal{X}}|\!)$ may be of interest for future work, while absorptive polynomials with coefficients are already used in Chapter 7 for a more elegant proof of the 0-1 laws for lattice semirings.

Fixed-point computation. An important aspect of the semiring semantics for LFP that has been left open is the question how the least and greatest fixed points for the evaluation of **lfp**- and **gfp**-formulae can be computed. A naive computation is trivial in standard semantics: for a finite structure on n elements, the iterative computation of an r-ary relation can take at most n^r steps, as there are only so many possible relations. This is, a priori, not true for semiring semantics. For example, recall the distance computation from the introduction:

$$\begin{pmatrix}\infty\\\infty\\\infty\\\infty\end{pmatrix} \mapsto \begin{pmatrix}0\\\infty\\\infty\\\infty\end{pmatrix} \mapsto \begin{pmatrix}0\\1\\\infty\\\infty\end{pmatrix} \mapsto \begin{pmatrix}0\\1\\4\\21\end{pmatrix} \mapsto \begin{pmatrix}0\\1\\4\\9\end{pmatrix} \circlearrowleft$$

Notice that the entry 21 is updated to 9. As the tropical semiring is infinite, there could in principle be an arbitrary number of such updates. This is clearly not the case in this example, as the values correspond to paths in the graphs, which are completely explored after n steps. (In fact, the n^r upper bound also holds for **lfp**-iterations in our setting, but this is not a trivial observation.)

Leaving fixed-point logic aside, the evaluation of a single **lfp**-formula can be formulated as the more general problem of computing the least solutions to a polynomial equation system of the form $(X_i = P_i(X_1, \dots, X_n))_{1 \le i \le n}$, where the P_i are polynomials with coefficients in an absorptive, fully-continuous semiring. Indeed, for a formula $[\mathbf{lfp}\, R\,\boldsymbol{x}.\ \varphi(R,\boldsymbol{x})](\boldsymbol{y})$, the polynomials P_i are induced by the evaluation of the inner FO-formula φ in a semiring interpretation π, with indeterminates for the annotations of R. Several techniques have been developed to solve such equation systems, see [HK99, EKL10], and also [GM08] for linear systems. Most relevant for us is the work of Esparza, Kiefer, and Luttenberger [EKL10] who generalise Newton's method (to approximate zeroes of real-valued functions) to ω-continuous semirings, resulting in a (possibly infinite) iteration scheme that is guaranteed to converge to the least solution of a given polynomial system. They introduce the concept of *derivation trees* of a polynomial system, inspired by formal grammars, and conclude by a detailed analysis of these trees that the Newton iteration always converges in n steps if the semiring is idempotent (and commutative). Interestingly, the Newton iteration in idempotent semirings is identical to the iteration defined by Hopkins and Kozen [HK99], but they were only able to prove an upper bound of $\mathcal{O}(3^n)$ on the number of iteration steps, while the improved bound of [EKL10] is tight. These results apply in particular to our setting of absorptive, fully-continuous semirings, thus answering our question for least fixed points.

The remaining question is how **gfp**-formulae can be evaluated or, equivalently, how the *greatest* solution to an equation system $(X_i = P_i(X_1, \dots, X_n))_{1 \le i \le n}$ over an absorptive, fully-continuous semiring can be computed. This has, to the best of my knowledge, not been considered before. Notice that we cannot just dualise the **lfp**-computation (since the semiring operations are not duals of another). Indeed, the following example illustrates that greatest fixed points behave differently:

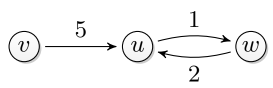

$$X_v = 5 + X_u \qquad X_u = 1 + X_w \qquad X_w = 2 + X_u$$

The polynomial system is induced by the formula $[\mathbf{gfp}\, R\,x.\, \exists y(Exy \wedge Ry)](v)$ that asserts the existence of an infinite path from v. The naive iterative computation of the greatest solution in the tropical semiring leads to the following infinite iteration,

$$\begin{pmatrix}0\\0\\0\end{pmatrix} \mapsto \begin{pmatrix}5\\1\\2\end{pmatrix} \mapsto \begin{pmatrix}6\\3\\3\end{pmatrix} \mapsto \begin{pmatrix}8\\4\\5\end{pmatrix} \mapsto \begin{pmatrix}9\\6\\6\end{pmatrix} \mapsto \dots$$

which converges to the greatest solution $X_v = X_u = X_w = \infty$. Notice that this iteration is infinite despite the finite graph (which is not possible for **lfp**-iterations), but also clearly exhibits a repeating pattern.

A key contribution of this thesis is the following result, which turns this pattern into a closed-form expression to compute greatest solutions, and hence to evaluate **gfp**-formulae. If F is the operator induced by a polynomial equation system in n indeterminates over an absorptive, fully-continuous semiring, or equivalently an operator induced by a FO-formula $\varphi(R, \boldsymbol{x})$, then

$$\mathbf{lfp}(F) = F^n(\mathbf{0}), \qquad \mathbf{gfp}(F) = F^n(F^n(\mathbf{1})^{\boldsymbol{\infty}}).$$

Here, a^∞ denotes the infinitary power of an element a that is defined in all absorptive, fully-continuous semirings as the infimum of the decreasing chain $a \ge a^2 \ge a^3 \ge \dots$. We note that this result applies to the general setting of polynomial equation systems $(X_i = P_i(X_1, \dots, X_n))_{1 \le i \le n}$ and may be of independent interest beyond the semiring semantics of LFP.

For the intuition behind this result, think of a greatest solution as an infinite object, such as an infinite path in a graph (in the above example, the path $v(uw)^\omega$). Because of absorption, it suffices to consider "lasso" paths, that consist of a reachability prefix (here v) followed by a single repeated loop (here uw). The expression $F^n(F^n(\mathbf{1})^{\boldsymbol{\infty}})$ matches this intuition: computing $F^n(\mathbf{1})$ explores all paths of length n, which covers in particular all loops, and $F^n(\mathbf{1})^{\boldsymbol{\infty}}$ is the value of repeating these loops infinitely often. This effectively collapses the infinite computation to a single application of the infinitary power (which is usually easy to compute, for instance $a^\infty = \infty$ for all $a \ne 0$ in the tropical semiring). The outer part $F^n(\dots)$ then accounts for the reachability prefix, which can again take up to n steps. It may appear as if a lot of double counting happens in this computation, but this all vanishes due to absorption. To turn this intuition into a proof, we follow [EKL10] and use derivation trees to represent $\mathbf{gfp}(F)$. Because of absorption, it then suffices to consider trees of a certain regular shape, corresponding to the lasso path intuition. As a byproduct of this argument, we reprove the observation of [EKL11] that least fixed points can be computed by just performing n steps of the native iteration.

This result has several interesting implications. The obvious application is the evaluation of LFP-formulae. If the arity of the fixed-point relations is bounded, this can now be done using a polynomial number of semiring operations (involving the infinitary power) for a single **lfp**- or **gfp**-operator. Whether these operations are efficient of course depends on the semiring; it is not difficult to find graphs on which a simple reachability formula evaluates to an absorptive polynomial of exponential length in $\mathbb{S}^\infty(X)$ or PosBool(X). Intuitively, the only overhead compared to the (naive) evaluation of LFP in classical semantics is the computation of the semiring operations, but not the number of iterations that have to be computed for fixed points. Although this may not be obvious, the fixed-point computation also works for nested fixed points, as these also induce (a slightly generalised form of) polynomial equation systems. An interesting question for future work is whether recent quasipolynomial-time algorithms for the evaluation of nested fixed points in the Boolean semiring [ANP21, HS21] can be generalised to our setting in absorptive semirings.

A second implication is that LFP can be embedded into infinitary logic, under semiring semantics in absorptive, fully-continuous semirings (Section 5.3). The fixed-point computation is needed for two reasons. First, it shows that the number of iterations does not depend on the cardinality of the semiring, so that countable infinitary conjunctions and disjunctions suffice for the embedding. The second point is more subtle: disjunctions are not interpreted as infima, but rather as (infinite) products. We thus cannot express the infimum of an iteration directly, but we can express its infinitary power, and this is all we need for $F^n(F^n(\mathbf{1})^{\boldsymbol{\infty}})$. We briefly sketch a third implication that may be interesting for applications: by turning the closed-form computation into a circuit, we can extend the polynomial-size circuit representations for semiring semantics of datalog [DMRT14] to stratified datalog (Section 5.4).

The fixed-point computation has been published at RAMiCS 2021 [Naa21a]. Chapter 4 is based on the corresponding full version [Naa21b]. The applications for the embedding of LFP and stratified datalog have not been published before.

Infinite games. Coming back to our original motivation of computing an explanation why Alice wins our example game, we present a detailed case study of semiring provenance for Büchi games, showcasing how the semiring semantics and the sum-of-strategies theorem for LFP of [Naa19, DGNT21] can be put to use for the provenance analysis of infinite games. The reason to focus on Büchi games is that these are the simplest infinite games with a nontrivial winning condition for infinite plays. As a consequence, the characterisation of the winning region (of either player) requires an alternation of a least and a greatest fixed point (see the definition of $\mathsf{win}_A(v)$), and our goal is to see to what extent the approach of semiring provenance is still applicable in this setting. Since Büchi games are relevant in formal verification, this may also be a first step for future applications of semiring provenance to verification. We consider the following three provenance applications: *strategy tracking*, where we annotate the edges and want to obtain detailed information about the winning strategies from a given position (as in the introduction); *minimal repairs*, where we want to find minimal modifications to the game arena so that Alice becomes the winner of the game;

and *cost computation*, where we associate costs with each edge and want to compute the minimal cost that it takes to win.

To track strategies, we use a semiring interpretation π_{strat} that annotates the edges by indeterminates, and then evaluate $\mathsf{win}_A(v)$ in the most informative provenance semiring $\mathbb{S}^\infty(X)$. We show that the sum-of-strategies theorem for LFP implies a sum-of-strategies result also for Büchi games, so that the semiring value $\pi_{\mathsf{strat}}[\![\mathsf{win}_A(v)]\!]$ of a position v is equal to the sum over the values of all winning strategies, where the value of a strategy is the product over all edges that it uses. Because of absorption, the values of some strategies are absorbed; we call the strategies that survive *absorption-dominant*. Interestingly, this is a rather natural notion of strategies that may be of independent interest: if we view the positional strategies (that make a decision for each position, without considering the history of a play) as the strategies that minimise the set of moves, then absorption-dominant strategies are those that minimise the *multiset* of moves. We provide a detailed analysis of these strategies and show that they lie between the positional and the persistent strategies, where the latter are strategies that make a unique choice for each position within a play (but not necessarily a unique choice across different plays). This shows that strategy tracking works as desired, telling us precisely which edges, and how often, the simplest winning strategies use. We should note, however, that the provenance information we get depends on choosing the right LFP-formula $\mathsf{win}_A(v)$. Indeed, formulae that are equivalent in classical semantics need not be equivalent in semiring semantics, and it depends on the semiring how robust this approach is to syntactic modifications of the formula.

For minimal repairs, we do not make use of the sum-of-strategies theorem. Instead, we consider a model-compatible interpretation π_{repair} in the dual-indeterminate semiring $\mathrm{PosBool}(X, \overline{X})$ that uses indeterminates to annotates those edges and target positions that can be modified as part of the repair. We then show, by using homomorphisms into the Boolean semiring, that the value $\pi_{\mathsf{repair}}[\![\mathsf{win}_A(v)]\!]$ correctly describes all minimal sets of modifications that let Alice win the repaired game. This argument via homomorphisms is more robust than the computation in $\mathbb{S}^\infty(X)$, in the sense that it produces correct results for all LFP-formulae that are equivalent to $\mathsf{win}_A(v)$ in classical semantics.

The computation of costs reveals a limitation of semiring provenance for infinite games. While the evaluation in the tropical semiring indeed produces values that, by the sum-of-strategies theorem, can be understood as the sum over the costs of all winning strategies, the induced notion of the cost of a strategy is not necessarily a natural one. In addition, we show that two more natural definitions of costs cannot directly be expressed through semiring semantics. This does not mean that all of the typical provenance applications fail: these issues do not arise for lattice semirings, which are absorptive semirings with an idempotent multiplication, so that applications such as the computation of the required access privilege to win a game with restricted access to the edges are still possible.

The techniques used in this case study also apply to parity games, by a straightforward adaptation of the proof of the sum-of-strategies theorem. Provenance information for further games can be defined, for instance, by memory reductions to parity games. The approach for minimal repairs in $\mathrm{PosBool}(X, \overline{X})$ is even more general and applies to all games with LFP-definable winning regions.

The case study for Büchi games is joint work with Erich Grädel and Niels Lücking. This work began with Niels Lücking's Bachelor's thesis that I co-supervised. It has been presented at GandALF 2021 [GLN21] and an extended version has been published in LMCS [GLN24]. All three of us contributed to the analysis of absorption-dominant strategies. I further contributed the proofs of the sum-of-strategies theorem and for minimal repairs. Section 6.3 is based on the extended version [GLN24].

Zero-one laws. On the logical side, we consider generalisations of the well-known 0-1 laws of classic model theory, concerning the behaviour of first-order logic (and also infinitary logic) on random structures. Given a fixed FO-sentence ψ in classical semantics, we can consider for each n the probability that ψ holds on a random graph on n nodes, where each edge is included with uniform probability p, or more generally whether ψ holds on a random relational structure on n nodes. The famous 0-1 law then states that for increasing n, this probability converges (exponentially fast) to either 0 or 1, so that each sentence ψ is either almost surely true or almost surely false on random structures. There is further a unique countable structure, called the *rado graph* in the case of random graphs, that satisfies exactly those sentences that are almost surely true on random structures. One implication of the 0-1 law is that properties with a different convergence behaviour are not expressible in first-order logic; a typical example is even cardinality, for which the probability oscillates between 0 and 1 as n grows. This result has been proved independently by Glebskii, Kogan, Liogon'kii, and Talanov [GKLT69] using a quantifier elimination argument, and in a more model-theoretic fashion by Fagin [Fag76] using the theory of extension axioms and compactness. We show that the 0-1 law can be generalised to semiring semantics, using a combination of these two proof strategies: we begin by adapting the extension axioms to semirings, but instead of Fagin's application of the compactness theorem, we then use a description of formulae by polynomials that facilitates a quantifier elimination argument.

We first need an appropriate notion of random semiring interpretations. The general idea is that for each edge Evw in a graph (or each atom in a relational structure), we first flip a coin to decide whether the positive atom Evw or its negation $\neg Evw$ should receive a positive value (the other one is mapped to 0, so that the interpretation is model-defining). This positive value is then chosen according to a fixed probability distribution $p\colon K \setminus \{0\} \to [0,1]$ on the semiring K, where we assume that each nonzero element occurs with a positive probability for a finite semiring (infinite semirings require further assumptions). Instead of asking for the probability that φ is true on a random graph with n nodes, we then ask for the probability that φ evaluates to a specific semiring value $j \in K$ on a random semiring interpretation with n nodes. We denote this probability as $\mu_{n,p}[\pi[\![\psi]\!] = j]$; the 0-1 law for a semiring K then states that the limit $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] = j]$ exists and is either 0 or 1. We are especially interested in the semiring values j for which the asymptotic probability 1 is assumed, which we call the *almost sure valuation* of ψ.

Our results show that a 0-1 law holds for all lattice semirings (which are absorptive semirings with idempotent multiplication, or, equivalently, bounded distributive lattices), with mild constraints on the probability distributions for infinite lattice semirings. This further implies a partial 0-1 law for absorptive semirings, and by a slight adaptation of our arguments we also get a 0-1 law for the natural semiring $\mathbb{N}$ and its completion $\mathbb{N}^\infty$. These 0-1 laws hold for both FO and, in cases where the semiring admits infinitary operations, also for the finite-variable fragment $L^\omega_{\infty\omega}$ of infinitary logic. For absorptive semiring or lattice semirings that are fully-continuous, they further hold for LFP via the embedding into $L^\omega_{\infty\omega}$. The general strategy to obtain these results is as follows. We first formulate suitable extension properties as an analogues of Fagin's extension axioms and show that these almost surely hold on random structures (due to the limitation that formulae cannot talk about semiring values, these are metatheoretic properties instead of FO-sentences). The new ingredient is a polynomial representation of formulae, or more precisely, of their almost sure valuations, which is essentially an algebraic encoding of a quantifier elimination argument. The 0-1 law is then obtained by showing inductively that these polynomials correctly describe the semiring value of the associated formula in any semiring interpretation with the appropriate extension property. What is interesting is the distribution of almost sure valuations. In classical semantics, both almost surely true and almost surely false sentences exist, but in semiring semantics this depends on the semiring: in $\mathbb{N}$ and $\mathbb{N}^\infty$, rather trivial constructions show that every semiring value can occur as almost sure valuation, while in lattice semirings only 0, the greatest element 1, and the smallest positive element $\varepsilon > 0$ (if it exists) can be almost sure valuations.

Following the classical theory, we propose several extensions of these results for future work, by considering stronger logics or more general probability measures (see Section 7.1.3 for a discussion). An interesting problem that remains open is whether a 0-1 law holds for semiring semantics of FO over all semirings, or at least over all finite semirings.

The results on 0-1 laws for first-order logic in semiring semantics have been obtained by joint work with Erich Grädel, Hayyan Helal, and Richard Wilke, which has appeared at LICS 2022 [GHNW22a]. These results were inspired by preliminary work in the Master's thesis of Hayyan Helal, co-supervised by Richard Wilke. The idea to combine polynomials and extension properties, the complexity considerations (Section 7.2.4), and the result about the random countable semiring interpretation (Section 7.5) are due to Erich Grädel. My contributions include the correctness proofs of the polynomial representations for finite and infinite lattice semirings (Theorems 7.2.9 and 7.3.7), as well as the generalisation to infinitary logic. Chapter 7 is based on the full version [GHNW22b] of the LICS paper.

The survey on model theory for semiring semantics (Section 7.1) includes a summary of joint work with Clotilde Bizière and Erich Grädel on locality properties [BGN23a, BGN23b]. Most of the results were obtained by Clotilde Bizière during an internship in our group, supervised by Erich Grädel. I was responsible for the presentation of the proof for Gaifman normal forms and for [BGN23a, Section 7] on the connection to classical semantics and the generalisation to lattice semirings. Also included is a summary of unpublished joint work with Anuj Dawar, Erich Grädel, and Lovro Mrkonjić on the compactness theorem in semiring semantics [DGMN24], to which I contributed the generalisation from finite to infinite min-max semirings (Section 7.1.2).

1.2 Outline

The first part of the thesis is devoted to the algebraic foundations and the fixed-point computation. Since these results may be of independent interest, no explicit connection to logic or semiring semantics is made in this part (other than as motivation). **Chapter 3** presents the necessary preliminaries on semirings, as well as the contributions on infinitary operations (Section 3.4) and on the variants of absorptive polynomials (Sections 3.6 and 3.7). The fixed-point computation is presented in **Chapter 4**, along with an additional symbolic computation tailored to the provenance semiring $\mathbb{S}^\infty(X)$.

The second part of the thesis is then concerned with semiring semantics for logics and games. **Chapter 5** provides an introduction to the semiring semantics for FO and LFP, including the connection to the fixed-point computation. New results are the semiring semantics for infinitary logic and the embedding of LFP into its finite-variable fragment (Section 5.3). The case study for Büchi games is presented in **Chapter 6**. For the sake of completeness, this chapter also includes a proof of the sum-of-strategies theorem for LFP (Section 6.2) that is used to derive the sum-of-strategies theorem for Büchi games. Extensions to more general infinite games are briefly discussed in Section 6.4. The final **Chapter 7** provides a survey on model theory for semiring semantics (Section 7.1) and the generalisations of the classical 0-1 laws to semiring semantics.

1.3 Acknowledgements

I would like to thank my supervisor Erich Grädel for his guidance and support during the five years that I worked on this thesis, for both giving me the freedom to work on my own ideas and also pushing me at the right times. It was through his engaging lectures, particularly *Mathematische Logik II* about set theory and fixed-point logics, that I became interested in logic. I am especially grateful that he supported my participation in two workshops at Schloss Dagstuhl and the Simons Institute in Berkeley. These opportunities allowed me to turn the names I had previously seen only on papers into personal interactions and discussions with many people working in this field. This is also where I met Javier Esparza, and I would like to thank him for agreeing to review this thesis.

Further, I am grateful to my colleagues at MGI, Richard, Benedikt, Matthias, Katrin, Lovro, and Sophie, and to everyone at i7, for creating a very enjoyable atmosphere in our group. I will miss all of the spontaneous conversations we had, about research, puzzles, and all other aspects of life. In particular, I thank Richard and Benedikt for welcoming me to the group as part of the former "Triforce office", with Patrick as an honorary office member, and also Lovro for managing much of our teaching infrastructure and for many chats about semirings. Research should not be done on an empty stomach, so I further thank Jan, Luca, Eva, Anton, and Silke (and everyone else who joined us) for shared lunch breaks, Chris and Tim for the best hummus, and everyone who tried (and sometimes succeeded) to bring delicious gluten-free cakes.

On a more serious note, I am grateful to Chris, Tim, Stefanie, and Nina for many meaningful conversations during the ups and downs of this journey, and to Chris also for his emotional support in a difficult time. Lastly, I would like to thank my family and friends, and everyone else who has supported and inspired me during these years.

Chapter 2

Preliminaries

This chapter introduces the notation that is used throughout this thesis. We assume that the reader is familiar with basic concepts of mathematical logic and set theory.

2.1 Notation

Sets, tuples, and indeterminates. The *power set* of a set A is denoted as $\mathcal{P}(A)$. Given a function $f\colon A \to B$ and a set $S \subseteq A$, we write $f(S)$ for the image of S under f. In particular, we use the notations $a \circ S = \{a \circ s \mid s \in S\}$ for binary operations, usually $\circ \in \{+, \cdot\}$. The notation $S \subset_{\text{fin}} A$ indicates that S is a (not necessarily proper) finite subset of A. Tuples are written in bold font, e.g. $\boldsymbol{a} = (a_1, \dots, a_n)$.

When working with polynomials, we usually use lowercase letters $x_1, x_2, \dots$ or $x, y, \dots$ as *indeterminates* and write X for a (usually finite) set of indeterminates, or $\boldsymbol{\mathcal{X}}$ for an infinite set. In some cases, we also use upper-case letters such as $X_1, X_2, \dots$ as indeterminates and then denote finite sets of indeterminates in bold font, such as $\boldsymbol{X}$. If we need a second set of indeterminates, we use the notation $\boldsymbol{A} = \{a, b, c, \dots\}$. As usual, *polynomials* in X with coefficients in K are denoted as $K[X]$, whereas $K[\![X]\!]$ denotes formal power series. We introduce more notation for polynomial semirings in Chapter 3.

Natural numbers and ordinals. We assume basic familiarity with ordinal numbers and transfinite induction (see, e.g., [Blu24, Chapter A3] for background). The main use of ordinal numbers in this thesis is as indices for transfinite sequences arising from fixed-point iterations. We denote the set of ordinals as On, the smallest infinite ordinal as ω, and the smallest uncountable ordinal as ω_1. Both ω and $\mathbb{N}$ denote the set of natural numbers, and both contain 0. We mainly use ω for indexing, usually writing $i < \omega$ to mean $i \in \omega$, and $\mathbb{N}$ for the (domain of the) semiring of natural numbers.

Graphs and trees. In this thesis, we only consider directed graphs. That is, by a *graph* we mean any structure (V, E) with nonempty node set V and edge relation $E \subseteq V \times V$. Paths can be finite or infinite, and may contain repetitions. A *tree* (with root ε) is a directed graph in which every node is reachable from ε by a unique path. The *depth of a node v* in a tree T (with root ε) is the length of the path from ε to v (in particular, ε is at depth 0), and the *height* of T is the maximal depth of any node (if this is finite). Given a node v in a tree T, the *subtree rooted at* v is the tree induced by the nodes reachable from v, with v as root. We occasionally write $v \sqsubseteq w$ to assert that w is reachable from v in a given tree, and $v \sqsubset w$ if additionally $v \neq w$. A tree is *finitely branching* if every node has only finitely many children. *Kőnig's lemma* states that every infinite tree that is finitely branching contains an infinite path (see, e.g., [Blu24, Chapter B2]).

Notation for semirings. A *monoid* is an algebraic structure $(M, \circ, 1)$ with a binary operation $\circ\colon M \times M \to M$ and a constant $1 \in M$, so that $\circ$ is associative and has 1 as (left- and right-)neutral element. Semirings are formally introduced in Chapter 3, here we only fix notation. We always denote an arbitrary semiring as $(K, +, \cdot, 0, 1)$ and use K to refer both to its domain and to the semiring itself (the notation with the letter K follows the introduction of K-relations in [GKT07]). We may add K as subscript of the operations, as in $+_K$ or 1_K, to avoid ambiguity.

2.2 Order Theory

Partial orders. As usual, a *partial order* $\leq$ on a set A is a binary relation that is reflexive, transitive, and antisymmetric. We use the symbols $\leq$, $\preceq$ and $\sqsubseteq$ for partial orders, and $<$, $\prec$, $\sqsubset$ for their strict versions. Two elements a, b of a partially ordered set A are *incomparable* if neither $a \leq b$ nor $a \geq b$ holds. If all elements are comparable, then the order is called *linear* or *total*.

We need the following notions of subsets of a partially ordered set $(A, \leq)$. An *antichain* is a subset $S \subseteq A$ of pairwise incomparable elements. A *chain* is a subset $C \subseteq A$ of pairwise comparable elements, i.e., that is totally ordered under $\leq$. An (ascending) ω-chain is a sequence $(a_i)_{i<\omega}$ with $a_i \in A$ and $a_i \leq a_j$ for $i < j$. Descending ω-chains are defined analogously, with $a_i \geq a_j$. An *upward-directed set* is a subset $D \subseteq A$ such that for all $a, b \in D$, there is a $c \in D$ with $a, b \leq c$. Downward-directed sets are defined analogously.

Lattices. The notation $a \sqcup b$ refers to the *supremum* (least upper bound) of $a, b \in A$ (if it exists), and $a \sqcap b$ denotes the *infimum* (greatest lower bound). The analogous notations for sets $S \subseteq A$ are $\bigsqcup S$ for the supremum and $\bigsqcap S$ for the infimum. We use the notations $\bigsqcup S$ and $\bigsqcup_{s \in S} s$ interchangeably.

A *lattice* is a partially ordered set A where $a \sqcup b$ and $a \sqcap b$ exist for all $a, b \in A$. A lattice is *distributive* if $a \sqcap (b \sqcup c) = (a \sqcap b) \sqcup (a \sqcap c)$ for all $a, b, c \in A$. It is a standard fact of lattice theory that this distributive law is equivalent to its dual: $a \sqcup (b \sqcap c) = (a \sqcup b) \sqcap (a \sqcup c)$. A lattice is *bounded* if least and greatest elements exist, which we denote by 0 and 1, respectively. We may also represent bounded lattices as algebraic structures $(A, \sqcup, \sqcap, 0, 1)$, leaving the underlying partial order implicit.

Complete orders. A partially-ordered set $(A, \leq)$ in which all chains $C \subseteq A$ have a supremum $\bigsqcup C \in A$ is called *chain-complete*. It is *fully chain-complete* if every chain also has an infimum $\bigsqcap C \in A$. Notice that our definition includes the empty chain, which implies that least and greatest elements exist. A *complete lattice* is a lattice in which every subset S has a (necessarily unique) supremum $\bigsqcup S$ and an infimum $\bigsqcap S$. In fact, the existence of only suprema (or only infima) suffices (indeed, we can define $\bigsqcap S = \bigsqcup\{a \mid a \le s \text{ for all } s \in S\}$). A complete lattice L is *completely distributive*, if it satisfies the infinitary distributivity law:

$$\bigsqcap_{i\in I} \bigsqcup_{a\in L_i} a = \bigsqcup_{f\in\mathcal{F}} \bigsqcap_{i\in I} f(i),$$

for every (possibly infinite) family $(L_i)_{i\in I}$ of subsets $L_i \subseteq L$, where $\mathcal{F}$ is the set of choice functions $f\colon I \to L$ with $f(i) \in L_i$ for all $i \in I$. We note that, as for finite distributivity, this is equivalent to its dual law. Moreover, any complete lattice induced by a total order is completely distributive (see, e.g., [Bir67, Chapter 5] for background).

Continuous functions. A function $f\colon A \to A$ on a partially-ordered set $(A, \le)$ is *monotone* if $a \le b$ implies $f(a) \le f(b)$ for all $a, b \in A$. If $(A, \le)$ is (fully) chain-complete, we further consider the following notions of continuity. We say that f is *ω-continuous*, if it commutes with suprema of ω-chains, i.e., $f(\bigsqcup_i a_i) = \bigsqcup_i f(a_i)$ for every ascending ω-chain $(a_i)_{i<\omega}$ in A. If it commutes with suprema of every nonempty chain $C \subseteq A$, i.e., $f(\bigsqcup C) = \bigsqcup f(C)$, then f is *chain-continuous*, and if the same also holds for infima of nonempty chains, then f is *fully chain-continuous*. If $(A, \le)$ is a complete lattice, we say that f commutes with *arbitrary suprema* if $f(\bigsqcup S) = \bigsqcup f(S)$ for every (possibly empty) set $S \subseteq A$, and f commutes with *arbitrary infima* if analogously $f(\bigsqcap S) = \bigsqcap f(S)$ for every set $S \subseteq A$.

Fixed points. Let $f\colon A \to A$ be a monotone operator on a partially-ordered set $(A, \le)$. An element $a \in A$ with $f(a) = a$ is a *fixed point* of f, and we write $\mathbf{lfp}(f)$ and $\mathbf{gfp}(f)$ for the *least* and *greatest* fixed points, if they exist. There are several conditions to ensure their existence. The most well-known one is the *Knaster-Tarski theorem*: if $(A, \le)$ is a complete lattice, then both $\mathbf{lfp}(f)$ and $\mathbf{gfp}(f)$ exist.

We can be a bit more general and only require that $(A, \le)$ is chain-complete and has a smallest element $0 \in A$. If f is ω-continuous, which means that it commutes with suprema of ω-chains, then *Kleene's fixed-point theorem* guarantees the existence of $\mathbf{lfp}(f)$. Indeed, writing $f^n(0)$ for the n-fold application of f, we see that $f(\bigsqcup_{n<\omega} f^n(0)) = \bigsqcup_{n<\omega} f^n(0)$ is a fixed point by ω-continuity of f, and it is the least one by monotonicity. The sequence

$$0,\ f(0),\ f(f(0)),\ \dots,\ f^n(0),\ \dots$$

is called *Kleene iteration* of f. Monotonicity of f implies, by induction, that this sequence forms a chain and that we have $f^n(0) \le a$ for any fixed-point a of f. This iteration can be extended to a transfinite sequence $(a_\alpha)_{\alpha\in\mathsf{On}}$, the *fixed-point iteration of f*, by setting $a_{\alpha+1} = f(a_\alpha)$ and $a_\gamma = \bigsqcup_{\alpha<\gamma} a_\alpha$ for limit ordinals γ. This again forms a well-defined chain by chain-completeness, and by monotonicity it is again bounded by $\mathbf{lfp}(f)$. We can thus conclude the existence of $\mathbf{lfp}(f)$ without requiring ω-continuity: since A is a set, the least fixed-point must be reached at some point, so $a_\alpha = \mathbf{lfp}(f)$ for some ordinal $\alpha \in \mathsf{On}$. The smallest such α is called the *closure ordinal* of the (least) fixed-point iteration for f. We refer to [DP02, Chapters 5 and 8] for further details and proofs of these fixed-point theorems.

This observation can be dualised to guarantee the existence of $\mathbf{gfp}(f)$. In particular, if $(A, \le)$ is fully chain-complete and has a greatest element 1, then the fixed-point iteration for $\mathbf{gfp}(f)$ is defined analogously by setting $a_0 = 1$, $a_{\alpha+1} = f(a_\alpha)$ as before, and $a_\gamma = \bigsqcap_{\alpha<\gamma} a_\alpha$ for limit ordinals. This sequence approaches $\mathbf{gfp}(f)$ from above, and since A is a set there is a closure ordinal α with $a_\alpha = \mathbf{gfp}(f)$.

2.3 Logic

We assume familiarity with first-order logic (see, e.g., [EF95] for background). In this thesis, we only consider signatures (or vocabularies) τ that are *finite* and *relational*. Relation symbols are usually denoted by the letters R and P, and $\mathrm{arity}(R)$ refers to the arity of the symbol R. We denote τ-structures by the letters $\mathfrak{A}$ and $\mathfrak{B}$, and the respective universes by A and B. We often refer to τ-structures as *classical* structures, and to the semantics of first-order logic as *classical* semantics, to clearly distinguish them from semiring interpretations and semiring semantics.

We mostly work with formulae in negation normal form and hence permit both $\vee$ and $\wedge$, as well as $\exists$ and $\forall$ in formulae (instead of defining them as abbreviations using negation). More precisely, we consider formulae built from atoms $R\boldsymbol{x}$ for $R \in \tau$, equalities $x = y$, inequalities $x \neq y$, the connectives $\neg$, $\vee$, $\wedge$, and the quantifiers $\exists$, $\forall$, with the usual semantics. A literal is an atom $R\boldsymbol{x}$ or its negation $\neg R\boldsymbol{x}$ (but not an (in)equality), and a formula is in negation normal form (nnf) if $\neg$ occurs only within literals. The notation $\equiv$ is used for (semantic) equivalence of formulae.

We write $\mathrm{FO}(\tau)$ for the set of first-order formulae of signature τ, which we denote by the letters $\varphi, \psi, \vartheta, \dots$. Free variables and sentences are defined as usual. The notation $\varphi(\boldsymbol{x})$ indicates that all free variables of φ are contained in the tuple $\boldsymbol{x}$ (but not every $x \in \boldsymbol{x}$ has to occur in φ). If A is the universe of a τ-structure (or of a semiring interpretation), $\varphi(\boldsymbol{x})$ a formula, and $\boldsymbol{a} \subseteq A$ a tuple of matching length, we write $\varphi(\boldsymbol{a})$ for the *instantiated formula* that results from replacing each variable $x \in \boldsymbol{x}$ by the corresponding element of $\boldsymbol{a}$.

We also consider extensions of first-order logic by fixed points, and by infinitary conjunctions and disjunctions. These are introduced in Section 5.2 and Section 5.3, respectively, and we refer to [GKL+07, EF95] for further background.

Chapter 3

Semirings

A central theme of this thesis is the combination of algebraic techniques and methods from logic and model theory. We begin with the algebraic side, focussing on semirings of polynomials, homomorphisms, and order theory. From the outset, these three aspects have been the cornerstone of the semiring framework for provenance analysis: semirings of polynomials provide most general provenance information; this information can then be specialised to application semirings via homomorphisms in order to model, e.g., costs or access privileges. Order theory comes into play when fixed-point logics such as datalog or LFP are considered, but is also relevant for the existence of well-defined infinitary addition and multiplication operations, needed for instance for the evaluation of infinitary logic.

A particular emphasis is on *absorptive* semirings, which turn out to be the appropriate semirings for provenance analysis of fixed-point logics that include both least and greatest fixed points. We show that absorptive semirings admit a natural definition of infinitary operations (under suitable continuity assumptions) that satisfies most of the properties we would expect, such as infinitary versions of associativity and commutativity. The only exception is infinitary distributivity, which holds in some, but not all absorptive semirings. Of central importance for the results in the forthcoming chapters is the semiring of (generalised) absorptive polynomials $\mathbb{S}^\infty(X)$, which is essentially a quotient of the polynomial semiring $\mathbb{N}[X]$ (extended by the exponent ∞) with respect to absorption. This semiring is on one hand quite well-behaved: absorptive polynomials are finite, all ascending chains are finite, and it satisfies the infinitary distributivity law. On the other hand, it is also quite powerful: it is the most general absorptive semiring that additionally satisfies desirable continuity properties. $\mathbb{S}^\infty(X)$ thus becomes a fruitful tool in our proofs, so we present a detailed analysis of $\mathbb{S}^\infty(X)$ and its extensions by coefficients and by infinitely many indeterminates.

3.1 Semiring Fundamentals

3.1.1 Definitions

Semirings are algebraic structures with two operations, usually denoted as addition $+$ and multiplication $\cdot$, with the key property that $\cdot$ distributes over $+$. Typical examples are the semiring of natural numbers $(\mathbb{N}, +, \cdot, 0, 1)$ and the absorptive Viterbi semiring $([0, 1], \max, \cdot, 0, 1)$ over the real interval $[0, 1]$. We summarise the relevant elementary properties of semirings.

Definition 3.1.1. A *commutative semiring* is an algebraic structure $(K, +, \cdot, 0, 1)$ with $0 \neq 1$ such that $(K, +, 0)$ and $(K, \cdot, 1)$ are commutative monoids, and

- multiplication distributes over addition: $a \cdot (b + c) = ab + ac$ for all $a, b, c \in K$,
- 0 is annihilating: $0 \cdot a = 0$ for all $a \in K$.

We only consider commutative semirings in this thesis, and henceforth just write *semiring* to refer to a commutative semiring. A *homomorphism* between two semirings K and K' is a function $h\colon K \to K'$ that preserves neutral elements and commutes with the semiring operations, i.e., $h(0_K) = 0_{K'}$, $h(1_K) = 1_{K'}$, and $h(a \circ_K b) = h(a) \circ_{K'} h(b)$ for all $a, b \in K$ and $\circ \in \{+, \cdot\}$. We associate with each semiring the order that is induced by addition. This means that addition leads to larger values, i.e., $a \leq a + b$ for all elements a, b, so that 0 is the least element. Moreover, this order is compatible with the semiring operations (for multiplication, this follows from distributivity) and with homomorphisms.

Definition 3.1.2 (Natural order)**.** With each semiring $(K, +, \cdot, 0, 1)$, we associate the binary relation $\leq$ on K, with $a \leq b$ if there is $c \in K$ with $a + c = b$. This relation is always reflexive and transitive. If it is also antisymmetric, and thus a partial order, we say that K is *naturally ordered* and refer to $\leq$ as the *natural order* of K.

Proposition 3.1.3. *In a naturally ordered semiring $(K, +, \cdot, 0, 1)$, both operations $+$ and $\cdot$ are monotone (in each argument). Moreover, if K and K' are naturally ordered, then for any homomorphism $h\colon K \to K'$, we have that $a \leq_K b$ implies $h(a) \leq_{K'} h(b)$.*

Since an order is required to speak about *least* or *greatest* fixed points, we always work with naturally-ordered semirings. This excludes rings such as $(\mathbb{R}, +, \cdot, 0, 1)$, due to their additive inverses. Unless stated otherwise, $\leq$ always refers to the natural order.

In some cases, it is desirable that the semiring operations preserve positive values. Naturally ordered semirings are always $+$-*positive*, i.e., $a + b = 0$ implies $a = 0$ and $b = 0$, but may have zero divisors.

Definition 3.1.4. A semiring $(K, +, \cdot, 0, 1)$ is *positive* if for all $a, b \in K$, we have that $a + b = 0$ implies $a = 0$ and $b = 0$, and $a \cdot b = 0$ implies $a = 0$ or $b = 0$.

Most of the results in this thesis apply to certain smaller classes of semirings, based on the following algebraic properties.

Definition 3.1.5. A semiring $(K, +, \cdot, 0, 1)$ is called

- *idempotent*, if $a + a = a$ for all $a \in K$;
- *multiplicatively idempotent*, if $a \cdot a = a$ for all $a \in K$,
- *fully idempotent*, if both of the above are satisfied;
- *absorptive*, if $a + ab = a$ for all $a, b \in K$.

3.1.2 A Hierarchy of Semirings

We next establish the relationship between these classes of semirings, resulting in the hierarchy shown in Fig. 3.1. Starting from the class of *naturally-ordered* semirings, the first subclass we consider are the *idempotent* semirings.

Proposition 3.1.6. *Every idempotent semiring is naturally ordered by* $a \leq b \Leftrightarrow a + b = b$, *and moreover* $a + b = a \sqcup b$, *for all elements* a, b.

Proof. Let K be an idempotent semiring and assume $a \leq b$, so there is $c \in K$ with $a + c \leq b$. Then $a + b = a + a + c = a + c = b$ by idempotence. If $b \leq a$, then analogously $a + b = a$, hence $\leq$ is antisymmetric and K naturally ordered. For the supremum, it is clear that $a + b$ is an upper bound. For minimality, assume that $a, b \leq c$. Then $a + c = c$ and $b + c = c$, and then by idempotence $(a + b) + c = c$, so $a + b$ is indeed the least upper bound. □

The natural order of an idempotent semiring can thus be defined as $a \leq b \Leftrightarrow a + b = b$. Going further down in the hierarchy, we consider two incomparable subclasses of idempotent semirings: the *fully-idempotent* and the *absorptive* semirings. Indeed, fully-idempotent semirings are idempotent by definition, and absorption $a + ab = a$ implies idempotence by choosing $b = 1$.

Absorptive semirings have also been called 1-*bounded* or 0-*closed* [Moh02], where the latter refers to the fact that $\sum_{i=0}^{n} a^i = a^0 = 1$ for any element a and $n < \omega$. These are the most important semirings in this thesis. The main reason is that absorption establishes a stronger connection (in addition to distributivity) between multiplication and the natural order, which implies that multiplication behaves, to some extent, dually to addition:

Proposition 3.1.7. *In a naturally-ordered semiring* $(K, +, \cdot, 0, 1)$, *the following are equivalent:*

(1) *absorption, i.e.,* $a + ab = a$ *for all* $a, b \in K$,

(2) 1 *is the greatest element, i.e.,* $a \leq 1$ *for all* $a \in K$,

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

Proof. Absorption implies $1 + b = 1$, and hence $b \leq 1$ for all $b \in K$. By monotonicity, this further implies $ab \leq a$ for all $a, b \in K$. And this implies $a + ab = a$ by Proposition 3.1.6. □

The class of *lattice semirings* consists of semirings that are both absorptive and multiplicatively idempotent. The name is due to the following observation.

Proposition 3.1.8. *The multiplicatively-idempotent absorptive semirings are exactly the bounded distributive lattices* $(K, \sqcup, \sqcap, 0, 1)$.

Proof. Let K be multiplicatively-idempotent and absorptive. We have $a + b = a \sqcup b$ by idempotence and $ab \leq a, b$ by absorption. To see that $ab = a \sqcap b$, assume that $c \leq a, b$. Then $c + ab = (c + a)(c + b) = ab$ by absorption and multiplicative idempotence, so $c \leq ab$ and ab is indeed the greatest lower bound. It follows that K is a distributive lattice, with least

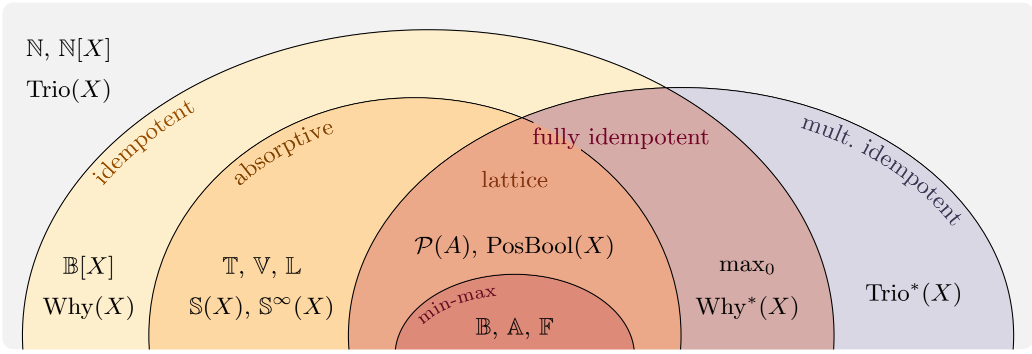

Figure 3.1: Venn diagram of the relevant classes of semirings.

element 0 and, by absorption, greatest element 1. Conversely, every lattice is absorptive and has idempotent operations. □

If the natural order of a lattice semiring is total, we write the semiring as $(K, \max, \min, 0, 1)$ and refer to it as *min-max semiring*. This is the smallest class of semirings we consider, and can be seen as the natural generalisation of the Boolean semiring to more than two values.

3.1.3 A List of Semirings

We present a non-exhaustive list of semirings relevant for our results or for provenance applications. Following the terminology of semiring provenance analysis, we split these into *application semirings*, which can be used to model properties such as costs, distances, multiplicities, or access rights, and into *provenance semirings* which are, for the most part, freely generated semirings of polynomials that can be used to represent general forms of provenance information (cf. [GKT07]). For most of the semirings listed here, their place in the hierarchy can be seen in 3.1.

Application Semirings

- The *Boolean* semiring $\mathbb{B} = (\{0, 1\}, \vee, \wedge, 0, 1)$ is the simplest semiring and the habitat of classical truth values. One of the guiding principles for defining semiring semantics is that semantics in $\mathbb{B}$ coincides with standard semantics.

- Every linear order $(A, \leq)$ with bounds 0 and 1 induces the semiring $(A, \max, \min, 0, 1)$. The *access control* semiring $\mathbb{A}$ (cf. [FGT08]) is the finite min-max semiring induced by the linear order $0 < \mathsf{T} < \mathsf{S} < \mathsf{C} < \mathsf{P}$. Its elements represent the access restrictions "inaccessible" (or false), "top secret", "secret", "confidential", and "public". An infinite min-max semiring is the *fuzzy* semiring $\mathbb{F} = ([0, 1], \max, \min, 0, 1)$ over real numbers.

- An example of a lattice semiring that is not a min-max semiring is the powerset semiring $(\mathcal{P}(A), \cup, \cap, \emptyset, A)$, for any nonempty set A. Examples of fully-idempotent semirings that are not absorptive are $(\{0, 1, \dots, n\}, \max, \max_0, 0, 1)$ for $n \geq 2$, where $\max_0(a, b) = 0$ if $a = 0$ or $b = 0$, and $\max_0(a, b) = \max(a, b)$ otherwise.

- The *natural* semiring $\mathbb{N} = (\mathbb{N}, +, \cdot, 0, 1)$ has applications in bag semantics for databases and for counting evaluation strategies. We also consider its (fully continuous, see below) completion $\mathbb{N}^\infty = (\mathbb{N} \cup \{\infty\}, +, \cdot, 0, 1)$ by an element representing infinity with the usual semantics ($n \cdot \infty = n \cdot \infty = \infty$ for $n > 0$).

- The *tropical* semiring $\mathbb{T} = (\mathbb{R}^\infty_{\geq 0}, \min, +, \infty, 0)$, over the nonnegative real numbers extended by ∞, is used for min-cost computations, such as distances in graphs. Notice that its natural order is the inverse of the standard order on real numbers, since semiring addition is given by min (instead of max).

  It is isomorphic to the *Viterbi* semiring $\mathbb{V} = ([0,1], \max, \cdot, 0, 1)$, whose values can be interpreted as confidence scores, by the isomorphism $x \mapsto e^{-x}$.

- The *Łukasiewicz* semiring $\mathbb{L} = ([0,1], \max, \star, 0, 1)$ with $a \star b = \max(a+b-1, 0)$ appears in many-valued logics.

- Examples of semirings we do not consider in this thesis are the ring $(\mathbb{R}, +, \cdot, 0, 1)$, which is not naturally ordered, and the non-commutative language semiring $(\mathcal{P}(\Sigma^*), \cup, \circ, \emptyset, \{\varepsilon\})$ consisting of all languages of finite words over the alphabet Σ, where ε is the empty word and $L \circ M = \{vw \mid v \in L, w \in M\}$ is element-wise concatenation.

We can obtain further semirings by product constructions. Given two semirings K_1, K_2, their *direct product* is defined as usual, with domain $K_1 \times K_2$ and pointwise operations (e.g., $(a_1, a_2) + (b_1, b_2) = (a_1 + a_2, b_1 + b_2)$). The induced natural order is the usual product order. A particular case of a (possibly infinite) direct product is the *function semiring* $(K^A, +, \cdot, \mathbf{0}, \mathbf{1})$ for a given semiring K and a nonempty set A, consisting of all functions $f : A \to K$. The operations and the natural order are defined pointwise (e.g., $(f+g)(a) = f(a) + g(a)$) and $\mathbf{0}$, $\mathbf{1}$ are constant functions. This construction is used to define semiring semantics of LFP in Chapter 5. Properties such as idempotence and absorption (and also full continuity, as defined below, but not positivity) are inherited for both constructions.

Notation for Polynomials

We next consider semirings of multivariate polynomials which are the most general semirings of their respective classes (in the sense of free objects). Let X be a finite set of indeterminates. We represent monomials, such as x^2y^3, as mappings $m : X \to E$ that map each indeterminate to an exponent in the semiring (or monoid) E. For standard polynomials, we set $E = \mathbb{N}$, but we also consider $E = \mathbb{B}$ and $E = \mathbb{N}^\infty$. We denote the set of monomials as $\mathsf{Mon}_X(E)$, omitting X and E if they are clear from the context. Monomials are written as finite products $x_1^{m(x_1)} \cdots x_n^{m(x_n)}$, for a monomial m in indeterminates $X = \{x_1, \dots, x_n\}$, and multiplication is defined as usual by adding exponents.

Several provenance semirings have been introduced, often named after the kind of provenance information they represent (see [GT17b] for a brief survey, but note the discussion of an inconsistency for the semirings $\mathrm{Why}(X)$ and $\mathrm{Trio}(X)$ below). In an attempt to unify the notation, we define the semirings $\mathrm{Poly}(K, E, X)$ and $\mathrm{Poly}^\omega(K, E, X)$ of formal power series and polynomials with coefficients in K, exponents in E, and indeterminates in X.

Definition 3.1.9. Let K be a semiring, $E \in \{\mathbb{B}, \mathbb{N}, \mathbb{N}^\infty\}$, and let X be a finite set. The semiring $\mathrm{Poly}(K, E, X)$ consists of all functions $P : \mathsf{Mon}_X(E) \to K$ with pointwise addition

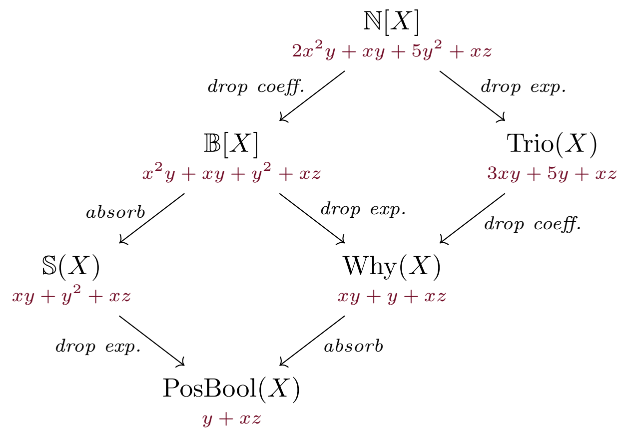

Figure 3.2: Overview on relevant provenance semirings.

and the usual polynomial multiplication, that is,

$$(P+Q)(m) = P(m) + Q(m),$$
$$(P \cdot Q)(m) = \sum_{m_1 \cdot m_2 = m} P(m_1) \cdot Q(m_2).$$

The neutral elements are the constant mapping $0 : m \mapsto 0_K$ and the mapping 1 with $1(m) = 1_K$ for $m = 1$, and $1(m) = 0_K$ otherwise. As usual, we denote elements $P \in \mathrm{Poly}(K, E, X)$ as formal sums $c_1 m_1 + c_2 m_2 + \dots$, where $c_i = P(m_i)$ for each i.

The semirings $\mathrm{Poly}^\omega(K, E, X)$ are defined in the same way, but only contain functions $P\colon \mathsf{Mon}_X(E) \to K$ with finite support.

Notice that the sum in the definition of $P \cdot Q$ can range over infinitely many monomials m_1, m_2 if $E = \mathbb{N}^\infty$. We only consider coefficient semirings where infinite sums are well-defined in this case (see Section 3.4), so this will not be an issue. We later introduce further semirings $\mathrm{AbsPoly}(K, E, X)$ of absorptive polynomials in Section 3.7.

Provenance Semirings

- The most general provenance semiring is the semiring $\mathbb{N}[X]$ of standard polynomials, in our notation $\mathrm{Poly}^\omega(\mathbb{N}, \mathbb{N}, X)$. This is the commutative semiring freely generated by X, in the sense of the following universal property: any assignment $f\colon X \to K$ into a commutative semiring K extends uniquely to a semiring homomorphism $h_f\colon \mathbb{N}[X] \to K$ (by polynomial evaluation).

- $\mathbb{B}[X]$, or $\mathrm{Poly}^\omega(\mathbb{B}, \mathbb{N}, X)$, is the freely generated idempotent semiring. It results from $\mathbb{N}[X]$ by dropping coefficients.

- The semiring $\mathrm{Trio}(X)$ was defined in [Gre11] to capture a specific form of provenance computation. It results from $\mathbb{N}[X]$ by dropping exponents. In our notation, this is $\mathrm{Poly}^\omega(\mathbb{N}, \mathbb{B}, X)$ or $\mathrm{Poly}(\mathbb{N}, \mathbb{B}, X)$ (there is no difference, as the set of monomials $\mathsf{Mon}_{\mathbb{B}}(X)$ is finite).

- The semiring $\mathrm{Why}(X)$ was first introduced in [BCTV08] to capture the notion of *why-provenance*. It was originally defined as $\mathrm{Why}(X) = (\mathcal{P}(\mathcal{P}(X)), \cup, \uplus, \emptyset, \{\emptyset\})$ with $S \uplus T = \{s \cup t \mid s \in S, t \in T\}$. It is isomorphic to $\mathrm{Poly}(\mathbb{B}, \mathbb{B}, X)$ in our notation, and results from $\mathbb{N}[X]$ by dropping both coefficients and exponents.

Remark 3.1.10. It was later claimed that $\mathrm{Trio}(X)$ and $\mathrm{Why}(X)$ result from $\mathbb{N}[X]$ and $\mathbb{B}[X]$, respectively, by making multiplication idempotent [GT17b, Figure 2]. However, this is not the case for the original definition. While dropping exponents indeed ensures that $x^2 = x$ for indeterminates $x \in X$, the same is not true for arbitrary polynomials: consider $(x+y)^2 = x^2 + xy + y^2$ in $\mathrm{Why}(X)$. We still have the property that any assignment into a semiring that is multiplicatively idempotent (for $\mathrm{Trio}(X)$) or fully idempotent (for $\mathrm{Why}(X)$) extends to a homomorphism by polynomial evaluation, but $\mathrm{Trio}(X)$ and $\mathrm{Why}(X)$ do not belong to the respective class themselves.

We note that free multiplicatively- and fully-idempotent semirings must exist (since these classes form a nontrivial variety in the sense of universal algebra [BS81]). We denote these semirings as $\mathrm{Trio}^*(X)$ and $\mathrm{Why}^*(X)$ in Fig. 3.1, but an explicit construction is likely more involved than for $\mathrm{Trio}(X)$ and $\mathrm{Why}(X)$, and we do not consider it here.

- Absorptive polynomials $\mathbb{S}(X)$ have been introduced for datalog provenance in [DMRT14]. This is the freely generated absorptive semiring, and it can be defined as a quotient of $\mathbb{N}[X]$ or $\mathbb{B}[X]$ by absorption. We are mostly interested in the generalised version $\mathbb{S}^\infty(X)$ that includes the exponent ∞, which is discussed in detail in Section 3.5. In the notation of Section 3.7, these are the semirings $\mathrm{AbsPoly}(\mathbb{B}, \mathbb{N}, X)$ and $\mathrm{AbsPoly}(\mathbb{B}, \mathbb{N}^\infty, X)$.
- Omitting the exponents of $\mathbb{S}(X)$ or $\mathbb{S}^\infty(X)$ results in the semiring $\mathrm{PosBool}(X)$, which is the distributive lattice (or lattice semiring) freely generated by X. It has originally been described as the set of positive Boolean formulae up to logical equivalence [GKT07]. Its operations are then denoted as $(\mathrm{PosBool}(X), \vee, \wedge, \mathrm{true}, \mathrm{false})$, but we stick to the polynomial notation $(\mathrm{PosBool}(X), +, \cdot, 0, 1)$ for consistency.

The construction of provenance semirings is illustrated in Fig. 3.2. We note that for simplicity, we refer to the elements of any of the above provenance semirings as *polynomials*, although strictly speaking only $\mathbb{N}[X]$ and $\mathbb{B}[X]$ are semirings of polynomials. This is also the reason why we use round brackets to denote the other semirings, such as $\mathrm{Why}(X)$.

We will revisit and extend this list of provenance semirings at the end of Section 3.2 to discuss completeness and continuity properties. We also point out that the appropriate provenance semirings for logics with negation are dual-indeterminate polynomials, such as $\mathbb{N}[X, \overline{X}]$, which are defined as quotients of the standard polynomial semirings, in this case $\mathbb{N}[X]$. These are discussed in Section 5.1.3 when semirings semantics for FO is introduced.

3.2 Fixed Points

Towards a well-defined semiring semantics for fixed-point logic, we have to ensure that fixed points (of monotone operators) always exist in the semirings we consider. These fixed points are always understood with respect to the natural order, and we thus impose additional completeness constraints on this order. We follow the exposition in [DGNT21, Naa19], where these concepts were first introduced.

We can think of this setting as the combination of two objects: a semiring with addition and multiplication operations, and a complete lattice (or chain-complete partial order). These are of course related, as the partial order is the natural order induced by addition, but it may be helpful to think of them as separate entities. In particular, the infimum operation of the partial order is, a priori, completely unrelated to the semiring multiplication. This is also the reason why we additionally assume continuity, to impose some relationship between the semiring and the lattice operations.

3.2.1 Continuous Semirings

The most convenient way to ensure the existence of fixed points in a semiring is to require that the natural order is a complete lattice. We are a bit more general and only require the natural order to be fully *chain*-complete (i.e., suprema and infima of totally-ordered subsets exist). However, we will see that for most of the semirings we are interested in, these two alternatives are in fact equivalent. We additionally assume full continuity, so that the semiring operations are compatible with suprema and infima. This is crucial for the proofs in the following chapters, but does not seem to be a strong restriction in practice: all natural examples[1] of fully chain-complete semirings known to me are in fact also fully continuous.

Definition 3.2.1. A naturally ordered semiring K is *fully chain-complete* if every chain $C \subseteq K$ has a supremum $\bigsqcup C$ and an infimum $\bigsqcap C$ in K. It is further *fully continuous* if its operations are fully chain-continuous in both arguments. That is,

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

for all $a \in K$, all nonempty chains $\emptyset \neq C \subseteq K$ and $\circ \in \{+, \cdot\}$.

Notice that fully chain-complete semirings are in particular bounded with least element $\bot = 0 = \bigsqcup \emptyset$ and greatest element $\top = \bigsqcap \emptyset$. By considering the transfinite fixed-point iteration (starting at either $\bot$ or $\top$), it follows that every monotone function has a least and a greatest fixed point.

Proposition 3.2.2. *For a monotone operator $f\colon K \to K$ on a fully chain-complete semiring, both* $\mathbf{lfp}(f)$ *and* $\mathbf{gfp}(f)$ *exist.*

We usually consider fully-continuous homomorphisms between fully-continuous semirings, which commute with suprema and infima of nonempty chains. Notice that we always have $h(0) = 0$, but not necessarily $h(\top) = \top$ (unless the semiring is absorptive).

If positive semiring values should be preserved by the semiring operations and by suprema and infima, we need the following property in addition to positivity.

Definition 3.2.3. A fully chain-complete semiring K is *chain-positive* if for each nonempty chain $C \subseteq K$ of nonzero elements, the infimum $\bigsqcap C$ is nonzero as well.

[1] A notable exception is the semiring $(\mathcal{P}(\mathbb{N} \times \mathbb{N}), \cup, \circ, \emptyset, \mathbb{N} \times \mathbb{N})$ of binary relations, where $\circ$ is composition. For $R_n = \{(1,k) \mid n \leq k < \omega\}$ and $P = \{(k,1) \mid k < \omega\}$, we have $\bigsqcap_n (R_n \circ P) = \{(1,1)\}$, but $(\bigsqcap_n R_n) \circ P = \emptyset \circ P = \emptyset$. However, this semiring is not commutative.

From chains to arbitrary sets. Full chain-completeness is more general than the common notion of complete lattices, as we only require suprema and infima of *chains* instead of arbitrary sets. However, based on results in [Mar76] it follows that the two notions are equivalent in many cases, in particular for the important class of absorptive semirings.

The idea is as follows: for a set $A = \{a_1, a_2, \dots\}$, we consider the chain of partial suprema a_1, $a_1 \sqcup a_2$, $a_1 \sqcup a_2 \sqcup a_3$, …, assuming they exist. This is always the case in idempotent semirings, where finite suprema coincide with summation. If A is infinite, this chain can be extended to an infinite chain by transfinite induction and chain-completeness, and it is easy to see that the supremum of the resulting infinite chain equals the supremum of A. This construction is also useful when we consider fully-continuous functions, such as the semiring operations or homomorphisms, as they preserve the supremum of the chain and thus also the supremum of A.

Lemma 3.2.4 ([Mar76])**.** *Let P, Q be fully chain-complete partial orders and let $f\colon P \to Q$ be fully chain-continuous.*

(1) *Then f commutes with infima of downward-directed and suprema of upward-directed sets (and these infima and suprema exist).*

(2) *Assume that suprema of finite subsets exist in P. Then P is a complete lattice, and if f commutes with suprema of finite subsets, then also with suprema of arbitrary sets.*

Corollary 3.2.5. *If K is an idempotent, fully chain-complete semiring, then its natural order is a complete lattice. If K is additionally fully continuous, then the operations $+$ and $\cdot$ commute with suprema of arbitrary sets, and with infima of downward-directed sets.*

Corollary 3.2.6. *Let K, K' be idempotent, fully-continuous semirings. Every fully continuous homomorphism $h\colon K \to K'$ commutes with suprema of arbitrary sets, and with infima of downward-directed sets.*

Countable chains. An alternative to the notion fully chain-complete semirings would be to require completeness and continuity only with respect to (ascending and descending) ω-chains. By Kleene's fixed-point theorem, this would guarantee the existence of fixed points for ω-continuous operators. If the semiring operations are ω-continuous, then operators induced by polynomial expressions in these operations are indeed ω-continuous. However, this is not necessarily the case when we consider nested fixed points (see Section 5.2.4). In fact, the situation is more intricate: although the operators induced by LFP-formulae (potentially with nested fixed points) are not ω-continuous in general, one can show that they always have closure ordinal at most ω (also Section 5.2.4). So in retrospect, the existence of ω-chains would suffice for a well-defined semantics of LFP, but to develop the general theory that leads to this observation, it is more convenient to assume completeness and continuity for chains of arbitrary cardinality. This is also needed to obtain the equivalence to complete lattices.

In the case of countable semirings, such as absorptive polynomials $\mathbb{S}(X)$ and $\mathbb{S}^\infty(X)$, it is easy to see that considering ω-chains is sufficient.

Lemma 3.2.7. *Let K, K' be fully chain-complete semirings, and $C \subseteq K$ a countable chain. Then there is a descending ω-chain $(x_i)_{i<\omega}$ such that $\bigsqcap C = \bigsqcap_i x_i$ and $\bigsqcap f(C) = \bigsqcap_i f(x_i)$ for every monotone function $f\colon K \to K'$. Analogously for suprema.*

Proof sketch. This is trivial for finite chains, so assume that C is infinite. Fix a bijection $g\colon \omega \to C$ and recursively define $x_0 \coloneqq g(0)$ and $x_{i+1} \coloneqq \min(g(i+1), x_i)$. This defines an ω-chain with $x_i \in C$ 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, $f(c) \geq f(x_i)$ and thus $\bigsqcap f(C) \geq \bigsqcap_i f(x_i)$. □

As a consequence, we frequently use the following observation about ω-chains in these semirings.

Lemma 3.2.8 (Chain splitting)**.** *Let K be a fully continuous semiring and let $(a_i)_{i<\omega}$ and $(b_i)_{i<\omega}$ be descending ω-chains in K. Then for $\circ \in \{+, \cdot\}$,*

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

Analogously for suprema of ascending ω-chains.

Proof sketch. $\bigsqcap_i a_i \circ \bigsqcap_j b_j = \bigsqcap_{i,j}(a_i \circ b_j)$ by full continuity. Clearly, $\bigsqcap_{i,j}(a_i \circ b_j) \leq \bigsqcap_k (a_k \circ b_k)$. For the other direction, observe that $a_i \circ b_j \geq a_k \circ b_k$ for $k = \max(i,j)$. □

Related concepts. There are several related concepts of semirings with completeness and continuity properties. The inspiration for our definition are the *ω-continuous semirings* that have, for instance, been used for provenance analysis of datalog [GKT07]. These require suprema of ascending ω-chains, which must be compatible with the semiring operations (i.e., the operations must be ω-continuous). It follows that least fixed points of operators built from the semiring operations exist by Kleene's fixed-point theorem. More similar to our definition are the *continuous semirings* in [DK09], which require suprema of upward-directed sets that are compatible with the semiring operations. This is equivalent to the analogous requirement for suprema of chains (cf. [Mar76]), but infima are not considered.

Also related are the *c-semirings* introduced in the context of constraint satisfaction problems in [BMR97]. These are equipped with a summation operation for (possibly infinite) sets that is compatible with multiplication and satisfies additional constraints, which imply that c-semirings are absorptive semirings in which arbitrary suprema exist and are compatible with the semiring operations. The last concept we mention here are the *topological dioids* considered [GM08]. These are dioids, i.e., naturally-ordered semirings, in which both operations are compatible with suprema of countable chains.

The common difference to our definition of fully-continuous semirings is that we also require completeness and continuity for infima, as we are interested not only in least, but also in greatest fixed points.

3.2.2 Provenance Semirings Revisited

All of the application semirings we considered in Section 3.1.3 are fully continuous, with the exception of $\mathbb{N}$, which has to be completed to $\mathbb{N}^\infty$. Finite semirings are trivially chain-positive, but out of the infinite application semirings, only $\mathbb{N}^\infty$ is chain-positive.

The situation becomes more interesting for the provenance semirings. We are now interested in freely generated (or universal) fully-continuous semirings, which means that assignments of the indeterminates must induce a unique homomorphism that is fully-continuous.

Example 3.2.9. The semiring $\mathbb{N}[X]$ is not chain-complete, both because of the coefficients $\mathbb{N}$ and the finite number of monomials. We instead consider the fully-continuous semiring $\mathbb{N}^\infty[\![X]\!]$ of formal power series, in our notation $\mathrm{Poly}(\mathbb{N}^\infty, \mathbb{N}, X)$.

However, $\mathbb{N}^\infty[\![X]\!]$ is not the free fully-continuous semiring. To see this, let $P_n = \sum_{n<i<\omega} x^i$ for $n < \omega$, and consider the instantiation $h(x) = 1$ into the fully-continuous semiring $\mathbb{N}^\infty$. This lifts to a unique semiring homomorphism $h\colon \mathbb{N}^\infty[\![x]\!] \to \mathbb{N}^\infty$, but this homomorphism is not fully-continuous: $\bigsqcap_n h(P_n) = \bigsqcap_n \infty = \infty$, whereas $h(\bigsqcap P_n) = h(0) = 0$.

To see that $\mathbb{N}^\infty[\![X]\!]$ is fully continuous, we make the following general observation.

Proposition 3.2.10. *Let K be a countable, fully-continuous semiring, let $E \in \{\mathbb{B}, \mathbb{N}\}$, and let X be a finite set. Then* $\mathrm{Poly}(K, E, X)$ *is fully continuous.*

Proof. Recall that $+$ is defined pointwise, by adding coefficients of each monomial. It follows that $\bigsqcup$ and $\bigsqcap$ also act pointwise. Full continuity of $+$ thus immediately follows from full continuity of K. Multiplication is more involved and needs Lemma 3.2.7 on countable chains and Lemma 3.2.8 (Chain splitting). Let $C \subseteq \mathrm{Poly}(K, E, X)$ be a chain. For each monomial m, we have:

$$\begin{aligned}\Big(Q \cdot \bigsqcap_{P \in C} P\Big)(m) &= \sum_{m_1 m_2 = m} \Big(Q(m_1) \cdot \bigsqcap_{P \in C} P(m_2)\Big) \\ &\overset{(1)}{=} \sum_{m_1 m_2 = m} \bigsqcap_{P \in C} (Q(m_1) \cdot P(m_2)) \\ &\overset{(2)}{=} \bigsqcap_{P \in C} \sum_{m_1 m_2 = m} (Q(m_1) \cdot P(m_2)) \overset{(3)}{=} \Big(\bigsqcap_{P \in C} Q \cdot P\Big)(m).\end{aligned}$$

Step (1) holds by full continuity of K. For (2), notice that for each $m \in \mathsf{Mon}_E(X)$, there are only finitely many m_1, m_2 with $m = m_1 \cdot m_2$ (this would no longer be the case for $E = \mathbb{N}^\infty$). Moreover, the values $Q(m_1) \cdot P(m_2)$ form a chain (for $P \in C$) in K, and since K is countable this chain can be reduced to an ω-chain by Lemma 3.2.7. Step (2) then follows by Lemma 3.2.8. The proof for suprema of chains is completely analogous. □

We recall the list of provenance semirings and discuss their continuity properties.

- The fully-continuous completion of $\mathbb{N}[X]$ is the semiring $\mathbb{N}^\infty[\![X]\!]$, or $\mathrm{Poly}(\mathbb{N}^\infty, \mathbb{N}, X)$, of formal power series. Example 3.2.9 shows that it is neither chain-positive nor the free fully-continuous semiring.

  An alternative is the semiring $\mathrm{Poly}(\mathbb{N}^\infty, \mathbb{N}^\infty, X)$, which we denote as $\mathbb{N}^\infty_\infty[\![X]\!]$. This semiring is not fully continuous, but it is universal for semirings with infinitary operations (see Section 3.4.3).

- Analogously, $\mathbb{B}[\![X]\!]$, or Poly$(\mathbb{B}, \mathbb{N}, X)$, is the fully-continuous completion of $\mathbb{B}[X]$. The counterexample 3.2.9 for $\mathbb{N}^\infty[\![X]\!]$ also applies to $\mathbb{B}[\![X]\!]$.
- The fully-continuous completion of Trio(X) is the semiring Poly$(\mathbb{N}^\infty, \mathbb{B}, X)$. Since $\mathbb{N}^\infty$ is chain-positive and $\mathsf{Mon}_{\mathbb{B}}(X)$ finite, this semiring is chain-positive. The semiring Why(X) is finite and hence trivially fully continuous.
- The absorptive semiring $\mathbb{S}(X)$ is already fully continuous (see Section 3.7 for a proof). However, it is not chain-positive and hence not the universal absorptive, fully-continuous semiring. To see this, consider the infimum $\bigsqcap_n x^n = 0$ and the assignment $h(x) = 1$ into the Viterbi semiring. Instead, *generalised* absorptive polynomials $\mathbb{S}^\infty(X)$ are the free absorptive, fully-continuous semiring (see Section 3.5).
- The semiring PosBool(X) is finite. It is thus trivially fully continuous and universal for all fully-continuous lattice semirings.

3.3 Infinitary Power

Given an element a of an absorptive semiring, what happens if we repeatedly multiply a with itself? Since multiplication is decreasing, the powers $a \geq a^2 \geq a^3 \geq \dots$ form a decreasing chain which, assuming full continuity, has an infimum. This infimum, which we denote as a^∞, plays a central role for the computation of greatest fixed points (see 4). We summarise its properties (cf. [DGNT19, Naa19]) and discuss a possible axiomatisation. A more general notion of infinite products is discussed in Section 3.4.

Definition 3.3.1. Let K be an absorptive, fully-continuous semiring. For each element $a \in K$, we define the *infinitary power* of a as

$$a^\infty := \bigsqcap_{n<\omega} a^n.$$

Alternatively, we can define a^∞ as the greatest fixed point of the operator $x \mapsto a \cdot x$. This is also the reason why use the notation a^∞ instead of a^ω.

Properties. We first observe that full continuity commutes with the semiring operations and fully-continuous homomorphisms.

Lemma 3.3.2. *Let K be absorptive and fully continuous. For all $a, b \in K$ and $A \subseteq K$,*

(1) $a^k \cdot a^\infty = a^\infty$ *and* $a^k + a^\infty = a^k$, *for all* $k < \omega$,

(2) $a^\infty \cdot a^\infty = a^\infty$ *and* $(a^\infty)^\infty = a^\infty$,

(3) *if* $a \geq b$, *then* $a^\infty \geq b^\infty$ *(monotonicity),*

(4) $h(a^\infty) = h(a)^\infty$ *for any fully-continuous homomorphism* $h\colon K \to K'$,

(5) $(ab)^\infty = a^\infty b^\infty$,

(6) $(a + b)^\infty = a^\infty + b^\infty$.

Proof. By full continuity, $a^k \cdot \prod_n a^n = \prod_n a^{n+k} = a^\infty$ for $k < \omega$. Further, $a^k \geq a^\infty$ by definition and hence $a^k + a^\infty = a^k$ by absorption. Claim (2) follows from (1), since $a^\infty \cdot \prod_n a^n = \prod_n a^\infty a^n = \prod_n a^\infty = a^\infty$, thus also $\prod_n (a^\infty)^n = \prod_n a^\infty = a^\infty$. Claims (3) and (4) follow directly from the definition, claim (5) from the Chain Splitting Lemma 3.2.8.

For (6), we clearly have $(a+b)^n \geq a^n + b^n$ for all $n < \omega$, and thus $(a+b)^\infty \geq a^\infty + b^\infty$. 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. □

Lemma 3.3.3. *Let K be absorptive and fully continuous. Then*

(1) $(\bigsqcup A)^\infty \geq \bigsqcup A^\infty$ *for every set $A \subseteq K$, where $A^\infty = \{a^\infty \mid a \in A\}$,*

(2) $(\prod_i a_i)^\infty = \prod_i a_i^\infty$ *for every decreasing ω-chain $(a_i)_{i<\omega}$ in K.*

(3) $(\prod D)^\infty = \prod D^\infty$ *for every downward-directed set.*

Proof. Claim (1) holds by monotonicity: for each $a \in A$, we have $\bigsqcup A \geq a$ and thus $(\bigsqcup A)^\infty \geq a^\infty$. So $(\bigsqcup A)^\infty$ is an upper bound for A^∞ and the claim follows. We note that (2) is implied by (3), but has a simpler proof by using the Chain Splitting Lemma 3.2.8:

$$\Big(\prod_{i<\omega} a_i\Big)^\infty = \prod_{n<\omega} \Big(\prod_{i<\omega} a_i\Big)^n \overset{(3.2.8)}{=} \prod_{i,n<\omega} a_i^n = \prod_{i<\omega} a_i^\infty.$$

For claim (3), recall that multiplication commutes with infima over downward-directed sets (Corollary 3.2.5). We further use that every set $\{d_1, \dots, d_n\}$ has a lower bound $d \in D$, so by monotonicity $d_1 \cdots d_n \geq d^n$. Thus,

$$\Big(\prod D\Big)^\infty = \prod_{n<\omega} \Big(\prod D\Big)^n = \prod_{n<\omega} \prod_{d_1 \in D} \cdots \prod_{d_n \in D} (d_1 \cdots d_n)$$
$$= \prod_{n<\omega} \prod_{d \in D} d^n = \prod_{d \in D} \prod_{n<\omega} d^n = \prod D^\infty. \qquad \square$$

We show in Section 3.5.1 that we even have equality in (1) for the semiring of absorptive polynomials. But this does not hold in general.

Example 3.3.4. Consider the ω-chain $a_i = \frac{i}{i+1}$ in the Viterbi semiring $([0,1], \max, \cdot, 0, 1)$. Then $(\bigsqcup_i a_i)^\infty = 1^\infty = 1$, but $\bigsqcup_i a_i^\infty = \bigsqcup_i 0 = 0$.

Axiomatisation in Kleene algebra. We have relied on full continuity to define the infinitary power and to derive its properties. Is there also a purely algebraic way to introduce such an operation,[2] without requiring completeness assumptions?

For a bit of context, let us briefly look at Kleene algebras. These are idempotent semirings with an additional star operation a^* and generalise regular expressions (see, e.g., [Koz94]). The intuition for the star operation is that $a^* = 1 + a + a^2 + a^3 + \cdots = \bigsqcup_{n<\omega} a^n$, but it can also be defined without an infinite supremum, by the following four axioms. In fact,

[2] This question was suggested to me by two anonymous reviewers at RAMiCS'21.

Kozen [Koz94] has shown that these axioms, together with the usual axioms for idempotent semirings, are complete for equalities between regular expressions:

$$1 + aa^* \le a^*, \qquad b + ax \le x \implies a^*b \le x,$$
$$1 + a^*a \le a^*, \qquad b + xa \le x \implies ba^* \le x.$$

The star operation is not interesting in our case, as $a^* = 1$ in absorptive semirings, but the axioms can serve as inspiration. More closely related to the infinitary power are Cohen's ω-algebras [Coh00] which extend Kleene algebras by an operation a^ω with additional axioms

$$a^\omega = aa^\omega, \qquad x \le ax + b \implies a^\omega + a^*b,$$

which are complete for equalities of ω-regular expressions. These axioms also hold for the infinitary power (i.e., if we define $a^\omega := a^\infty$ and, by absorption, $a^* = 1$). We show that by slightly strengthening the axioms, we can capture all of the properties in Lemma 3.3.2 and obtain a similar completeness result for expressions including the infinitary power.

Definition 3.3.5. An ${}^\infty$-*algebra* is an algebraic structure $(K, +, \cdot, {}^\infty, 0, 1)$ with an underlying absorptive semiring $(K, +, \cdot, 0, 1)$ such that the following two ${}^\infty$-axioms hold:

$$a^\infty = aa^\infty \;\; ({}^\infty\text{-abs}), \qquad x \le ax + b \implies x \le a^\infty x + b \;\; ({}^\infty\text{-ind}),$$

where $\le$ denotes the natural order.

For a complete axiomatisation, we can use the axioms (3) – (13) of [Koz94] for idempotent semirings together with commutativity $ab = ba$, absorption $1 + a = 1$, and the two new axioms (${}^\infty$-abs), (${}^\infty$-ind). It is easy to see that every ${}^\infty$-algebra induces an ω-algebra by setting $a^\omega = a^\infty$ and $a^* = 1$. As expected, Definition 3.3.1 satisfies both axioms.

Lemma 3.3.6. *Every absorptive, fully-continuous semiring is an* ${}^\infty$*-algebra (by defining* a^∞ *as in Definition 3.3.1).*

Proof. Axiom (${}^\infty$-abs) holds by Lemma 3.3.8. For (${}^\infty$-ind), let $x \le ax + b$. Then

$$x \;\le\; ax + b \;\le\; a(ax + b) + b \;=\; a^2x + b$$

by absorption, and inductively $x \le a^n x + b$ for all $n \ge 1$. Using full continuity,

$$x \;\le\; \prod_{n\ge1}(a^n x + b) = \Big(\prod_{n\ge1} a^n\Big)x + b = a^\infty x + b. \qquad \square$$

Conversely, we show how to derive the properties of the infinitary power from the new axioms. The first lemma is a stronger form of monotonicity (consider $c = 1$).

Lemma 3.3.7. *In every* ${}^\infty$*-algebra,* $ac \le bc \implies a^\infty c \le b^\infty c$.

Proof. Assume $ac \le bc$. Using (${}^\infty$-abs), we get $a^\infty c = a^\infty ac \le a^\infty bc = b(a^\infty c)$. Then $a^\infty c \le b(a^\infty c) + b^\infty c$. Finally, (${}^\infty$-ind) and absorption give $a^\infty c \le b^\infty(a^\infty c) + b^\infty c = b^\infty c$. $\square$

We can also derive the remaining properties in Lemma 3.3.2 (except for homomorphisms, which have to commute with $^\infty$ by definition).

Lemma 3.3.8. *Every $^\infty$-algebra satisfies the following properties:*

(1) $a^k \cdot a^\infty = a^\infty$ *and* $a^k + a^\infty = a^k$*, for all* $k < \omega$*,*

(2) $a^\infty a^\infty = a^\infty$*,*

(3) $(a^\infty)^\infty = a^\infty$*,*

(4) $(ab)^\infty = a^\infty b^\infty$*,*

(5) $(a+b)^\infty = a^\infty + b^\infty$*.*

Proof. (1) Follows inductively from ($^\infty$-abs) and absorption.

(2) Trivially $a^\infty a^\infty \leq a^\infty$ by absorption. For the other direction, we have $a^\infty \leq aa^\infty$ by ($^\infty$-abs) and thus $a^\infty \leq a^\infty a^\infty$ by ($^\infty$-ind).

(3) By ($^\infty$-abs) and absorption, $(a^\infty)^\infty = a^\infty \cdot (a^\infty)^\infty \leq a^\infty$. Conversely, $a^\infty \leq a^\infty a^\infty$ by (2) and thus $a^\infty \leq (a^\infty)^\infty a^\infty \leq (a^\infty)^\infty$ by ($^\infty$-ind) and absorption.

(4) We have $a^\infty b^\infty \leq (ab) a^\infty b^\infty$ by ($^\infty$-abs) and thus $a^\infty b^\infty \leq (ab)^\infty a^\infty b^\infty \leq (ab)^\infty$ by ($^\infty$-ind) and absorption. For the other direction, we observe that $(ab)^\infty = ab(ab)^\infty \leq a(ab)^\infty$ by ($^\infty$-abs) and absorption. Hence $(ab)^\infty \leq a^\infty (ab)^\infty$ by ($^\infty$-ind) and symmetrically, $(ab)^\infty \leq b^\infty (ab)^\infty$. Combining both, we get $(ab)^\infty \leq a^\infty b^\infty (ab)^\infty \leq a^\infty b^\infty$, with absorption in the last step.

(5) By monotonicity, $a + b \geq a, b$ implies $(a+b)^\infty \geq a^\infty, b^\infty$ and thus $(a+b)^\infty \geq a^\infty + b^\infty$ by idempotence (addition is supremum). Conversely, we have $(a+b)^\infty = a(a+b)^\infty + b(a+b)^\infty$ by ($^\infty$-abs). Thus $(a+b)^\infty \leq a^\infty(a+b)^\infty + b^\infty(a+b)^\infty \leq a^\infty + b^\infty$ by applying ($^\infty$-ind) twice, with absorption in the last step. □

The above lemmas provide a sanity check for our axiomatisation. It is further easy to see that neither of ($^\infty$-abs) and ($^\infty$-ind) is redundant, as the remaining axiom would be satisfied by either setting $a^\infty := a$ or $a^\infty := 0$, which, in general, violates the property $a^\infty = \prod_n a^n$ we want to axiomatise. To further justify our axioms, we prove a completeness result.

Proposition 3.3.9. *Let α, β be two terms constructed from variables $x_1, \dots, x_n$, the constants* 0 *and* 1*, and the operations* $+$*,* $\cdot$ *and* $^\infty$*. Then, $\alpha = \beta$ holds in every absorptive, fully-continuous semiring if, and only if, it holds in every $^\infty$-algebra.*

Proof. Let $X = \{x_1, \dots, x_n\}$. We use the semiring $\mathbb{S}^\infty(X)$ in this proof, which is formally defined in Section 3.5. One implication is trivial, since every absorptive, fully-continuous semiring is an $^\infty$-algebra. For the converse, assume that $\alpha = \beta$ holds in every absorptive, fully-continuous semiring. We show that we can derive $\alpha = \beta$ from the $^\infty$-axioms and the properties of absorptive semirings.

Using Lemma 3.3.8 and distributivity, we can rewrite each term into a sum of monomials (with exponents from $\mathbb{N}^\infty$). We can thus derive equalities of the form $\alpha = m_1 + \dots + m_k$

and $\beta = m'_1 + \dots + m'_l$. We view $P_\alpha \coloneqq m_1 + \dots + m_k$ and $P_\beta \coloneqq m'_1 + \dots + m'_l$ as absorptive polynomials in $\mathbb{S}^\infty(X)$. Then $P_\alpha = P_\beta$ by assumption, so $\mathsf{Maximals}\{m_1, \dots, m_k\} = \mathsf{Maximals}\{m'_1, \dots, m'_l\}$.

Let m, m' be two monomials over X (with exponents in $\mathbb{N}^\infty$) such that $m \succeq m'$ (m absorbs m'). We claim that we can derive $m \geq m'$ from the axioms. Since $\cdot$ is monotone and X finite, we can reason about each variable x in m, m' separately. So assume that $x^a \succeq x^b$, thus $a \leq b$. If $a, b < \infty$, we can derive $x^a \geq x^b$ by absorption. If $b = \infty$, then we can derive $x^a \geq x^b$ using Lemma 3.3.8.

With the claim established, it follows that we can derive $\alpha = \sum \mathsf{Maximals}\{m_1, \dots, m_k\}$, and analogously for β. Hence $\alpha = \beta$ holds in every $^\infty$-algebra. □

3.4 Infinitary Summation and Product

The goal of this section is to extend the addition and multiplication operations of semirings to infinite sums and products. These operations should behave in a natural way, satisfying infinitary versions of the standard algebraic laws such as associativity and distributivity. To make this precise, we compile a "wish list" of properties that one might expect from such infinitary operations, and we identify a subset of these properties that we propose as indispensable requirements for semirings with infinitary operations. We show that these requirements can be met for all absorptive, fully-continuous semirings, and further that a universal semiring with infinitary operations exists.

These results lay the foundations to extend semiring semantics to infinite structures (which we do not consider in this thesis), and to infinitary logic (see Section 5.3). The material of this section is joint work with Sophie Brinke, Erich Grädel, and Lovro Mrkonjić and will appear in [BGMN24].

3.4.1 Infinitary Operations Wish List

Notation. Let K be a fixed semiring. For a family $(a_i)_{i \in I}$ of elements $a_i \in K$ with arbitrary index set I, we denote the infinitary sum and product operations over this family as

$$\widehat{\sum_{i \in I}} a_i, \qquad \text{and} \qquad \widehat{\prod_{i \in I}} a_i, \qquad \text{respectively.}$$

We use the symbols $\widehat{\sum}$ and $\widehat{\prod}$ throughout this thesis to distinguish between infinite and finite sums and products. Notice that we allow repetitions (i.e., $a_i = a_j$ for $i \neq j$) so that non-idempotent operations can be defined. However, as we are using an index *set* I, we implicitly assume commutativity of the infinitary operations; this is justified by the commutativity requirement on the following list. We always use the letters I and J for index sets, which can be infinite unless explicitly mentioned otherwise. The letters F, G, H always refer to finite index sets, and a_i, b_i denote semiring elements.

Infinite operations wish list. We begin with two essential requirements. Our definition also covers finite index sets I, and in this case the infinitary operations should coincide with the semiring operations. We also require commutativity to justify our notation.

- Compatibility with the finite: for finite index sets I,

$$\widehat{\sum_{i\in I}} a_i = \sum_{i\in I} a_i \qquad \text{and} \qquad \widehat{\prod_{i\in I}} a_i = \prod_{i\in I} a_i. \tag{F}$$

- Commutativity (or bijection invariance): for any bijection $f\colon J \to I$,

$$\widehat{\sum_{i\in I}} a_i = \widehat{\sum_{j\in J}} a_{f(j)} \qquad \text{and} \qquad \widehat{\prod_{i\in I}} a_i = \widehat{\prod_{j\in J}} a_{f(j)}. \tag{C}$$

We next formulate infinitary versions of associativity and distributivity. These properties turn out to be very powerful, so we include a weak and a strong version for both.

- Weak associativity: for disjoint index sets I and J,

$$\widehat{\sum_{i\in I\cup J}} a_i = \widehat{\sum_{i\in I}} a_i + \widehat{\sum_{j\in J}} a_j \qquad \text{and} \qquad \widehat{\prod_{i\in I\cup J}} a_i = \widehat{\prod_{i\in I}} a_i \cdot \widehat{\prod_{j\in J}} a_j. \tag{$\mathbf{A}_0$}$$

- Strong associativity (or partition invariance): for any partition $(I_j)_{j\in J}$ of I,

$$\widehat{\sum_{i\in I}} a_i = \widehat{\sum_{j\in J}}\widehat{\sum_{i\in I_j}} a_i \qquad \text{and} \qquad \widehat{\prod_{i\in I}} a_i = \widehat{\prod_{j\in J}}\widehat{\prod_{i\in I_j}} a_i. \tag{A}$$

- Weak distributivity: for every element $b \in K$,

$$b \cdot \widehat{\sum_{i\in I}} a_i = \widehat{\sum_{i\in I}} b a_i. \tag{$\mathbf{D}_0$}$$

- Strong distributivity: for any family $(I_j)_{j\in J}$ of subsets $I_j \subseteq I$,

$$\widehat{\prod_{j\in J}}\widehat{\sum_{i\in I_j}} a_i = \widehat{\sum_{f\in\mathcal{F}}}\widehat{\prod_{j\in J}} a_{f(j)}, \tag{D}$$

  where $\mathcal{F}$ is the set of choice functions, i.e., $f\colon J \to \bigcup_j I_j$ with $f(j) \in I_j$. This property is the analogue of complete distributivity in lattices, and the two notions are indeed equivalent if the semiring is a complete lattice (with $\widehat{\sum} \coloneqq \bigsqcup$ and $\widehat{\prod} \coloneqq \bigsqcap$).

We note that strong associativity implies commutativity. Indeed, for a given bijection $f\colon J \to I$ we can consider the partition $(I_j)_{j\in J}$ into singleton sets $I_j = \{f(j)\}$ (tacitly assuming that J is ordered in some way).

We also want infinitary operations to be monotone. This follows in a similar ways as for finite addition and multiplication: writing $a_i \geq b_i$ as $a_i = b_i + c_i$ for some c_i, monotonicity of infinite sums follows from (**A**) and (**F**); for multiplication we need (**D**) and ($\mathbf{A}_0$). However, strong distributivity (**D**) is a powerful property that may not always hold, whereas monotonicity appears to be weaker, so we list it separately:

- Monotonicity: if $a_i \geq b_i$ for all $i \in I$, then also

$$\widehat{\sum_{i\in I}} a_i \geq \widehat{\sum_{i\in I}} b_i \qquad \text{and} \qquad \widehat{\prod_{i\in I}} a_i \geq \widehat{\prod_{i\in I}} b_i. \tag{M}$$

Further natural properties concern the neutral elements:

- Annihilation (implied by (**F**)+(**A**$_0$)):

$$\text{if } a_i = 0 \text{ for some } i \in I, \text{ then also } \widehat{\prod_{i\in I}} a_i = 0. \qquad (\mathbf{Z})$$

- Neutral elements:

$$\widehat{\sum_{i\in I}} a_i = \widehat{\sum_{\substack{i\in I,\\ a_i\neq 0}}} a_i \qquad \text{and} \qquad \widehat{\prod_{i\in I}} a_i = \widehat{\prod_{\substack{i\in I,\\ a_i\neq 1}}} a_i. \qquad (\mathbf{N})$$

We now consider repeated addition and multiplication of a single value a. An important requirement is that infinite sums or products of a should produce a unique value, not depending on the cardinality of the index set. This is vital for the existence of a universal semiring with infinitary operations, as such a semiring would otherwise need representations for all infinite cardinalities, which is not possible. It further seems natural to assume that infinitary operations respect idempotent elements (in fact, this turns out to be rather subtle; we come back to this in Section 3.4.3).

- Idempotent elements: for every element $a \in K$,

$$a + a = a \implies \widehat{\sum_{i\in I}} a = a, \qquad a \cdot a = a \implies \widehat{\prod_{i\in I}} a = a. \qquad (\mathbf{I})$$

- Unique infinite powers: for every element $a \in K$ and all *infinite* sets I, J:

$$\widehat{\sum_{i\in I}} a = \widehat{\sum_{j\in J}} a \qquad \text{and} \qquad \widehat{\prod_{i\in I}} a = \widehat{\prod_{j\in J}} a. \qquad (\mathbf{U})$$

  We then write $\infty \cdot a$ and a^∞ for the unique infinite sum and product of a.

If we assume (**A**) and (**F**), then the implication (**I**) $\implies$ (**N**) holds. Indeed, neutral elements are idempotent, and infinite sums or products of neutral elements are the only reason why (**N**) might be violated. Moreover (**I**) $\implies$ (**U**), again assuming (**A**). The idea is as follows: because of (**C**), powers for countable sets $|I| = |J| = \omega$ must be equal, and these ω-powers are thus idempotent elements; larger index sets I, J can be partitioned into countable sets by (**A**) (see [BGMN24] for more details).

A noteworthy alternative to the last two items is the following property, which intuitively asserts that sums (or products) over infinite families can only be different if there is already a difference for finite subsets.

- Compactness: for two (finite or infinite) families $(a_i)_{i\in I}$, $(b_j)_{j\in J}$,

$$\Big\{\sum_{i\in F} a_i \;\Big|\; F \subset_{\text{fin}} I\Big\} = \Big\{\sum_{j\in G} a_j \;\Big|\; G \subset_{\text{fin}} J\Big\} \implies \widehat{\sum_{i\in I}} a_i = \widehat{\sum_{j\in J}} a_j, \qquad (\mathbf{Cp})$$

  and analogously also for infinitary products. Recall that $F \subset_{\text{fin}} I$ denotes the finite subsets of I.

Compactness seems to be a strong constraint. Together with (**F**), it implies that idempotent elements are respected (by choosing $a_i = b_j = a$ and $|J| = 1$). We refer to [BGMN24, Mrk24] for more details, including a stronger version of compactness and further implications.

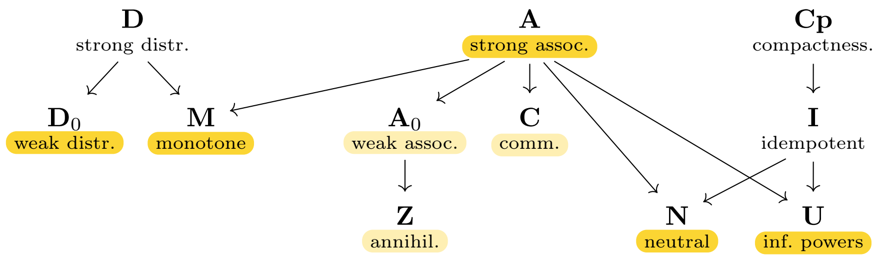

Figure 3.3: Implications between properties of infinitary operations, always assuming compatibility with the finite (**F**). For example, (**D**), (**A**), (**F**) together imply (**M**). The properties that hold in infinitary semirings are highlighted in yellow.

Infinitary semirings. The wish list, as per usual, has become quite long (see Fig. 3.3 for an overview). We propose the following definition of semirings with infinitary operations, based on the most relevant properties on our list. This is, of course, not the only possible definition. Our choice is guided by the analysis of infinitary operations in absorptive, fully-continuous semirings (Section 3.4.2) and by the goal of having a universal semiring (Section 3.4.3).

Definition 3.4.1. An *infinitary semiring* is a naturally-ordered semiring $(K, +, \cdot, 0, 1)$ extended by infinitary operations $\widehat{\sum}$ and $\widehat{\prod}$ that satisfy (**F**), (**A**), ($\mathbf{D}_0$), (**N**), (**U**), and (**M**).

A homomorphism $h\colon K_1 \to K_2$ between infinitary semirings is called *infinitary* if it commutes with the infinitary operations, i.e., $h\big(\widehat{\sum}_{i\in I} a_i\big) = \widehat{\sum}_{i\in I} h(a_i)$ and $h\big(\widehat{\prod}_{i\in I} a_i\big) = \widehat{\prod}_{i\in I} h(a_i)$ for all families $(a_i)_{i\in I}$ in K_1.

3.4.2 Absorptive semirings

Are there examples of infinitary semirings, or even semirings that fulfil our entire wish list? It turns out that for absorptive, fully continuous semirings, the answer is mostly positive; only strong distributivity fails in several semirings. The idea is to look at sums (or products) over finite subfamilies of $(a_i)_{i\in I}$. Summing more elements leads to larger values (by definition of the natural order) and multiplying more elements leads to smaller values (by absorption), which motivates the following definition of infinitary operations.

Theorem 3.4.2. *Every absorptive, fully continuous semiring can be extended to an infinitary semiring, which additionally satisfies* (**Cp**) *and* (**I**), *by*

$$\widehat{\sum_{i\in I}} a_i := \bigsqcup \Big\{ \sum_{i\in F} a_i \mid F \subset_{\text{fin}} I \Big\}, \qquad \textit{and} \qquad \widehat{\prod_{i\in I}} a_i := \bigsqcap \Big\{ \prod_{i\in F} a_i \mid F \subset_{\text{fin}} I \Big\},$$

In the following, we always assume this definition of infinitary operation when we work with absorptive, fully-continuous semirings. Before we prove the theorem, we make a number of observations. Infinitary summation can be equivalently written as $\widehat{\sum}_{i\in I} a_i = \bigsqcup_{i\in I} a_i$, as summation is just supremum in absorptive semirings. For products, we have $\widehat{\prod}_{i\in I} a_i \le \bigsqcap_{i\in I} a_i$ as in the finite case. To see this, consider the singleton sets $F = \{i\}$ for each $i \in I$.

Since the natural order of an absorptive, fully-continuous semiring is always a complete lattice, arbitrary suprema and infima exist and both operations are well-defined. In fact, we can be a bit more precise: the set $\{\sum_{i\in F} \mid F \subset_{\text{fin}} I\}$ of finite sub-sums is *upward directed*, the set of finite sub-products is *downward directed*. Indeed, for two finite subsets $F, F' \subset_{\text{fin}} I$ we can always take $F \cup F' \subset_{\text{fin}} I$. Hence Corollary 3.2.6 immediately implies that fully-continuous homomorphisms commute with infinite operations[3].

Proposition 3.4.3. *Let $h\colon K_1 \to K_2$ be a fully-continuous homomorphism on absorptive, fully-continuous semirings. Then h is an infinitary homomorphism.*

The infinitary operations are further compatible with the infinitary power. In fact, they have unique infinite powers (**U**) given by $\infty \cdot a = a$ (due to idempotence) and by the infinitary power $a^\infty = \bigsqcap_{n<\omega} a^n$.

Lemma 3.4.4. *Let K be an absorptive, fully-continuous semiring. Then $a^\infty = \widehat{\prod}_{i<\omega} a$ and $(\widehat{\prod}_{i\in I} a_i)^\infty = \widehat{\prod}_{i\in I} a_i^\infty$, for every element a and family $(a_i)_{i\in I}$ of K.*

Proof. The first claim follows directly from the definition, as $\widehat{\prod}_{i\in F} a = a^{|F|}$ for all $F \subset_{\text{fin}} \omega$. For the second claim, we apply Lemmas 3.3.2 and 3.3.3:

$$\Big(\widehat{\prod_{i\in I}} a_i\Big)^\infty = \Big(\bigsqcap_{F\subset_{\text{fin}} I} \widehat{\prod_{i\in F}} a_i\Big)^\infty = \bigsqcap_{F\subset_{\text{fin}} I} \widehat{\prod_{i\in F}} a_i^\infty = \widehat{\prod_{i\in I}} a_i^\infty. \qquad \square$$

We now prove Theorem 3.4.2. Most of the properties follow immediately from the definition, the only interesting case is strong associativity.

Proof of Theorem 3.4.2. To simplify notation, we use the abbreviation $a_F := \prod_{i\in F} a_i$ for a finite set $F \subset_{\text{fin}} I$. Hence $\widehat{\prod}_{i\in I} a_i = \bigsqcap_{F\subset_{\text{fin}} I} a_F$.

Recall that multiplication is decreasing due to absorption. Hence for finite I, the infimum of $\{a_F \mid F \subset_{\text{fin}} I\}$ is assumed for $F = I$ (and analogously for supremum and increasing addition), implying compatibility with the finite (**F**). Compactness (**Cp**) holds by definition, for both $\widehat{\sum}$ and $\widehat{\prod}$. Weak distributivity ($\mathbf{D}_0$) follows from finite distributivity and the fact that multiplication commutes with arbitrary suprema (by Corollary 3.2.5). We further have monotonicity (**M**), as all operators appearing in the definition are monotone.

It remains to prove strong associativity (the other properties are implied, cf. Fig. 3.3). Infinitary summation is the same as supremum and thus associative, but infinite products need more attention. Let I be an index set and $(I_j)_{j\in J}$ an arbitrary partition of I. We have to show:

$$\widehat{\prod_{i\in I}} a_i = \bigsqcap_{H\subset_{\text{fin}} I} a_H \overset{!}{=} \bigsqcap_{F\subset_{\text{fin}} J} \Big(\widehat{\prod_{j\in F}} \bigsqcap_{G\subset_{\text{fin}} I_j} a_G\Big) = \widehat{\prod_{j\in J}} \widehat{\prod_{i\in I_j}} a_i.$$

We prove both directions. First fix a finite set $H \subset_{\text{fin}} I$. Since $(I_j)_{j\in J}$ is a partition, there is a finite set $F \subset_{\text{fin}} J$ such that $H \subseteq \bigcup_{j\in F} I_j$. Moreover, for each $i \in I_j$ we clearly have

[3]This property is quite important, as it means that we can combine semiring semantics on infinite structures and homomorphisms in the same way as on finite structures. Perhaps, this property should have been part of our wish list in the first place.

$a_i \geq \bigsqcap_{G \subset_{\text{fin}} I_j} a_G$ by considering $G = \{i\}$. By monotonicity and absorption,

$$a_H = \widehat{\prod_{i \in H}} a_i \overset{(\text{abs})}{\geq} \widehat{\prod_{\substack{i \in I_j \\ j \in F}}} a_i \overset{(\text{mon})}{\geq} \widehat{\prod_{j \in F}} \bigsqcap_{G \subset_{\text{fin}} I_j} a_G \geq \bigsqcap_{F \subset_{\text{fin}} J} \widehat{\prod_{j \in F}} \bigsqcap_{G \subset_{\text{fin}} I_j} a_G.$$

For the other direction, fix a finite set $F \subset_{\text{fin}} J$ and write $F = \{j_1, \ldots, j_k\}$. Recall that a_F is a finite product and thus associative. Together with continuity of multiplication, we get

$$\begin{aligned}
\bigsqcap_{H \subset_{\text{fin}} I} a_H \leq \bigsqcap_{H \subset_{\text{fin}} \bigcup_{j \in F} I_j} a_H &= \bigsqcap_{H_{j_1} \subset_{\text{fin}} I_{j_1}} \cdots \bigsqcap_{H_{j_k} \subset_{\text{fin}} I_{j_k}} a_{(H_{j_1} \cup \cdots \cup H_{j_k})} \\
&\overset{(\text{assoc})}{=} \bigsqcap_{H_{j_1} \subset_{\text{fin}} I_{j_1}} \cdots \bigsqcap_{H_{j_k} \subset_{\text{fin}} I_{j_k}} (a_{H_{j_1}} \cdot \ldots \cdot a_{H_{j_k}}) \\
&\overset{(\text{cont})}{=} \left(\bigsqcap_{H_{j_1} \subset_{\text{fin}} I_{j_1}} a_{H_{j_1}} \right) \cdots \left(\bigsqcap_{H_{j_k} \subset_{\text{fin}} I_{j_k}} a_{H_{j_k}} \right) = \widehat{\prod_{j \in F}} \bigsqcap_{G \subset_{\text{fin}} I_j} a_G,
\end{aligned}$$

which closes the proof. □

The only property which does not hold in all absorptive, fully-continuous semirings is strong distributivity (**D**), as seen in the following example.

Example 3.4.5. Consider the Boolean algebra of regular[4] open subsets of $(0, 1) \subseteq \mathbb{R}$. It is known to be complete, but not completely distributive [HG09, Chapter 10].

By forgetting about the complement operation, we can view any Boolean algebra as a distributive, bounded lattice – and thus as an absorptive semiring. In this particular case, the lattice is complete, but not completely distributive. One can further show that it is fully continuous (see [HG09, Chapter 8]). The infinite operations of Theorem 3.4.2 are simply supremum and infimum, so that strong distributivity coincides with complete distributivity, which does not hold in this case.

The above example is based on a counterexample for complete distributivity. This leads to the follow-up question whether complete distributivity (of the natural order) implies strong distributivity (of the infinitary operations). This is not the case:

Example 3.4.6. Consider the Viterbi semiring $\mathbb{V} = ([0, 1], \max, \cdot, 0, 1)$ and the family $a_i = \frac{i}{i+1}$ for $i < \omega$. Let further J be an *uncountable* index set and let $I_j = \omega$ for all $j \in J$. That is, we multiply the same sum uncountably many times, which gives

$$\widehat{\prod_{j \in J}} \widehat{\sum_{i \in I_j}} \frac{i}{i+1} = \widehat{\prod_{j \in J}} \bigsqcup_{i \in I_j} \frac{i}{i+1} = \widehat{\prod_{j \in J}} 1 = 1.$$

On the other hand, we claim that

$$\widehat{\sum_{f \in \mathcal{F}}} \widehat{\prod_{j \in J}} a_{f(j)} = 0.$$

[4] An open set is *regular* if it is equal to the interior of its closure. Intuitively, this means that the set has no holes of only a single element.

To see this, fix an arbitrary choice function f (here: any function $f\colon J \to \omega$). Since J is uncountable and ω countable, there must be an index $i < \omega$ that is hit infinitely (even uncountably) often. But multiplying any a_i infinitely often with itself gives $a_i^\infty = 0$. Indeed, if $J' = f^{-1}(\{i\})$ are the indices in J mapped to i, then

$$\widehat{\prod_{j\in J}} a_{f(j)} \overset{(*)}{\leq} \widehat{\prod_{j\in J'}} a_{f(j)} = \bigsqcap_{H\subset_{\text{fin}} J'} a_i^{|H|} = \bigsqcap_{n<\omega} a_i^n = 0,$$

where $(*)$ holds by $(\mathbf{A}_0)$ and absorption. As f was arbitrary, this proves the claim, violating strong distributivity.

The Vitberi semiring is completely distributive (due to its total order), so this gives a negative answer to our question. However, this counterexample is somewhat unsatisfactory, as uncountable products in the Viterbi semiring are not very interesting. To see this, consider an uncountable family $(a_i)_{i\in I}$. Partition the domain $[0,1]$ of the Viterbi semiring in countably many intervals, say $[\frac{n}{n+1}, \frac{n+1}{n+2})$ for $n < \omega$, and $\{1\}$. Assuming that there are uncountably many i with $a_i \neq 1$, there must be some n and an infinite set $J \subseteq I$ such that $a_i \in [\frac{n}{n+1}, \frac{n+1}{n+2})$ for all $i \in J$. But then $\widehat{\prod}_{i\in I} a_i \leq \widehat{\prod}_{i\in J} a_i \leq \left(\frac{n+1}{n+2}\right)^\infty = 0$. In other words, products of uncountably many values (different from 1) always equal 0.

For countable products, the following argument (cf. [BGMN24]) shows that strong distributivity holds in the tropical semiring.

Proposition 3.4.7 ([BGMN24]). *The tropical semiring satisfies strong distributivity for countable index sets.*

Proof. Recall the operations of the tropical semiring: semiring addition is minimum, semiring multiplication is addition of real numbers, and the natural order is the inverse of the standard order. For this example, we work only with the standard operations on real numbers. Since all index sets are countable, we can assume w.l.o.g. (due to $(\mathbf{C})$) that all index sets are ω. We thus have to prove:

$$\sum_{j<\omega} \inf_{i<\omega} a_{i,j} = \inf_{f\in\mathcal{F}} \sum_{j<\omega} a_{f(j),j}.$$

Here, $\sum$ refers to the standard (infinitary) addition on real numbers, where diverging sums have value ∞. Notice that we added j as second subscript in $a_{i,j}$ to account for the fact that $(\mathbf{D})$ allows different index sets I_j for each j.

Direction "$\leq$" is easy (in general): for each choice function f and each j, we have $\inf_i a_{i,j} \leq a_{f(j),j}$, and we can then apply monotonicity of $\sum$.

For direction "$\geq$", note that there is nothing to prove if the left-hand side is ∞. So assume that $b_j \coloneqq \inf_i a_{i,j} < \infty$ for all j, and also $c \coloneqq \sum_j b_j < \infty$. Let $\varepsilon > 0$. We construct a choice function f_ε so that $\sum_j a_{f_\varepsilon(j),j} \leq c + \varepsilon$. Since ε is arbitrary, this proves that the infimum on right-hand side is at most c.

To define $f_\varepsilon(j)$, consider the real number $b_j + \varepsilon \cdot 2^{-(j+1)} > b_j$. Since $b_j = \inf_{i<\omega} a_{i,j}$, there must be an index i such that $b_j + \varepsilon \cdot 2^{-(j+1)} > a_{i,j}$. We set $f_\varepsilon(j) = i$ for this index. Then

$$\sum_{j<\omega} a_{f_\varepsilon(j),j} \leq \sum_{j<\omega} \left(b_j + \varepsilon \cdot 2^{-(j+1)}\right) = \sum_{j<\omega} b_j + \varepsilon \cdot \sum_{j<\omega} 2^{-(j+1)} = c + \varepsilon. \qquad \square$$

This result also applies to the isomorphic Viterbi semiring, and a similar argument works for the Łukasiewicz semiring. Another example where strong distributivity holds (even without the cardinality restriction) is the semiring of generalised absorptive polynomials $\mathbb{S}^\infty(X)$, see Section 3.5.3.

3.4.3 Universal Semirings

The class of absorptive, fully-continuous semirings admits a uniform definition of infinitary operations. The universal absorptive and fully-continuous infinitary semiring is given by $\mathbb{S}^\infty(X)$, as fully-continuous homomorphisms are also infinitary. There are further semirings that can be equipped with, more specifically defined, infinitary operations. We show that a general (non-absorptive) universal semiring exists: the semiring $\mathrm{Poly}(\mathbb{N}^\infty, \mathbb{N}^\infty, X)$ of formal power series with coefficients and exponents in $\mathbb{N}^\infty$, which we denote as $\mathbb{N}^\infty_\infty[\![X]\!]$.

As a first step, we observe that the (extended) semiring of natural numbers $\mathbb{N}^\infty$ becomes an infinitary semiring under the natural infinitary operations. That is, $\widehat{\sum}_{i\in I} a_i = \infty$ whenever $a_i \neq 0$ for infinitely many i, and otherwise we take the finite sum over the nonzero entries. For the product, we set $\widehat{\prod}_{i\in I} a_i = 0$ if there is some $a_i = 0$. If $a_i > 1$ for infinitely many i, then $\widehat{\prod}_{i\in I} a_i = \infty$, and otherwise we take the finite product over all $a_i > 1$. One can show by a simple (but somewhat tedious) case distinction (see [BGMN24]) that these operations additionally satisfy strong distributivity, and also compactness.

Proposition 3.4.8 ([BGMN24])**.** *The semiring $\mathbb{N}^\infty$ is an infinitary semiring with strong distributivity and compactness (with the natural infinitary operations).*

We now consider the infinitary extension of the semiring $\mathbb{N}[X]$. We must be able to infinitely often add the same monomial (e.g., $x + x + x + \dots$), add infinitely many different monomials (e.g., $x + x^2 + x^3 + \dots$), and multiply the same indeterminate infinitely often (e.g., $x \cdot x \cdot x \cdot \dots$). It is not clear how to define these operations in $\mathbb{N}[X]$, but we can extend the semiring to address these issues, respectively, by allowing coefficients in $\mathbb{N}^\infty$, using formal power series instead of polynomials, and allowing exponents in $\mathbb{N}^\infty$. This results in the semiring of *generalised power series* $\mathrm{Poly}(\mathbb{N}^\infty, \mathbb{N}^\infty, X)$, or $\mathbb{N}^\infty_\infty[\![X]\!]$ for short. Notice that the definition of multiplication (cf. Definition 3.1.9) of two power series P, Q now becomes

$$(P \cdot Q)(m) = \widehat{\sum_{m_1 \cdot m_2 = m}} P(m_1) \cdot Q(m_2),$$

where $\widehat{\sum}$ is the infinitary summation in $\mathbb{N}^\infty$. Indeed, a monomial with exponents in $\mathbb{N}^\infty$, such as $m = x^\infty$, can have infinitely many splits $m = m_1 m_2$, for instance $m_1 = x^n$ and $m_2 = x^\infty$ for all n.

We equip $\mathbb{N}^\infty_\infty[\![X]\!]$ with the following infinitary operations. Infinite summation acts pointwise and infinite products can be defined by considering the sum over all possible factorisations of a given monomial, similar to the finite case. Formally,

$$\widehat{\sum_{i\in I}} P_i \coloneqq \Big(m \mapsto \widehat{\sum_{i\in I}} P_i(m)\Big), \qquad \widehat{\prod_{i\in I}} P_i \coloneqq \Big(m \mapsto \widehat{\sum_{(m_i)_i \in \mathsf{splits}(m)}} \widehat{\prod_{i\in I}} P_i(m_i)\Big).$$

Here, $\mathsf{splits}_I(m)$ is the set of sequences $(m_i)_{i\in I}$ of monomials with $m = \widehat{\prod}_{i\in I} m_i$ (this infinite product is defined by adding exponents in the infinitary semiring $\mathbb{N}^\infty$). We may omit the index I if it is clear from the context.

We next study the properties of these infinitary operations and prove that they satisfy all requirements of an infinitary semiring. However, compactness does not hold:

Example 3.4.9. Consider the power series $Q = \widehat{\sum}_{i<\omega} \infty \cdot x^i$ in $\mathbb{N}^\infty_\infty[\![x]\!]$. Then $Q^2 = Q$, but $Q^\omega = \widehat{\prod}_{i<\omega} Q = \infty \cdot x^\infty + Q$, which is different from Q. This shows that infinite products in $\mathbb{N}^\infty_\infty[\![X]\!]$, even in the case $|X| = 1$, do not respect idempotent elements (**I**), and hence compactness (**Cp**) cannot hold.

This example is one of the reasons why we did not include compactness in Definition 3.4.1. In order to prove that, despite this counterexample, $\mathbb{N}^\infty_\infty[\![X]\!]$ has unique infinite powers, we need the following technical lemma.

Lemma 3.4.10. *Let $P \in \mathbb{N}^\infty_\infty[\![X]\!]$ and I an infinite index set. The only possible coefficients of $\widehat{\prod}_{i\in I} P$ are 0, 1, and ∞.*

Proof. Let $Q = \widehat{\prod}_{i\in I} P$, fix a monomial m and assume that $Q(m) \neq 0$. Consider a sequence $(m_i)_{i\in I} \in \mathsf{splits}(m)$ with value $\widehat{\prod}_{i\in I} P(m_i) > 0$. If there is a monomial v that occurs only finitely often (but at least once) in $(m_i)_i$, then by permuting the occurrences of v with other monomials in the sequence, we obtain infinitely many pairwise different sequences in $\mathsf{splits}(m)$ with the same value. Since $Q(m)$ is the sum over all sequences in $\mathsf{splits}(m)$, this implies $Q(m) = \infty$. Similarly, if two different monomials v, v' occur infinitely often in $(m_i)_i$, we can again obtain infinitely many different permutations of the sequence and $Q(m) = \infty$.

So the only possibility for $Q(m) < \infty$ is that all sequences $(m_i)_i$ with $\widehat{\prod}_{i\in I} P(m_i) > 0$ consist of only one monomial, i.e. $m_i = m_j$ for all $i, j \in I$. Assume towards a contradiction that two such sequences exist, say $(v)_{i\in I} \in \mathsf{splits}(m)$ and $(w)_{i\in I} \in \mathsf{splits}(m)$. Since $|\omega| + |I| = |I|$, we can construct a sequence $(u_i)_{i\in I}$ that consists of countably many repetitions of v, and $|I|$ many repetitions of w. All indeterminates occurring in v or w must have exponent ∞ in m, and thus $(u_i)_i \in \mathsf{splits}(m)$. But this sequence uses two monomials infinitely often which implies $Q(m) = \infty$, contradiction.

Hence $Q(m) < \infty$ implies that there is only one sequence $(m_i)_i \in \mathsf{splits}(m)$ with value $\widehat{\prod}_{i\in I} P(m_i) > 0$, and since $m_i = m_j$ for all i, j this value must be either 1 or ∞. ☐

Lemma 3.4.11. *$\mathbb{N}^\infty_\infty[\![X]\!]$ has unique infinite powers.*

Proof. Infinite summation is defined pointwise, so the unique power property is inherited from $\mathbb{N}^\infty$. Multiplication requires more work. Notice that the set of monomials $\mathsf{Mon}_{\mathbb{N}^\infty}(X)$ is countable, as X is finite. Let $P \in \mathbb{N}^\infty_\infty[\![X]\!]$ and I, J two infinite index sets. It suffices to prove that

$$Q := \widehat{\prod_{i\in I}} P \ \leq \ \widehat{\prod_{j\in J}} P =: R$$

due to symmetry. Fix a monomial m and a sequence $(m_i)_{i\in I} \in \mathsf{splits}_I(m)$. Consider the set of monomials $\{m_i \mid i \in I\}$ occurring in this sequence. We partition this set into a set M_0 of monomials that occur only finitely often and a set M_∞ of monomials that occur infinitely often. Let $I_0 \subseteq I$ be the set of indices i with $m_i \in M_0$. Since M_0 is countable and each $m \in M_0$ occurs finitely often, I_0 is countable as well. We further fix an enumeration of the set M_∞ (which is countable as well).

We now construct a sequence $(v_j)_{j\in J} \in \mathsf{splits}_J(m)$ by first constructing a sequence $(w_l)_{l\in L} \in \mathsf{splits}_L(m)$ for some index set L and then applying a bijection $f\colon J \to L$. We define L as follows:

$$L := I_0 \ \dot\cup \ (\omega \times \omega) \ \dot\cup \ J.$$

We next define the elements of the sequence $(w_l)_{l\in L}$. We first need a "padding element" m_∞ (in case $|J| > |I|$). If $M_\infty \neq \emptyset$, we choose an arbitrary but fixed $m_\infty \in M_\infty$. If $M_\infty = \emptyset$, then M_0 contains infinitely many monomials (each of which occurs finitely often in $(m_i)_{i\in I}$). Recall that $m = \widehat{\prod}_{i\in I} m_i$ and consider the indeterminates $X_0 \subseteq X$ with *finite* exponents in m. Since the exponents are finite, there can be only finitely many monomials in M_0 containing an indeterminate in X_0 (recall that X is finite). Consider the infinitely many monomials in M_0 *not* containing an indeterminate in X_0. Among these, we choose m_∞ so that $P(m_\infty)$ is minimal. Notice that when $P(m_\infty) > 1$, then by minimality there are infinitely many monomials $m' \in M_0$ with $P(m') > 1$, and hence $\widehat{\prod}_{i\in I} P(m_i) = \infty$ (†).

We are now ready to define the sequence $(w_l)_{l\in L}$.

- if $l \in I_0$, we set $w_l = m_l$ (i.e., we copy all finitely often occurring monomials),
- if $l = (n, m) \in \omega \times \omega$, we set w_l to the n-th monomial in M_∞ (i.e., we repeat each monomial in M_∞ infinitely often); if $M_\infty = \emptyset$ we set $w_l = m_\infty$ (or we omit the $(\omega \times \omega)$-part),
- if $l \in J$, we set $w_l = m_\infty$ (this is only for padding so that L has the right cardinality).

By construction, the monomials in $(w_l)_{l\in L}$ match the monomials in $(m_i)_{i\in I}$ in the following sense: each monomial appears either infinitely often in both sequences, or the same finite number of times in both sequences. The only exception is m_∞ in the case $M_\infty = \emptyset$, but in this case we know that m_∞ contains only indeterminates that have exponent ∞ in m. It follows that the products of the two sequences result in the same monomial m, so $(w_l)_{l\in L} \in \mathsf{splits}_L(m)$ as claimed.

It further follows (using compactness of $\mathbb{N}^\infty$) that the values of the sequences are also equal: $\widehat{\prod}_{i\in I} P(m_i) = \widehat{\prod}_{l\in L} P(w_l)$. Again, the case $M_\infty = \emptyset$ needs special attention. If $P(m_\infty) = 0$, then both sequences have value 0 and are equal. If $P(m_\infty) = 1$, then repeating m_∞ does not affect the value of the sequence. If $P(m_\infty) > 1$, then we have $\widehat{\prod}_{l\in L} P(w_l) = \infty$ due to the padding, but in this case also $\widehat{\prod}_{i\in I} P(m_i) = \infty$ by (†) and the equality still holds.

We can finally define the sequence $(v_j)_{j\in J} \in \mathsf{splits}_J(m)$. Notice that $|L| = |J|$ since J is infinite and both I_0 and $\omega \times \omega$ are countable. We thus have a bijection $f\colon J \to L$ and can set $v_j = w_{f(j)}$. Since f is bijective, $(v_j)_{j\in J}$ has the same product and value as $(w_l)_{l\in L}$, and thus also as $(m_i)_{i\in I}$.

We still have to prove that $Q(m) \leq R(m)$. (This does not immediately follow from the above argument, since we have to sum over all sequences, but different sequences $(m_i)_{i\in I}$ could be mapped to the same sequence $(v_j)_{j\in J}$.) We proceed by a case distinction using Lemma 3.4.10. The case $Q(m) = 0$ is trivial. If $Q(m) = 1$, then the construction of

$(v_j)_{j \in J}$ witnesses $R(m) \geq 1$. The only other possibility is $Q(m) = \infty$. If there is a sequence $(m_i)_{i \in I} \in \mathsf{splits}_I(m)$ with value ∞, then by the above construction also $R(m) = \infty$. Otherwise there must be a sequence (in fact infinitely many) $(m_i)_{i \in I} \in \mathsf{splits}_I(m)$ with value $1 < s < \infty$. At least two distinct monomials must occur in $(m_i)_{i \in I}$ (otherwise the value would be 0, 1, or ∞), and, by construction, these monomials must also be contained in the sequence $(v_j)_{j \in J}$. It follows that there are infinitely many pairwise different permutations of $(v_j)_{j \in J}$, and since all of these sequences occur in the summation we have $R(m) = \infty$. □

The other requirements for infinitary semirings mostly follow by applying the properties of infinitary operations in $\mathbb{N}^\infty$ to coefficients and exponents.

Theorem 3.4.12. $\mathbb{N}^\infty_\infty[\![X]\!]$ *is an infinitary semiring with strong distributivity.*

Proof. It follows directly from the definition that addition and multiplication in $\mathbb{N}^\infty_\infty[\![X]\!]$ are compatible with finite operations and respect neutral elements, and we have already considered infinite powers in Lemma 3.4.11. It then suffices to prove strong associativity and strong distributivity, as these imply all remaining properties.

Strong associativity of addition follows immediately from the respective property of $\mathbb{N}^\infty$, as addition is defined by adding coefficients. For multiplication, fix a partition $(I_j)_{j \in J}$ of I. Using strong distributivity of $\mathbb{N}^\infty$, it remains to prove for each monomial m:

$$\begin{aligned}
\Big(\widehat{\prod_{j \in J}} \widehat{\prod_{i \in I_j}} P_i\Big)(m) &= \widehat{\sum_{(m_j)_j \in \mathsf{splits}_J(m)}} \widehat{\prod_{j \in J}} \widehat{\sum_{(v_i)_i \in \mathsf{splits}_{I_j}(m_j)}} \widehat{\prod_{i \in I_j}} P_i(v_i) \\
&= \widehat{\sum_{(m_j)_j \in \mathsf{splits}_J(m)}} \widehat{\sum_{f \in F}} \widehat{\prod_{j \in J}} \widehat{\prod_{i \in I_j}} P_i(f(j,i)) \\
&\stackrel{!}{=} \widehat{\sum_{(u_i)_i \in \mathsf{splits}_I(m)}} \widehat{\prod_{i \in I}} P_i(u_i).
\end{aligned}$$

Here, F is the set of choice functions that choose $(v_i)_i \in \mathsf{splits}_{I_j}(m_j)$ for each j. To simplify the presentation, we let $f(j,i) = v_i$ for the chosen sequence (i.e., we include the index i as argument).

We prove both directions of the last equality. First let $(u_i)_i \in \mathsf{splits}_I(m)$. Set $m_j = \widehat{\prod}_{i \in I_j} u_i$. Then $(m_j)_j \in \mathsf{splits}_J(m)$ by comparing exponents and strong associativity of $\mathbb{N}^\infty$. For $i \in I_j$, we define $f(j,i) = u_i$. Since each i occurs in exactly one I_j, we have $\widehat{\prod}_{j \in J} \widehat{\prod}_{i \in I_j} P_i(f(j,i)) = \widehat{\prod}_{i \in I} P_i(u_i)$ by strong associativity of $\mathbb{N}^\infty$. Notice that our construction $(u_i)_i \mapsto ((m_j)_j, f)$ is injective, so direction $\geq$ holds (by strong associativity of addition in $\mathbb{N}^\infty$).

For the other direction, let $(m_j)_j \in \mathsf{splits}_J(m)$ and $f \in F$. For each $i \in I$, pick the unique j with $i \in I_j$ and set $u_i = f(j,i)$. Then (by strong associativity of $\mathbb{N}^\infty$ in each exponent):

$$\widehat{\prod_{i \in I}} u_i = \widehat{\prod_{j \in J}} \widehat{\prod_{i \in I_j}} u_i = \widehat{\prod_{j \in J}} \widehat{\prod_{i \in I_j}} f(j,i) = \widehat{\prod_{j \in J}} m_j = m,$$

so $(u_i)_i \in \mathsf{splits}_I(m)$. Applying the same argument to the coefficients yields

$$\widehat{\prod_{i \in I}} P_i(u_i) = \widehat{\prod_{j \in J}} \widehat{\prod_{i \in I_j}} P_i(u_i) = \widehat{\prod_{j \in J}} \widehat{\prod_{i \in I_j}} P_i(f(j,i)).$$

Again, the mapping $((m_j)_j, f) \mapsto (u_i)_i$ we construct is injective, so direction $\leq$ holds as well.

To prove strong distributivity, let $(I_j)_{j\in J}$ be a partition of I and F be the set of choice functions f with $f(j) \in I_j$. Strong distributivity then follows from strong distributivity and associativity of $\mathbb{N}^\infty$:

$$\begin{aligned}
\Big(\widehat{\prod_{j\in J}}\widehat{\sum_{i\in I_j}} P_i\Big)(m) &= \widehat{\sum_{(m_j)_j\in \mathsf{splits}_J(m)}} \widehat{\prod_{j\in J}} \widehat{\sum_{i\in I_j}} P_i(m_j) \\
&= \widehat{\sum_{(m_j)_j\in \mathsf{splits}_J(m)}} \widehat{\sum_{f\in F}} \widehat{\prod_{j\in J}} P_{f(j)}(m_j) \\
&= \widehat{\sum_{f\in F}} \widehat{\sum_{(m_j)_j\in \mathsf{splits}_J(m)}} \widehat{\prod_{j\in J}} P_{f(j)}(m_j) \\
&= \Big(\widehat{\sum_{f\in F}} \widehat{\prod_{j\in J}} P_{f(j)}\Big)(m).
\end{aligned}$$

□

We now establish a universal property, showing that $\mathbb{N}^\infty_\infty[\![X]\!]$ is the freely generated infinitary semiring with strong distributivity.

Theorem 3.4.13 (Universality of $\mathbb{N}^\infty_\infty[\![X]\!]$)**.** *Let K be an infinitary semiring with strong distributivity. Every mapping $h\colon X \to K$ extends uniquely to an infinitary homomorphism $h\colon \mathbb{N}^\infty_\infty[\![X]\!] \to K$.*

Proof. For every element $s \in K$ there exist unique elements $\infty \cdot s = \widehat{\sum}_{i\in I} s$ and $s^\infty = \widehat{\prod}_{i\in I} s$ for all infinite index sets I. Thus, $n \cdot s$ and s^n are well-defined for all $s \in K$ and $n \in \mathbb{N}^\infty$. We first lift h to monomials m by setting $h(m) := \widehat{\prod}_{x\in X} h(x)^{m(x)}$. By strong associativity, h commutes with (finite and infinitary) products of monomials:

$$h\Big(\widehat{\prod_{i\in I}} m_i\Big) = \widehat{\prod_{x\in X}} h(x)^{\widehat{\sum}_{i\in I} m_i(x)} = \widehat{\prod_{x\in X}}\widehat{\prod_{i\in I}} h(x)^{m_i(x)} = \widehat{\prod_{i\in I}}\widehat{\prod_{x\in X}} h(x)^{m_i(x)} = \widehat{\prod_{i\in I}} h(m_i).$$

We write power series $P \in \mathbb{N}^\infty_\infty[\![X]\!]$ as $P = \widehat{\sum}_{m\in M}(P(m)\cdot m)$. Then h is uniquely defined by $h(P) := \widehat{\sum}_{m\in M}(P(m)\cdot h(m))$. We need to show that h commutes with the (finite and infinitary) semiring operations. Since infinitary operations are compatible with the finite ones, it suffices to prove that h commutes with the infinitary sum and product. This is easy for summation due to strong associativity (notice that $\widehat{\sum}_i P_i(m) \in \mathbb{N}^\infty$ and the multiplication with $h(m)$ is just an abbreviation for repeated addition):

$$h\Big(\widehat{\sum_{i\in I}} P_i\Big) = \widehat{\sum_{m\in M}}\Big(\widehat{\sum_{i\in I}} P_i(m)\Big)\cdot h(m) = \widehat{\sum_{m\in M}}\widehat{\sum_{i\in I}}(P_i(m)\cdot h(m)) = \widehat{\sum_{i\in I}} h(P_i).$$

For products, we use strong associativity and strong distributivity of K. Let F be the set of (unrestricted) functions $f\colon I \to M$. Then,

$$\widehat{\prod_{i\in I}} h(P_i) = \widehat{\prod_{i\in I}}\widehat{\sum_{m\in M}}(P_i(m)\cdot h(m)) \overset{(\mathbf{D})}{=} \widehat{\sum_{f\in F}}\widehat{\prod_{i\in I}} P_i(f(i))\cdot h(f(i))$$

$$\overset{(\mathbf{A})}{=} \widehat{\sum_{m \in M}} \widehat{\sum_{(m_i)_i \in \mathsf{splits}(m)}} \Big(\widehat{\prod_{i \in I}} (P_i(m_i) \cdot h(m_i)) \Big)$$
$$\overset{(\mathbf{A})}{=} \widehat{\sum_{m \in M}} \widehat{\sum_{(m_i)_i \in \mathsf{splits}(m)}} \Big(\widehat{\prod_{i \in I}} P_i(m_i) \cdot \widehat{\prod_{i \in I}} h(m_i) \Big)$$
$$= \widehat{\sum_{m \in M}} \widehat{\sum_{(m_i)_i \in \mathsf{splits}(m)}} \Big(\Big(\widehat{\prod_{i \in I}} P_i(m_i) \Big) \cdot h(m) \Big)$$
$$\overset{(\mathbf{A})}{=} \widehat{\sum_{m \in M}} \Big(\widehat{\sum_{(m_i)_i \in \mathsf{splits}(m)}} \prod_{i \in I} P_i(m_i) \Big) \cdot h(m) = h\Big(\prod_{i \in I} P_i \Big). \qquad \square$$

Remark 3.4.14. The universal property can be generalised to infinitary semirings K in which strong distributivity only holds for index sets of a certain cardinality. More precisely, we say that K is *κ-distributive* for a cardinal κ (cf. [BGMN24]), if for any set I and any family $(I_j)_{j \in J}$ of subsets with $|J| < \kappa$,

$$\widehat{\prod_{j \in J}} \widehat{\sum_{i \in I_j}} a_i = \widehat{\sum_{f \in \mathcal{F}}} \widehat{\prod_{j \in J}} a_{f(j)}.$$

In this case, an analysis of the above proof shows that the homomorphism $h \colon \mathbb{N}^{\infty}_{\infty}[\![X]\!] \to K$ is no longer infinitary, but still commutes with arbitrary infinite sums, and with infinite products over families I of cardinality $|I| < \kappa$. This may be interesting, for instance, to cover cases such as the ω-distributive tropical semiring.

Remark 3.4.15. Notice that Theorem 3.4.13 does not hold for the semiring $\mathbb{N}^{\infty}[\![X]\!]$ without the exponent ∞. It is not clear how the infinite power $x \cdot x \cdot x \cdots$ is then defined, but we can argue by case distinction that in any case, the universal property is violated.

- If $\widehat{\prod}_{i<\omega} x = 0$, then $h(\widehat{\prod}_{i<\omega} x) = 0 \neq 1 = \widehat{\prod}_{i<\omega} h(x)$ for $h(x) = 1$ (say into the Viterbi semiring).
- If $P = \widehat{\prod}_{i<\omega} x \neq 0$, then there must a monomial m (with finite exponents) and coefficient $P(m) = n > 0$. For $h(x) = \frac{1}{2}$ into the Viterbi semiring, we then get $h(P) \geq h(n \cdot m) = n \cdot h(m) = h(m) > 0$, but also $\widehat{\prod}_{i<\omega} h(x) = (\frac{1}{2})^{\infty} = 0$, contradiction.

On the other hand, $\mathbb{N}^{\infty}[\![X]\!]$ is fully continuous (by Proposition 3.2.10), but the following example shows that $\mathbb{N}^{\infty}_{\infty}[\![X]\!]$ is not. This means that $\mathbb{N}^{\infty}_{\infty}[\![X]\!]$ is not well-suited as a provenance semiring for fixed-point logics, but may be of interest for a provenance analysis involving infinite structures.

Example 3.4.16. Let $P_i = \widehat{\sum}_{i<j<\omega} x^j$ for each $i < \omega$, and observe that this defines a decreasing chain in $\mathbb{N}^{\infty}_{\infty}[\![X]\!]$. As the monomial x^i disappears in the i-th step, this chain has infimum 0. We thus have $x^{\infty} \cdot \bigsqcap_i P_i = 0$, whereas $\bigsqcap_i (x^{\infty} \cdot P_i) = \bigsqcap_i \infty \cdot x^{\infty} = \infty \cdot x^{\infty}$. Hence multiplication does not commute with infima of chains.

Universal infinitary semirings also exist for smaller classes of semirings. In the idempotent case, the above proofs can easily be adapted to show that $\mathrm{Poly}(\mathbb{B}, \mathbb{N}^{\infty}, X)$ is the universal idempotent infinitary semiring with strong distributivity. We have already considered the absorptive case, where the universal semiring is $\mathbb{S}^\infty(X)$. For lattice semirings, PosBool(X) is trivially infinitary (with strong distributivity) due to its finiteness.

A challenging question is whether these semirings can be generalised to infinite indeterminate sets. We study this in detail for $\mathbb{S}^\infty(X)$ and PosBool(X) in Section 3.6. Here we only note that for an infinite set $\boldsymbol{\mathcal{X}}$, the semiring $\mathbb{N}^\infty_\infty[\![\boldsymbol{\mathcal{X}}]\!]$ is no longer infinitary, as it does not have unique infinite powers.

Example 3.4.17. Let $\boldsymbol{\mathcal{X}} = \{x_1, x_2, \dots\}$ be a countably infinite set of indeterminates. Consider $P = \widehat{\sum}_{i<\omega} x_i \in \mathbb{N}^\infty_\infty[\![\boldsymbol{\mathcal{X}}]\!]$. Then $P^\omega = \widehat{\prod}_{i<\omega} P$ is different from $P^{\omega_1} = \widehat{\prod}_{\alpha<\omega_1} P$. Indeed, since ω_1 is uncountable, one of the countably many monomials x_i of P has to be multiplied infinitely often for each monomial of P^{ω_1}, and hence all monomials of P^{ω_1} have ∞ as exponent. In contrast, P^ω contains the monomial $x_1 \cdot x_2 \cdot x_3 \cdots$ with 1 as only exponent.

3.5 Absorptive Polynomials

Semirings that are absorptive and fully continuous play a major role in this thesis, and we are thus naturally interested in the universal semiring with these properties. The universal absorptive semiring has first been introduced as Sorp(X) in [DMRT14] (also called $\mathbb{S}[X]$ in [GT20]), with two equivalent definitions: as a quotient of $\mathbb{N}[X]$ by absorption, and as antichain polynomials with respect to monomial absorption.

The idea is that "shorter" monomials absorb "longer" ones (more precisely, we compare the exponents of each indeterminate). For example, x^2y absorbs both x^2yz and x^2y^2, but not xy^3. Absorptive polynomials are then antichain polynomials, that is, *sums of monomials such that none of them absorbs another* [DMRT14]. From a provenance perspective, absorption means that only *minimal* witnesses are represented.

In our context, Sorp(X) has the issue that it is universal only for ω-continuous absorptive semirings, but not for fully-continuous absorptive semirings (cf. Section 3.2.2). This can be solved by including the exponent ∞, thus switching to the semiring $\mathbb{S}^\infty(X)$ of *generalised* absorptive polynomials, which was introduced in [GT20] and shown to be universal in [Naa19, DGNT21]. In this thesis, we focus only on the generalised version $\mathbb{S}^\infty(X)$ and, for simplicity, refer to its elements simply as *absorptive polynomials*.

We summarize the main properties of $\mathbb{S}^\infty(X)$ shown in [Naa19, DGNT21] in this section, most notably its universal property. New contributions are the proof of strong distributivity and the extensions of $\mathbb{S}^\infty(X)$ discussed in Sections 3.6 and 3.7.

3.5.1 Definition via Antichains

Let X be a finite set of indeterminates. Throughout this section, let Mon refer to the set $\mathsf{Mon}_X(\mathbb{N}^\infty)$ of monomials $m\colon X \to \mathbb{N}^\infty$. The semiring $\mathbb{S}^\infty(X)$ can be defined as a quotient of the power series semiring Poly$(\mathbb{N}, \mathbb{N}^\infty, X)$ by absorption (cf. Section 3.6), but here we follow the original explicit definition based on antichains.

Definition 3.5.1 (Absorption)**.** Given two monomials $m_1, m_2 \in \mathsf{Mon}$, we say that m_1 *absorbs* m_2, denoted $m_1 \succeq m_2$, if m_1 has smaller exponents than m_2, i.e., $m_1(x) \leq m_2(x)$ for all $x \in X$. We further write $\mathsf{Maximals}(M)$ for the set of $\succeq$-maximal monomials in $M \subseteq \mathsf{Mon}$.

Absorption can be seen as the pointwise partial order induced by the reverse order on $\mathbb{N}^\infty$. The set Mon of monomials inherits a lattice structure from $\mathbb{N}^\infty$ and is, of course, infinite. However, it has some crucial finiteness properties.

Lemma 3.5.2. *Every antichain of monomials (w.r.t. absorption $\succeq$) is finite. Further, while there are infinitely descending chains of monomials, such as $1 = x^0 \succ x^1 \succ x^2 \succ \dots$, no such chain is infinitely ascending.*

To see this, consider the set of k-tuples $(\mathbb{N}^\infty)^k$ with the component-wise standard order on $\mathbb{N}^\infty$. This is a well-partial order and therefore has no infinite descending chains and no infinite antichains. The set of monomials is isomorphic to $(\mathbb{N}^\infty)^k$ with $k = |X|$, but with the order reversed. It follows that all ascending chains and antichains of monomials are finite.

Definition 3.5.3 (Absorptive polynomials)**.** The semiring $(\mathbb{S}^\infty(X), +, \cdot, 0, 1)$ consists of all antichains of monomials in $\mathsf{Mon}_X(\mathbb{N}^\infty)$. We write such antichains as formal sums of their monomials and call them *absorptive polynomials*. Addition and multiplication of polynomials proceed as usual, but keeping only the maximal monomials (w.r.t. $\succeq$) in the result (and disregarding coefficients); 0 and 1 correspond to the antichains $\emptyset$ and $\{1\}$, respectively.

For example, $(xy^2 + x^2y) \cdot x^\infty = x^\infty y^2 + x^\infty y = x^\infty y$. Since antichains of monomials are finite, there is no difference between polynomials and power series here, and $\mathbb{S}^\infty(X)$ is of countable cardinality. Regarding notation, we usually denote polynomials by the letters P, Q and monomials by m, v. When working with sets $S \subseteq \mathbb{S}^\infty(X)$, we may switch to lowercase letters p, q to avoid confusion with S.

The natural order on $\mathbb{S}^\infty(X)$ can be characterised by monomial absorption: $P \leq Q$ if, and only if, for each $m \in P$ there is $m' \in Q$ with $m' \succeq m$. With Lemma 3.5.2, it follows that there are no infinitely ascending chains of polynomials, and further that the supremum of $S \subseteq \mathbb{S}^\infty(X)$ is $\bigsqcup S = \mathsf{Maximals}(\bigcup S)$, i.e., the set of $\succeq$-maximal monomials occurring anywhere in S. We thus know that the natural order is a complete lattice. One can further show that it is completely distributive [Mrk20] (see Corollary 3.6.8 for a more general result), and that $\mathbb{S}^\infty(X)$ is fully continuous [DGNT21]. Due to the exponent ∞ and the finite number of indeterminates, there is a smallest monomial $m_\infty \neq 0$ with $m_\infty(x) = \infty$ for all $x \in X$. This ensures chain-positivity of $\mathbb{S}^\infty(X)$.

Proposition 3.5.4. $(\mathbb{S}^\infty(X), +, \cdot, 0, 1)$ *is absorptive, fully continuous, and chain-positive. The natural order is a completely distributive lattice.*

Recall that addition and multiplication commute with arbitrary sets (not just chains) by Corollary 3.2.5. The same does not hold for infima:

Example 3.5.5. Let $p = x + y \in \mathbb{S}^\infty(X)$ and $S = \{x, y\} \subseteq \mathbb{S}^\infty(X)$. Then

$$p \cdot \bigsqcap S = p \cdot xy = x^2y + xy^2 \quad < \quad xy = \bigsqcap\{x^2 + xy, y^2 + xy\} = \bigsqcap(p \cdot S).$$

Following Section 3.3, we can equip $\mathbb{S}^\infty(X)$ with the infinitary power operation. Due to the properties shown in Lemma 3.3.2, taking the infinitary power of a polynomial amounts to replacing all (nonzero) exponents by ∞. This can lead to additional absorption, e.g.,

$(x^2y + y^3 + z^\infty)^\infty = x^\infty y^\infty + y^\infty + z^\infty = y^\infty + z^\infty$. A distinctive feature of $\mathbb{S}^\infty(X)$ is that the infinitary power commutes with arbitrary suprema (cf. Lemma 3.3.3), as suprema can be described by $\bigsqcup S = \mathsf{Maximals}(\bigcup S)$ due to the finiteness properties.

Lemma 3.5.6. *Let $S \subseteq \mathbb{S}^\infty(X)$. Then $\left(\bigsqcup S\right)^\infty = \bigsqcup S^\infty$, where $S^\infty = \{p^\infty \mid p \in S\}$.*

3.5.2 Universality

The central property of $\mathbb{S}^\infty(X)$ is the following universal property which says that it is the absorptive, fully-continuous semiring freely generated by X (w.r.t. fully-continuous homomorphisms). We refer to the more general Theorem 3.7.13 for a proof. The main difficulty is to establish the continuity requirement for infima of chains, for which Kőnig's lemma is used.

Theorem 3.5.7 (Universality)**.** *Every mapping $h\colon X \to K$ into an absorptive, fully continuous semiring K uniquely extends to a fully-continuous semiring homomorphism $h\colon \mathbb{S}^\infty(X) \to K$.*

The fact that the universal property guarantees fully-continuous homomorphisms should not be taken lightly: we have seen in Example 3.2.9 that this is not the case for formal power series $\mathbb{N}^\infty[\![X]\!]$. Kőnig's lemma can also be applied to describe infima of polynomials in terms of monomial chains, which comes in handy when we prove infinite distributivity.

Proposition 3.5.8 (Characterisation of Infima)**.** *Let $(P_i)_{i<\omega}$ be a descending ω-chain in $\mathbb{S}^\infty(X)$. Let further $\mathcal{M}$ be the set of descending ω-chains $(m_i)_{i<\omega}$ of monomials with the property that $m_i \in P_i$ for all i. Then,*

$$\bigsqcap_{i<\omega} P_i = \bigsqcup \Big\{ \bigsqcap_{i<\omega} m_i \mid (m_i)_{i<\omega} \in \mathcal{M} \Big\}.$$

Proof [DGNT21]. Direction "$\geq$" is easy, as $m_i \leq P_i$ implies $\bigsqcap_i m_i \leq \bigsqcap_i P_i$ for every chain $(m_i)_{i<\omega} \in \mathcal{M}$. For the other direction, let $P_\omega = \bigsqcap_i P_i$. To prove "$\leq$", it suffices to show that for every monomial $m_\omega \in P_\omega$, there is a monomial chain $(m_i)_{i<\omega} \in \mathcal{M}$ with $\bigsqcap_i m_i \succeq m_\omega$.

We use a similar argument as in the proof of the universal property. Fix a monomial $m_\omega \in P_\omega$ and, for the moment, an $i < \omega$. We have $P_\omega \leq P_i$, so there is a monomial $m_i \in P_i$ with $m_\omega \preceq m_i$. As $P_i \leq P_{i-1} \leq \cdots \leq P_0$, we obtain a sequence $m_i \preceq m_{i-1} \preceq \cdots \preceq m_0$ of monomials with $m_\omega \preceq m_j$ and $m_j \in P_j$ (for all $j \leq i$).

Repeating this construction for each i produces sequences of unbounded finite length, so by Kőnig's lemma (recall that all polynomials are finite), there must be an infinite monomial chain $(m_i)_{i<\omega}$ such that $m_i \in P_i$ and $m_i \succeq m_\omega$ for all i. Hence $(m_i)_{i<\omega} \in \mathcal{M}$ with infimum $\bigsqcap_{i<\omega} m_i \succeq m_\omega$ as claimed. □

3.5.3 Strong Distributivity

Infinite summations and products can be defined in $\mathbb{S}^\infty(X)$ just as in any other absorptive, fully-continuous semiring, by taking the supremum or infimum over all finite subsets (see Section 3.4.2). In addition to the properties that hold in all such semirings, $\mathbb{S}^\infty(X)$ satisfies the strong infinite distributivity law (property (**D**) in Section 3.4). The proof requires some preparation; we first describe products in a similar way as infima.

Theorem 3.5.9. *Let $(P_i)_{i<\omega}$ be a countable family of polynomials in $\mathbb{S}^\infty(X)$. Then,*

$$\widehat{\prod_{i<\omega}} P_i = \bigsqcup \Big\{ \widehat{m}_f \;\Big|\; f \colon \omega \to \bigcup_i P_i \text{ such that } f(i) \in P_i \text{ for all } i \Big\},$$

where the monomial $\widehat{m}_f$ is the infinite product over all $f(i)$. That is, for each $x \in X$,

$$\widehat{m}_f(x) = \widehat{\sum_{i<\omega}} f(i)(x).$$

Proof. Using Proposition 3.5.8, we can rewrite the product as follows:

$$\widehat{\prod_{i<\omega}} P_i = \bigsqcap_{n<\omega} P_0 \cdots P_n = \bigsqcup \Big\{ \bigsqcap_{i<\omega} m_i \;\Big|\; (m_i)_{i<\omega} \text{ is a monomial chain and } m_i \in P_0 \cdots P_i \Big\}.$$

Let $v_j^{(i)} \in P_j$ be the monomial that m_i chooses from P_j, for $j \leq i$.

We first show the easier direction "$\geq$" of the theorem. Consider a monomial $\widehat{m}_f$. We define a chain $(m_i)_{i<\omega}$ by setting $m_i \coloneqq f(0) \cdots f(i)$. This is clearly a chain and we have $\bigsqcap_i m_i = m_f$ by comparing exponents. Hence $\widehat{\prod}_i P_i \geq \bigsqcap_i m_i = \widehat{m}_f$ for each f.

For the other direction, consider a monomial chain $(m_i)_{i<\omega}$. Recall that each polynomial P_i is finite. Hence some monomial in P_0 must be chosen by infinitely many m_i. We set $f(0)$ to be such a monomial, so that $f(0) = v_0^{(i)}$ for infinitely many i. For $f(1)$, we consider only those i with $f(0) = v_0^{(i)}$ and choose $f(1)$ in the same way, so that there are infinitely many i such that $f(0) = v_0^{(i)}$ and $f(1) = v_1^{(i)}$. Proceed in the same way to define $f(i)$ for all $i < \omega$. We claim that $\bigsqcap_i m_i \leq \widehat{m}_f$ which ends the proof.

We prove the claim by comparing exponents, so let $x \in X$. First consider the case that $\widehat{m}_f(x) = k$ is finite. Then there is a j such that the the exponent is reached, i.e., $(f(0) \cdots f(j))(x) = k$. Choose an i_* such that $f(0), \ldots, f(j)$ are chosen by m_{i_*} (there are infinitely many such i_* by construction). Then $\widehat{m}_f(x) = k \leq m_{i_*}(x) \leq \bigsqcup_i m_i(x)$.

Now consider the case $\widehat{m}_f(x) = \infty$. Then there are indices $j_0, j_1, j_2, \ldots$ such that $(f(0) \cdots f(j_k))(x) \geq k$, for each k. Now for each k, proceed as in the previous case by picking an index i_k such that $f(0), \ldots, f(j_k)$ are chosen by m_{i_k}. Then $\bigsqcup_i m_i(x) \geq m_{i_k}(x) \geq k$. As this holds for each k, we have $\bigsqcup_i m_i(x) = \infty$. □

We are now ready to prove strong distributivity. Notice that we still assume X to be a *finite* set of indeterminates, so that the result of an infinite summation or product must always be a finite polynomial.

Corollary 3.5.10. $\mathbb{S}^\infty(X)$ *satisfies strong distributivity.*

Proof. Recall that infinite summation is the same as supremum. We want to prove

$$\widehat{\prod_{j\in J}} \widehat{\sum_{i\in I_j}} P_i = \widehat{\sum_{f\in\mathcal{F}}} \widehat{\prod_{j\in J}} P_{f(j)},$$

where $\mathcal{F}$ is the set of choice functions $f\colon J \to \bigcup_j I_j$ with $f(j) \in I_j$ for all j. We first rewrite both sides using Theorem 3.5.9:

$$\widehat{\prod_{j\in J}}\left(\widehat{\sum_{i\in I_j}} P_i\right) = \bigsqcup\left\{m_h \;\middle|\; \text{for all } j\colon\ h(j) \in P_i \text{ for some } i \in I_j\right\} =: A,$$

(We remark that a function h may choose monomials $h(j)$ that are absorbed in the summation $\widehat{\sum}_{i\in I_j} P_i$. This is not a problem, as the corresponding monomial m_h is then similarly absorbed by the monomial m_g for a different function g that chooses the absorbing monomial.)

$$\widehat{\sum_{f\in F}}\widehat{\prod_{j\in J}} P_{f(j)} = \bigsqcup\left\{m_{f,g} \;\middle|\; f \in F, \text{ for all } j\colon\ g(j) \in P_{f(j)}\right\} =: B.$$

The monomial $m_{f,g}$ is defined in the same way as m_g, i.e., $m_{f,g}(x) = \widehat{\sum}_j g(j)(x)$. We write $m_{f,g}$ to emphasise that g depends on f, and to distinguish the monomials of A and B.

It is now easy to show that $A = B$, which proves the claim. First consider $m_h \in A$. We define f, g so that $m_h = m_{f,g} \in B$. For each $j \in J$, set $f(j) = i$ with i so that $h(j) \in P_i$. Further set $g(j) = h(j)$. Then $g(j) \in P_{f(j)}$ holds by construction, thus $m_{f,g} \in B$. Conversely, let $m_{f,g} \in B$. We obtain h so that $m_{f,g} = m_h \in A$ simply by setting $h = g$, since $h(j) = g(j) \in P_i$ for some i (namely for $i = f(j)$). $\square$

This result generalises to any absorptive, fully-continuous semiring by means of the universal property. However, this only works under the restriction that the families $(a_i)_{i\in I_j}$ for $j \in J$ assume only finitely many values, which arguably excludes the more interesting examples. The reason is that we can then choose X large enough to represent each assumed value by a different indeterminate, apply strong distributivity in $\mathbb{S}^\infty(X)$, and then instantiate the indeterminates by the original values. This instantiation is fully continuous by the universal property and preserves infinite sums and products by Proposition 3.4.3.

Corollary 3.5.11. *Let K be an absorptive, fully-continuous semiring. If the set of values $\{a_i \mid i \in I_j, j \in J\}$, is finite, then strong distributivity holds, i.e.,*

$$\widehat{\prod_{j\in J}}\widehat{\sum_{i\in I_j}} a_i = \widehat{\sum_{f\in\mathcal{F}}}\widehat{\prod_{j\in J}} a_{f(j)}.$$

3.6 Absorptive Polynomials with More Indeterminates

In this section, we discuss an extension of generalised absorptive polynomials $\mathbb{S}^\infty(X)$ to infinitely many indeterminates, which we denote as $\mathbb{S}^\infty(\!(\boldsymbol{\mathcal{X}})\!)$. One motivation for this is semiring semantics for fixed-point logics on infinite structures, where we might want to track infinitely many atomic facts with pairwise different indeterminates. Moreover, we have just seen in Corollary 3.5.11 how strong distributivity of $\mathbb{S}^\infty(X)$ can be generalised to all absorptive, fully-continuous semirings – assuming we sum and multiply over a finite range. A version of $\mathbb{S}^\infty(X)$ with infinitely many indeterminates could, in principle, be used to lift the range restriction and prove strong distributivity in general, assuming that it has a similar universal property as $\mathbb{S}^\infty(X)$. We already know that this reasoning cannot go through, as the Viterbi semiring does not have strong distributivity (cf. Section 3.4.2), but it may be worthwhile to see how far we can get on this path.

We show that the semiring $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ can be defined by a quotient construction, similar to the original definition of $\mathbb{S}(X)$ as a quotient of $\mathbb{N}[X]$. The main observation is that the resulting semiring still satisfies some, but not all continuity requirements, and is thus no longer fully continuous. It does, however, have universality properties for suitably defined infinitary operations.

3.6.1 Quotient Construction

Let $\boldsymbol{\mathcal{X}} = \{x_0, x_1, x_2, \dots\}$ be a countably infinite set of indeterminates. Monomials $\mathsf{Mon} = \mathsf{Mon}_{\boldsymbol{\mathcal{X}}}(\mathbb{N}^\infty)$ are defined as usual, as mappings $m\colon \boldsymbol{\mathcal{X}} \to \mathbb{N}^\infty$. We use the convenient notation $x_{[i,j]} = x_i x_{i+1} \cdots x_j$ and $x_{[i,\infty)} = x_i x_{i+1} \cdots$, where the latter denotes an infinite monomial.

We define $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ as a quotient of $\mathrm{Poly}(\mathbb{B}, \mathbb{N}^\infty, \boldsymbol{\mathcal{X}})$ by absorption, but we use a more convenient representation of its elements as sets of monomials (as there are no coefficients). That is, we define $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ as quotient of the *monomial set semiring* $(\mathcal{P}(\mathsf{Mon}), +, \cdot, 0, 1)$, with addition $P + Q = P \cup Q$, multiplication $P \cdot Q = \{m_1 \cdot m_2 \mid m_1 \in P, m_2 \in Q\}$ for $P, Q \subseteq \mathsf{Mon}$, and neutral elements $0 = \emptyset$ and $1 = \{1\}$.

Absorption $\succeq$ on monomials is defined as usual by comparing exponents. We extend this to a congruence on monomial sets as follows.

Definition 3.6.1 (Absorption equivalence)**.** Let $P, Q \subseteq \mathsf{Mon}$. We say that P is absorbed by Q, denoted $P \preceq Q$, if for each $m \in P$ there is $m' \in Q$ with $m \preceq m'$. Further, *absorption equivalence* is the equivalence relation defined by

$$P \sim Q \quad \iff \quad P \preceq Q \text{ and } P \succeq Q.$$

We write $P \prec Q$ if $P \preceq Q$ and $P \nsim Q$.

Notice that absorption $P \preceq Q$ is the same as the natural order $P \leq Q$ of $\mathbb{S}^\infty(X)$ in the finite-indeterminate setting. Here we need to define the order first, before constructing $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$. In the end, the natural order of the quotient semiring will coincide with $\preceq$.

Intuitively, $P \sim Q$ holds if P and Q are identical up to monomials that are absorbed, e.g. $x^2 + x^2 y \sim x^2$. In general, $P \sim \mathsf{Maximals}(P)$ whenever P is finite, but maximal monomials need not exist if P is infinite, e.g. $P = \{x_{[i,\infty)} \mid i < \omega\}$. It is straightforward to verify that $\sim$ is a congruence so that the quotient construction is well-defined.

Definition 3.6.2 ($\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ as quotient)**.** We define $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!) = \mathcal{P}(\mathsf{Mon})/_\sim$, whose elements are equivalence classes $[P] = \{Q \subseteq \mathsf{Mon} \mid P \sim Q\}$ called *absorptive polynomials (over $\boldsymbol{\mathcal{X}}$).*

If $\boldsymbol{\mathcal{X}}$ is finite, this construction is isomorphic to the definition of $\mathbb{S}^\infty(X)$, by identifying each class $[P]$ with its minimal (w.r.t. set inclusion) representative $\mathsf{Maximals}(P)$. In the infinite case, we can no longer identify polynomials with antichains. More precisely, we say that $[P]$ is an *antichain* if it has a representative $Q \sim P$ that is an antichain. We observe that the antichain representative, if it exists, is always given by $\mathsf{Maximals}(P)$. Examples of non-antichain polynomials are discussed in the next section.

Lemma 3.6.3. *Let* $[P] \in \mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$. *Then* $[P]$ *is an antichain if, and only if,* $P \sim \mathsf{Maximals}(P)$.

Proof. The direction "⇐" is trivial, as $\mathsf{Maximals}(P)$ is always an antichain. For "⇒", assume $Q \sim P$ is an antichain, so $Q = \mathsf{Maximals}(Q)$. Trivially $\mathsf{Maximals}(P) \preceq P$, so it remains to show $P \preceq \mathsf{Maximals}(P)$. We first claim that $Q \subseteq P$. To see this, let $m \in Q$. Since $P \sim Q$, there are $m' \in P$, $m'' \in Q$ such that $m \preceq m' \preceq m''$. But Q is an antichain, hence $m = m' = m'' \in P$.

Now let $m \in P$. There is $m' \in Q$ with $m \preceq m'$, and by the claim $m' \in P$. We show that m' is maximal. Assume we have $m' \preceq m'' \in P$. Then there is $m''' \in Q$ with $m' \preceq m'' \preceq m'''$. But Q is an antichain still, so $m' = m'' = m'''$ and we indeed have $m' \in \mathsf{Maximals}(P)$. □

We consider the essential properties of the constructed semiring. The neutral elements are $0 = [\emptyset]$ and $1 = [\{1\}] = [\mathsf{Mon}]$. By construction of $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ as quotient of $\mathcal{P}(\mathsf{Mon})$, addition is idempotent. The natural order can thus be defined as $[P] \leq [Q]$ if $[P] + [Q] = [Q]$ (cf. Proposition 3.1.6). As expected, this is the same as $\preceq$ and $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ is absorptive.

Lemma 3.6.4. *The natural order of* $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ *coincides with absorption* $\preceq$.

Proof. We show that $[P] + [Q] = [Q]$ iff $P \preceq Q$. Recall that $[P] + [Q] = [Q]$ is defined as $(P \cup Q) \sim Q$. Now if $P \preceq Q$, then also $(P \cup Q) \preceq Q$ and trivially $(P \cup Q) \succeq Q$ by definition of absorption, hence $[P] + [Q] = [Q]$. Conversely, $(P \cup Q) \sim Q$ implies $P \preceq (P \cup Q) \preceq Q$. □

Corollary 3.6.5. $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ *is an absorptive semiring.*

Proof. By Proposition 3.1.7, as $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ is naturally ordered and $P \preceq 1$ for all $P \subseteq \mathsf{Mon}$. □

How is completeness of the natural order affected? Recall that in the finite setting, $\bigsqcup S = \mathsf{Maximals}(\bigcup S)$ due to the finiteness of ascending chains. We show that suprema are preserved by the quotient construction. Since the supremum in the $\mathcal{P}(\mathsf{Mon})$-semiring is simply union, this leads to the same characterisation also in the infinite setting.

Proposition 3.6.6. *Let* $S \subseteq \mathcal{P}(\mathsf{Mon})$ *be a set of monomial sets and* $[S] := \{[P] \mid P \in S\}$. *The supremum of* $[S]$ *is given by*

$$\bigsqcup [S] = \left[\bigcup S\right].$$

In particular, the natural order of $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ *is a complete lattice.*

Proof. For each $P \in S$, we have $P \subseteq \bigcup S$ and hence $[\bigcup S]$ is an upper bound. It remains to prove minimality. Let $[Q]$ be any upper bound for $[S]$, so $P \preceq Q$ for all $P \in S$. Consider a monomial $m \in \bigcup S$. Then $m \in P$ for some $P \in S$, so there is $m' \in Q$ with $m \preceq m'$. Applying this argument to all monomials implies $\bigcup S \preceq Q$, so $[\bigcup S] \leq [Q]$. □

Infima are not preserved. For example, $\bigsqcap\{x^n \mid n < \omega\} = \emptyset$ in the $\mathcal{P}(\mathsf{Mon})$-semiring, but $\bigsqcap\{[x^n] \mid n < \omega\} = [x^\infty]$ since $[x^\infty]$ and $[x^n]$ become comparable in $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$. We can explicitly describe infima via choice functions as follows.

Proposition 3.6.7. *Let $S \subseteq \mathcal{P}(\mathsf{Mon})$ be a set of monomial sets and $[S] := \{[P] \mid P \in S\}$. The infimum of $[S]$ is*

$$\bigsqcap[S] = \big[\{m_f \mid f\colon S \to \bigcup S \text{ with } f(P) \in P \text{ for all } P \in S\}\big],$$

where the monomial m_f is defined by

$$m_f(x) = \bigsqcup_{P \in S} f(P)(x), \quad \text{for all } x \in \boldsymbol{\mathcal{X}}.$$

Proof. We first prove that $\bigsqcap[S] = \big[\{m \in \mathsf{Mon} \mid m \preceq P \text{ for all } P \in S\}\big]$ by showing absorption in both directions. Recall that we can express the infimum as the supremum over all lower bounds: $\bigsqcap[S] = \bigsqcup\big\{[Q] \bigm| [Q] \preceq [P] \text{ for all } P \in S\big\}$. The set $R = \{m \mid m \preceq P, P \in S\}$ is one possible lower bound, so direction $\geq$ is trivial. For the other direction, consider any lower bound $[Q]$. By definition of absorption, each $m \in Q$ must satisfy $m \preceq P$ for all $P \in S$. Hence $Q \subseteq R$ and thus $[Q] \leq [R]$ and direction $\leq$ follows.

We prove the proposition by comparing elements. Fix a choice function f with $f(P) \in P$ for $P \in S$. By comparing exponents, it is clear that $m_f \preceq f(P) \preceq P$ for all $P \in S$, and hence $m_f \in R$. Conversely, fix $m \in R$. For each $P \in S$, we have $m' \in P$ with $m \preceq m'$, so we can choose f with $f(P) = m'$. Then $m \preceq m_f$, again by comparing exponents. □

In the finite-indeterminate setting, we can describe the infimum $\bigsqcap[S]$ by infima of monomial chains (cf. Proposition 3.5.8). Example 3.6.10 below shows that this is no longer possible. The reason is that our proof relied on Kőnig's lemma, which is not applicable for infinite polynomials. However, the above characterisation suffices to prove that the natural order remains completely distributive.

Corollary 3.6.8. *The natural order of $\mathbb{S}^\infty⟅\boldsymbol{\mathcal{X}}⟆$ is a completely distributive lattice.*

Proof. To simplify notation, we do not distinguish between polynomials (i.e., equivalence classes) and monomial sets. Let $(S_i)_{i \in I}$ be a family of sets $S_i \subseteq \mathbb{S}^\infty⟅\boldsymbol{\mathcal{X}}⟆$. We need to prove:

$$\bigsqcap_{i \in I} \bigsqcup S_i \stackrel{!}{=} \bigsqcup_{h \in \mathcal{H}} \bigsqcap_{i \in I} h(i),$$

where $\mathcal{H}$ is the set of choice functions $h\colon I \to \bigcup_i S_i$ with $h(i) \in S_i$ for all i. Using the above characterisations of suprema (1) and infima (2),

$$\begin{aligned}\bigsqcap_{i \in I} \bigsqcup S_i \stackrel{(1)}{=} \bigsqcap_{i \in I} \bigcup S_i &\stackrel{(2)}{=} \{m_f \mid f\colon I \to \bigcup_{i \in I} \bigcup S_i,\ f(i) \in \bigcup S_i \text{ for all } i\} \\ &= \{m_f \mid f(i) \in P_i \text{ for some } P_i \in S_i,\ \text{for all } i\} \\ &= \bigcup_{h \in \mathcal{H}} \{m_f \mid f(i) \in h(i),\ \text{for all } i\} \stackrel{(2,1)}{=} \bigsqcup_{h \in \mathcal{H}} \bigsqcap_{i \in I} h(i)\end{aligned}$$

□

These results may serve as a sanity check that the quotient construction of $\mathbb{S}^\infty⟅\boldsymbol{\mathcal{X}}⟆$ leads to reasonable properties for infinite $\boldsymbol{\mathcal{X}}$. To simplify the presentation, we drop the explicit notation $[P]$ from now on and no longer distinguish between equivalence classes $[P]$ and monomial sets $P \subseteq \mathsf{Mon}$; this is not an issue, as the choice of the representative P does not matter. Consequently, we use $\leq$ to compare polynomials and reserve $\preceq$ for monomials.

3.6.2 Examples

The first example shows that, unlike in the finite-indeterminate setting, ascending chains need not be finite and polynomials need not be antichains.

Example 3.6.9 (Ascending chain)**.** The polynomials $P_i = x_{[i,\infty)}$ form the chain

$$x_{[0,\infty)} \prec x_{[1,\infty)} \prec x_{[2,\infty)} \prec x_{[3,\infty)} \prec \dots$$

which is infinite and has no maximal monomials, so $\mathsf{Maximals}(\{x_{[i,\infty)} \mid i < \omega\}) = \emptyset$. The supremum is given by $\bigsqcup_i P_i = \{x_{[i,\infty)} \mid i < \omega\}$. This representative is clearly not an antichain, and in fact there is no antichain representative (by Lemma 3.6.3).

Even if a polynomial is an antichain, it can obviously be infinite – in fact, even uncountable. To see this, label the nodes of an infinite binary tree with $\boldsymbol{\mathcal{X}}$ (recall that $\boldsymbol{\mathcal{X}}$ is countable). The monomials corresponding to the infinite paths from the root form an uncountable antichain.

We next consider two examples of decreasing chains.

Example 3.6.10 (Decreasing chain I)**.** Let $P_i = \sum_{j \geq i} x_j$, so that

$$\begin{aligned} P_0 &= x_0 + x_1 + x_2 + x_3 + x_4 + x_5 + \dots \\ P_1 &= \quad\;\; x_1 + x_2 + x_3 + x_4 + x_5 + \dots \\ P_2 &= \qquad\qquad x_2 + x_3 + x_4 + x_5 + \dots \\ P_3 &= \qquad\qquad\qquad x_3 + x_4 + x_5 + \dots \\ \dots & \qquad\qquad\qquad\qquad \dots \end{aligned}$$

is an infinite descending chain. As each monomial x_j is removed at some point, the infimum cannot contain the monomial x_j. Instead, $x_{[0,\infty)}$ is a lower bound, and so is $x_{[1,\infty)}$ or $x_{[42,\infty)}$. It is not hard to see that a single monomial m is a lower bound precisely if m contains indeterminates of unbounded index (if the index is bounded, then for some P_i all indeterminates have been removed, and hence m is no lower bound for P_i). We thus have

$$\bigsqcap_{i<\omega} P_i = \Big\{ \prod_{j \in J} x_i \mid J \subseteq \omega \text{ is infinite} \Big\}.$$

There is no maximal such monomial, so there is again no antichain representative.

We further note that there are arbitrarily long monomial chains through $(P_i)_{i<\omega}$: simply choose the monomial x_j from each P_i with $i < j$. But unlike in the proof of Proposition 3.5.8, we cannot apply Kőnig's lemma as the P_i are infinite. And indeed, there is no infinite monomial chain, as each monomial disappears after finitely many steps.

Example 3.6.11 (Decreasing chain II)**.** Let $P = \{x_i \mid i < \omega\}$ and consider the powers $(P^n)_{n<\omega}$ which form a decreasing chain. The infimum is given by

$$\bigsqcap_{n<\omega} P^n = \{x_i^\infty \mid i < \omega\} \cup \left\{ \prod_{i\in I} x_i \mid I \subseteq \omega \text{ is infinite} \right\}.$$

To see why this is the infimum, we apply Proposition 3.6.7:

- By setting f to choose $x_i^n \in P^n$ for each n, we get $m_f = x_i^\infty$. For an infinite $I \subseteq \omega$, we choose the monomial from P^n that consists of the first n indeterminates mentioned in I, so that the resulting monomial m_f is the product of all indeterminates in I.
- Conversely, each monomial m_f contains either an indeterminate with exponent ∞ or infinitely many different indeterminates. So it is absorbed by a monomial of the form x_i^∞ or $\prod_{i\in I} x_i$.

3.6.3 Continuity

We show that $\mathbb{S}^\infty⦇\boldsymbol{\mathcal{X}}⦈$ satisfies exactly three of the four requirements for full continuity. This suggests that there is no universal property similar to the finite-indeterminate case – and indeed, we turn Example 3.6.10 into a counterexample.

Theorem 3.6.12 (3/4-continuous)**.** *Let $S \subseteq \mathbb{S}^\infty⦇\boldsymbol{\mathcal{X}}⦈$ be a nonempty set of polynomials and let $p \in \mathbb{S}^\infty⦇\boldsymbol{\mathcal{X}}⦈$ be a single polynomial. Then the following equations hold:*

$$p \cdot \bigsqcup S = \bigsqcup (p \cdot S), \quad p + \bigsqcup S = \bigsqcup (p + S), \quad p + \bigsqcap S = \bigsqcap (p + S).$$

Proof. We prove the statements from left to right.

- The first equation follows from Proposition 3.6.6, since

$$p \cdot \bigcup S = \{m \cdot v \mid m \in p,\ v \in q \text{ for some } q \in S\} = \bigcup (p \cdot S).$$

- Both $+$ and $\bigsqcup$ are simply unions, making the second equation trivial (if S is nonempty!).
- For the last equation, observe that $p + \bigsqcap S \leq \bigsqcap(p + S)$ by monotonicity of addition. For the other direction, consider a monomial $m \in \bigsqcap(p + S)$. If $m \preceq p$, then also $m \leq p + \bigsqcap S$ and we are done. Otherwise, we have $m \leq p + q$ for all $q \in S$. As m is a single monomial and $m \not\preceq p$, we must then have $m \preceq q$ for all $q \in S$, and thus $m \leq \bigsqcap S \leq p + \bigsqcap S$. □

It remains to consider the case $p \cdot \bigsqcap S \stackrel{?}{=} \bigsqcap(p \cdot S)$. We know from the finite-indeterminate setting that we must assume S to be a chain (cf. Example 3.5.5). The key insight for the following counterexample is that for each $q \in S$ we can choose a monomial $m \in p$ in a clever way so that the corresponding choice monomial m_f for the right-hand side is at least as large as the infimum over S. As m_f is contained in $\bigsqcap(p \cdot S)$, this infimum is larger than $p \cdot \bigsqcap S$ due to the additional multiplication by p in the latter expression.

Example 3.6.13. Let $S = \{x_{[1,i]} \mid 1 \le i < \omega\}$ and $p = \{x_i \mid 1 \le i < \omega\}$. Then $\bigsqcap S = x_{[1,\infty)}$ and we have

$$p \cdot \bigsqcap S = p \cdot x_{[1,\infty)} \quad < \quad x_{[1,\infty)} = \bigsqcap (p \cdot S)$$

To see why the inequality holds, observe that every monomial in $p \cdot x_{[1,\infty)}$ contains an indeterminate with exponent 2, and hence $x_{[1,\infty)}$ is not absorbed. To see why $x_{[1,\infty)}$ is the infimum of $p \cdot S$, observe that for each $x_{[1,i]} \in S$, we can choose $x_{i+1} \in p$ and obtain the monomial $x_{[1,i+1]}$. The corresponding choice monomial is then $x_{[1,\infty)}$ and this is the largest choice monomial we can get.

Since $\mathbb{S}^\infty(\!\!|\boldsymbol{\mathcal{X}}|\!\!)$ is not fully continuous itself, we would expect that $\mathbb{S}^\infty(\!\!|\boldsymbol{\mathcal{X}}|\!\!)$ is no longer universal (cf. Theorem 3.5.7). And indeed, this is witnessed by Example 3.6.10.

More precisely, let $h \colon \boldsymbol{\mathcal{X}} \to K$ be an instantiation of the indeterminates with values in an absorptive, fully-continuous semiring K. Recall that we can equip K with the infinitary power and infinitary sums and products (Sections 3.3 and 3.4.2), which allow us to naturally extend h to $\mathbb{S}^\infty(\!\!|\boldsymbol{\mathcal{X}}|\!\!)$. We first show that this defines a semiring homomorphism.

Lemma 3.6.14. *Let $h \colon \boldsymbol{\mathcal{X}} \to K$ with K as above. We obtain a semiring homomorphism $h \colon \mathbb{S}^\infty(\!\!|\boldsymbol{\mathcal{X}}|\!\!) \to K$ by extending h to polynomials $P \in \mathbb{S}^\infty(\!\!|\boldsymbol{\mathcal{X}}|\!\!)$ as follows:*

$$h(P) = \widehat{\sum_{m \in P}} h(m), \qquad h(m) = \widehat{\prod_{x \in \boldsymbol{\mathcal{X}}}} h(x)^{m(x)}.$$

Proof. We remark that h is well-defined, since $m \preceq m'$ implies $h(m) \le h(m')$ for monomials (by monotonicity **M**), and thus $P \sim Q$ implies $h(P) = h(Q)$ for monomial sets $P, Q \subseteq \mathsf{Mon}$. It is further clear that $h(0) = 0$ and $h(1) = 1$. Now let $P, Q \in \mathbb{S}^\infty(\!\!|\boldsymbol{\mathcal{X}}|\!\!)$. Recall that infinite summation in K is defined as supremum, so we have

$$h(P) + h(Q) = \bigsqcup_{m \in P} h(m) + \bigsqcup_{m \in Q} h(m) = \bigsqcup_{m \in P \cup Q} h(m) = h(P + Q).$$

For multiplication, we first show that h commutes with multiplication of monomials due to associativity **A**. That is, for $m_1, m_2 \in \mathsf{Mon}$,

$$\begin{aligned} h(m_1) \cdot h(m_2) &= \widehat{\prod_{x \in \boldsymbol{\mathcal{X}}}} h(x)^{m_1(x)} \cdot \widehat{\prod_{x \in \boldsymbol{\mathcal{X}}}} h(x)^{m_2(x)} \\ &= \widehat{\prod_{x \in \boldsymbol{\mathcal{X}}}} \Big(h(x)^{m_1(x)} \cdot h(x)^{m_2(x)} \Big) \\ &= \widehat{\prod_{x \in \boldsymbol{\mathcal{X}}}} h(x)^{m_1(x) + m_2(x)} = h(m_1 m_2). \end{aligned}$$

Recall that multiplication in K commutes with suprema (Corollary 3.2.5). Altogether,

$$h(P) \cdot (Q) = \bigsqcup_{m_1 \in P} h(m_1) \cdot \bigsqcup_{m_2 \in Q} h(m_2) = \bigsqcup_{\substack{m_1 \in P \\ m_2 \in Q}} h(m_1) \cdot h(m_2) = h(P \cdot Q). \qquad \square$$

However, these homomorphisms are not fully continuous in general, so $\mathbb{S}^\infty(\!(\boldsymbol{\mathcal{X}})\!)$ is not universal for absorptive, fully-continuous semirings.

Example 3.6.15. Recall $P_i = \{x_j \mid j \geq i\}$ from Example 3.6.10. Since each indeterminate disappears at some point, we have argued that the infimum is $\bigsqcap_i P_i = \{\prod_{j \in J} x_j \mid J \subseteq \omega \text{ is infinite}\}$.

Now let $0 < c < 1$ and consider the assignment $h(x_j) = c$ into the Viterbi semiring. Then $h(P_i) = \bigsqcup_{j \geq i} h(x_j) = c$ and $\bigsqcap_i h(P_i) = c$. On the other hand, $h(\prod_{j \in J} x_j) = \widehat{\prod}_{j \in J} h(x_j) = c^\infty = 0$ for each infinite $J \subseteq \omega$. Hence $h(\bigsqcap_i P_i) = 0$ and h is not fully continuous.

3.6.4 Infinitary Operations

Given that $\mathbb{S}^\infty(\!(\boldsymbol{\mathcal{X}})\!)$ is not fully continuous, the usual definitions of the infinitary power (Definition 3.3.1) and infinitary sums and products (Theorem 3.4.2) may no longer be suitable. We show that a reasonable notion of infinitary operations can instead be defined via choice functions. In particular, we show that polynomial evaluation into strongly distributive semirings induces infinitary homomorphisms.

We first look at the infinitary power $P^\infty \coloneqq \bigsqcap_n P^n$ of a polynomial $P \in \mathbb{S}^\infty(\!(\boldsymbol{\mathcal{X}})\!)$. We have already seen in Example 3.6.11 that we no longer have $P^\infty = \{m^\infty \mid m \in P\}$. A more elaborate counterexample shows that also $P \cdot P^\infty = P$ no longer holds, which was the initial motivation for the infinitary power.

Example 3.6.16. For each $I \subseteq \mathbb{N}$, we consider the monomial $x_I = \prod_{i \in I} x_i$. We say that x_I and x_J are disjoint if I and J are disjoint. For this example, we consider arithmetic progressions $I = p\mathbb{N} + r$, where p is a prime number and $r < p$. Notice that for primes $p \neq q$, we always have $(p\mathbb{N} + r) \cap (q\mathbb{N} + s) \neq \emptyset$ by the Chinese remainder theorem.

The intuition behind this choice is that for any k, choosing $p > k$ yields a collection $x_{p\mathbb{N}}$, $x_{p\mathbb{N}+1}$, $\dots$, $x_{p\mathbb{N}+k}$ of more than k pairwise disjoint monomials. On the other hand, once we fix a monomial $x_{p\mathbb{N}+r}$ we can extend this to a collection of at most p pairwise disjoint monomials, as we are forced to always use the same prime number p.

Now let $P = \{x_{p\mathbb{N}+r} \mid p \text{ is prime and } 0 \leq r < p\}$ and consider $P^\infty = \bigsqcap_n P^n$. We define the choice function f by selecting the following monomials from each P^n:

$$\begin{aligned}
&x_{2\mathbb{N}} && \in P \\
&x_{2\mathbb{N}} \cdot x_{2\mathbb{N}+1} && \in P^2 \\
&x_{3\mathbb{N}} \cdot x_{3\mathbb{N}+1} \cdot x_{3\mathbb{N}+2} && \in P^3 \\
&x_{5\mathbb{N}} \cdot x_{5\mathbb{N}+1} \cdot x_{5\mathbb{N}+2} \cdot x_{5\mathbb{N}+3} && \in P^4 \\
&x_{5\mathbb{N}} \cdot x_{5\mathbb{N}+1} \cdot x_{5\mathbb{N}+2} \cdot x_{5\mathbb{N}+3} \cdot x_{5\mathbb{N}+4} && \in P^5 \\
&\dots && \dots
\end{aligned}$$

That is, from each P^n we select a product of pairwise disjoint monomials by choosing the smallest prime $p \geq n$. This results in $m_f = x_\mathbb{N}$ and, by Proposition 3.6.7, $x_\mathbb{N} \leq P^\infty$.

Assuming towards a contradiction that $P \cdot P^\infty = P^\infty$, there must be some monomial $v \in P \cdot P^\infty$ such that $x_\mathbb{N} \preceq v$. By Proposition 3.6.7, this monomial is of the form $v = m \cdot m_g$, with $m = x_{p\mathbb{N}+r}$ and g a choice function with $g(n) \in P^n$. But now our intuition applies. For a prime $q > p$, the monomial $g(q) \in P^q$ is a product of q monomials from P. But then $g(q)$ cannot be disjoint from m, as there are at most $p - 1$ monomials disjoint from m in P. It follows that m and m_g are not disjoint, and hence v must have an exponent ≥ 2. But this contradicts $x_{\mathbb{N}} \preceq v$, so we must have $P \cdot P^\infty \neq P^\infty$.

Instead of the definition as infimum, we may define the infinitary power directly as $P^\infty := \{m^\infty \mid m \in P\}$ which also satisfies $P \cdot P^\infty = P^\infty$ (the proof is straightforward). Since the main application of the infinitary power is the fixed-point computation in Chapter 4 which is no longer applicable as it needs full continuity, we do not investigate this definition further and instead focus on the other infinitary operations.

Infinite summation can be defined simply as supremum, just as in the finite case. We define infinite products via choice functions to ensure distributivity.

Definition 3.6.17. Let $(P_i)_{i \in I}$ be a family of monomial sets with arbitrary index set I. The infinite operations are defined as

$$\widehat{\sum_{i \in I}} P_i := \bigcup_{i \in I} P_i, \qquad \widehat{\prod_{i \in I}} P_i := \Big\{ \widehat{m}_f \mid f \colon I \to \bigcup_{i \in I} P_i \text{ with } f(i) \in P_i \text{ for all } i \in I \Big\},$$

where the product monomial $\widehat{m}_f$ is defined as product over all $f(i)$, so for each $x \in \boldsymbol{\mathcal{X}}$,

$$\widehat{m}_f(x) = \sum_{i \in I} f(i)(x).$$

We write $\widehat{m}_f$ to avoid confusion with the infimum monomials m_f in Proposition 3.6.7.

It is easy to see that both operations are compatible with $\sim$. That is, if $P_i \sim Q_i$ for all $i \in I$, then also $\widehat{\sum}_i P_i \sim \widehat{\sum}_i Q_i$ and $\widehat{\prod}_i P_i \sim \widehat{\prod}_i Q_i$. We can thus lift them to the quotient $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$. Notice that summation coincides with supremum (cf. Proposition 3.6.6).

We note that $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ does not become an infinitary semiring under these operations, as the unique infinite powers property is violated. Indeed, the same counterexample $P = \{x_i \mid i < \omega\}$ with $P^\omega \neq P^{\omega_1}$ as for $\mathbb{N}^\infty_\infty[\![\boldsymbol{\mathcal{X}}]\!]$ applies (cf. Example 3.4.17 and Example 3.6.21 below). All other requirements of infinitary semirings are met:

Proposition 3.6.18. *The infinitary operations of Definition 3.6.17 satisfy strong associativity, strong distributivity, and respect neutral elements.*

Proof. It follows directly from the definition and the properties of $\mathbb{N}^\infty$ that neutral elements are respected. Strong associativity is trivial for summation (i.e., supremum). For products, we have to show that for any partition $(I_j)_{j \in J}$ of I,

$$\widehat{\prod_{i \in I}} P_i \;=\; \widehat{\prod_{j \in J}} \widehat{\prod_{i \in I_j}} P_i.$$

On the left-hand side, we get the monomials $\widehat{m}_f$ with $f(i) \in P_i$ for $i \in I$. That is, $\widehat{m}_f(x) = \sum_{i \in I} f(i)(x)$. For $\widehat{\prod}_{i \in I_j} P_i$, we get monomials $\widehat{m}_g$ with $g(i) \in P_i$ for $i \in I_j$. The outer product $\widehat{\prod}_{j \in J} \dots$ then consists of monomials $\widehat{m}_h$, where h chooses a monomial $\widehat{m}_g$ for each $j \in J$. To simplify notation, we identify $\widehat{m}_g$ and g, so that $h(j) = g$. The monomial $\widehat{m}_h$ is then of the form $\widehat{m}_h(x) = \sum_{j\in J}\sum_{i\in I} h(j)(i)(x)$.

It is now easy to show equality. For a monomial $\widehat{m}_h$, define f as $f(i) = h(j)(i)$ for the unique j with $i \in I_j$. Conversely, for each $\widehat{m}_f$ consider $g_j = f\restriction_{I_j}$ and set $h(j) = g_j$, for each $j \in J$. The monomials $\widehat{m}_f$ and $\widehat{m}_h$ are equal by comparing exponents, using strong associativity of summation in $\mathbb{N}^\infty$.

To prove strong distributivity, we first rewrite both sides by applying Definition 3.6.17. Recall from **D** that $\mathcal{F}$ is the set of choice functions f with $f(j) \in I_j$ for $j \in J$.

$$\widehat{\prod_{j\in J}}\Big(\widehat{\sum_{i\in I_j}} P_i\Big) = \Big\{\widehat{m}_{f'} \mid f'\colon J \to \bigcup_{\substack{j\in J,\\ i\in I_j}} P_i \text{ with } f'(j) \in P_i \text{ for some } i \in I_j\Big\} =: A$$

$$\widehat{\sum_{f\in\mathcal{F}}}\Big(\widehat{\prod_{j\in J}} P_{f(j)}\Big) = \Big\{\widehat{m}_{g_f} \mid f \in \mathcal{F},\ g_f\colon J \to \bigcup_{j\in J} P_{f(j)} \text{ with } g_f(j) \in P_{f(j)}\Big\} =: B$$

Let $\widehat{m}_{f'} \in A$. We define f, g so that $\widehat{m}_{f'} = \widehat{m}_{g_f} \in B$. For each $j \in J$, set $f(j) = i$ for some (not necessarily unique) i with $f'(j) \in P_i$. Further set $g(j) = f'(j)$. Then $g(j) \in P_{f(j)}$ holds and hence $\widehat{m}_{g_f} \in B$.

Conversely, let $\widehat{m}_{g_f} \in B$. We define f' so that $\widehat{m}_{g_f} = \widehat{m}_{f'} \in A$. To this end, we simply set $f' = g_f$ and observe that $f'(j) = g_f(j) \in P_i$ for some $i \in I_j$ (namely for $i = f(j)$). ☐

Despite the failure of the unique infinite powers property, we can prove that homomorphisms induced by polynomial evaluation commute with infinitary operations.

Theorem 3.6.19. *Let K be an absorptive, fully-continuous semiring with the usual infinitary operations, which additionally satisfy strong distributivity. Then every mapping $h\colon \boldsymbol{\mathcal{X}} \to K$ extends uniquely to a semiring homomorphism $h\colon \mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!) \to K$ that commutes with infinitary sums and products.*

Proof. We extend h to a semiring homomorphism by Lemma 3.6.14. To prove uniqueness, we rewrite polynomials as infinite sums and products. For x^∞, we have $h(x^\infty) = h(\widehat{\prod}_{i<\omega} x) = \widehat{\prod}_{i<\omega} h(x) = h(x)^\infty$ (using Lemma 3.4.4). For monomials m, we get $h(m) = h(\widehat{\prod}_{x\in\boldsymbol{\mathcal{X}}} x^{m(x)}) = \widehat{\prod}_{x\in\boldsymbol{\mathcal{X}}} h(x)^{m(x)}$, and $h(P) = h(\widehat{\sum}_{m\in P} m) = \widehat{\sum}_{m\in P} h(m)$ for polynomials, matching the definition in Lemma 3.6.14.

For summation $\widehat{\sum}_i P_i = \bigcup_i P_i$ we have $h(\bigcup_i P_i) = \widehat{\sum}_i \widehat{\sum}_{m\in P_i} h(m) = \widehat{\sum}_i h(P_i)$. For multiplication, we first rewrite the left-hand side:

$$\begin{aligned} h\left(\widehat{\prod_{i\in I}} P_i\right) = \widehat{\sum_{f\in\mathcal{F}}} h(\widehat{m}_f) &= \widehat{\sum_{f\in\mathcal{F}}}\widehat{\prod_{x\in\boldsymbol{\mathcal{X}}}} h(x)^{\sum_i f(i)(x)} \\ &= \widehat{\sum_{f\in\mathcal{F}}}\widehat{\prod_{x\in\boldsymbol{\mathcal{X}}}}\widehat{\prod_{i\in I}} h(x)^{f(i)(x)} = \widehat{\sum_{f\in\mathcal{F}}}\widehat{\prod_{i\in I}} h(f(i)). \end{aligned}$$

To see that this equals the right-hand side, we need strong distributivity:

$$\widehat{\prod_{i\in I}} h(P_i) = \widehat{\prod_{i\in I}}\widehat{\sum_{m\in P_i}} h(m) = \widehat{\sum_{f\in\mathcal{F}}}\widehat{\prod_{i\in I}} h(f(i)).$$ ☐

Remark 3.6.20. The assumption of strong distributivity is necessary. Indeed, combining Example 3.4.6 in the Viterbi semiring and Example 3.6.21 below yields a counterexample.

We close this section by comparing the definition of infinite products via choice functions to the one via finite subproducts (Theorem 3.4.2). The first example shows how the cardinality of the index set I can lead to different results. We then show that Example 3.6.16 yields a counterexample for countable products.

Example 3.6.21. Let $P = \{x_i \mid i < \omega\}$ and let I be an uncountable index set. For the product via finite subproducts, we get:

$$\bigsqcap \Big\{ \prod_{i \in F} P \mid F \subset_{\text{fin}} I \Big\} = \bigsqcap_{n<\omega} P^n.$$

For each n, we can choose $x_1 \cdots x_n \in P^n$. Hence $x_{[1,\infty)} \leq \bigsqcap_n P^n$ by Proposition 3.6.7.

For $\widehat{\prod}_i P$, consider any choice function f with $f(i) \in P$ for $i \in I$. Since P is countable and I uncountable, some $x \in \boldsymbol{\mathcal{X}}$ must be hit infinitely often. But then $\widehat{m}_f(x) = \infty$ and $x_{[1,\infty)} \not\leq \widehat{m}_f$. Since f was arbitrary, this shows that the two definitions do not coincide. More specifically, we get

$$\widehat{\prod_{i \in I}} P = \{x_i^\infty \mid i < \omega\},$$

since choosing the same monomial for each i clearly absorbs any other choice.

Example 3.6.22. Recall the polynomial $P = \{x_{p\mathbb{N}+r} \mid p \text{ is prime and } 0 \leq r < p\}$ of Example 3.6.16 and consider the countable product $\widehat{\prod}_{i<\omega} P$. The infimum over finite subproducts is, again, $\bigsqcap_n P^n = P^\infty$ and we know from Example 3.6.16 that $x_\mathbb{N} \leq P^\infty$.

Now consider any choice function f with $f(i) \in P$ for $i < \omega$. By construction of P, it is not possible to choose infinitely many pairwise disjoint monomials in P. Hence $\widehat{m}_f$ must have an exponent ≥ 2 and $x_\mathbb{N} \not\leq \widehat{m}_f$, so the two definitions give different results.

3.6.5 Restrictions

We briefly discuss two restrictions we can impose on $\mathbb{S}^\infty\llparenthesis\boldsymbol{\mathcal{X}}\rrparenthesis$ to obtain full continuity. The first restriction is a finiteness assumption for polynomials. The second option is to drop exponents, so that we obtain $\mathrm{PosBool}\llparenthesis\boldsymbol{\mathcal{X}}\rrparenthesis$ with infinitely many indeterminates.

Finite polynomials. We say that a polynomial P is finite if it contains only finitely many monomials (more precisely, there is a finite representative of P). Notice that this is not the same as requiring $\boldsymbol{\mathcal{X}}$ to be finite, as we still allow infinite monomials. Finite polynomials are always antichains (cf. Lemma 3.6.3). We prove full continuity by induction on the number of monomials.

Lemma 3.6.23. *Let $m \in \mathsf{Mon}$ and $\emptyset \neq S \subseteq \mathbb{S}^\infty\llparenthesis\boldsymbol{\mathcal{X}}\rrparenthesis$. Then, $m \cdot \bigsqcap S = \bigsqcap(m \cdot S)$.*

Proof. It suffices to prove $m \cdot \bigsqcap S \geq \bigsqcap(m \cdot S)$ due to monotonicity. We apply Proposition 3.6.7. On the left-hand side, we get

$$m \cdot \bigsqcap S = \{m \cdot m_f \mid f(q) \in q \text{ for all } q \in S\}.$$

On the right-hand side, we get

$$\bigsqcap(m \cdot S) = \{m_g \mid g(mq) \in mq \text{ for all } q \in S\}.$$

Fix a choice function g (on the set mS). It suffices to construct a choice function f (on the set S) such that $m_g \preceq m \cdot m_f$. For each $q \in S$, let $v_q \in q$ be a monomial such that $g(mq) = mv_q$ (notice that v_q may not be unique, but this is not an issue). We then set $f(q) = v_q$ so that $m \cdot f(q) = g(mq)$. Then $m_g = m \cdot m_f$ by comparing exponents. □

For the induction step, we need the following simple observation. This is the only place where the chain requirement is needed.

Lemma 3.6.24. *Let $S \subseteq \mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ be a chain and $p, q \in \mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ polynomials. Then,*

$$\bigsqcap(pS + qS) = \bigsqcap(p + q)S.$$

Proof. Every $(p+q)s \in (p+q)S$ can be written as $ps + qs \in pS + qS$, so direction $\leq$ is clear. For the other direction, let $ps_1 + qs_2 \in pS + qS$. As S is a chain, s_1 and s_2 are comparable, say $s_1 \leq s_2$. Then $ps_1 + qs_2 \geq (p+q)s_1 \geq \bigsqcap(p+q)S$ and direction $\geq$ follows. □

Proposition 3.6.25. *Let $p \in \mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ be finite and let $S \subseteq \mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ be a nonempty chain. Then,*

$$p \cdot \bigsqcap S = \bigsqcap(p \cdot S)$$

Proof. Let $\hat{p} \subseteq \mathsf{Mon}$ be a finite representative of p. We proceed by induction on $|\hat{p}|$. The statement is trivial for $\hat{p} = \emptyset$. For the induction step, assume that the claim holds for $\hat{p}$ and consider the polynomial $m + \hat{p}$ for some monomial m. Using Lemma 3.6.24 and continuity of addition (Theorem 3.6.12),

$$\bigsqcap(m + \hat{p})S = \bigsqcap(mS + \hat{p}S) = \bigsqcap mS + \bigsqcap \hat{p}S = (m + \hat{p}) \cdot \bigsqcap S,$$

where the last step holds by induction hypothesis and Lemma 3.6.23. □

Working with finite polynomials further implies that the two notions of infinite products (via choice functions and via infima) coincide in the countable case.

Lemma 3.6.26. *Let $(P_i)_{i<\omega}$ be a countable family of finite polynomials. Then,*

$$\widehat{\prod_{i<\omega}} P_i = \bigsqcap\Big\{\prod_{i \in F} P_i \mid F \subset_{\text{fin}} \omega\Big\}.$$

Proof. By Proposition 3.6.7, we can rewrite the right-hand side as $\bigsqcup_{g \in \mathcal{G}} m_g$, where $\mathcal{G}$ is the set of choice functions g with $g(F) \in \prod_{i \in F} P_i$ for all $F \subset_{\text{fin}} \omega$. We first show direction "$\leq$". Let $\widehat{m}_f$ be a product monomial of $\widehat{\prod}_i P_i$, so $f(i) \in P_i$ for all $i < \omega$. Define $g(F) = \prod_{i \in F} f(i)$ for all $F \subset_{\text{fin}} \omega$. Then $g \in \mathcal{G}$ and $\widehat{m}_f \preceq m_g$ by comparing exponents.

The other direction is more difficult. Consider a monomial m_g for $g \in \mathcal{G}$. We use Kőnig's lemma to construct f so that $m_g \preceq \widehat{m}_f$. To this end, we consider the finitely-branching tree with nodes $V = \{(m_0, m_1, \ldots, m_n) \mid n < \omega,\, m_i \in P_i \text{ for all } i\}$ and edges $(m_0, \ldots, m_n) \to (m_0, \ldots, m_n, m_{n+1})$. We prune this tree and retain only the nodes

$$V' = \{(m_0, \ldots, m_n) \in V \mid m_0 \cdots m_n \succeq g(F) \text{ for some } F \subset_{\text{fin}} \omega \text{ with } \{0, \ldots, n\} \subseteq F\}.$$

The pruned tree remains connected and infinite. To see this, note that if $(m_0, \ldots, m_n) \in V'$, then also $(m_0, \ldots, m_k) \in V'$ for all $k \leq n$ (using the same witnessing set F). Moreover, for each level n consider $F = \{0, \ldots, n\}$. As g must make some choice for F, say $g(F) = m_0 \cdots m_n$, there is at least one node $(m_0, \ldots, m_n)$ on level n.

Now by Kőnig's lemma, there is an infinite path from the root, corresponding to a sequence $(m_i)_{i<\omega}$ of monomials with $m_i \in P_i$. We set $f(i) = m_i$ and show that $m_g \preceq \widehat{m}_f$ by comparing exponents for each $x \in \boldsymbol{\mathcal{X}}$.

- First assume that all occurrences of x in the sequence happen in $m_0, \ldots, m_k$ for some $k < \omega$. Since $(m_0, \ldots, m_k) \in V'$, we have $m_0 \cdots m_k \succeq g(F)$ for some $F \subset_{\text{fin}} \omega$, so $m_g(x) \geq \sum_{i \leq k} m_i(x) = \widehat{m}_f(x)$.
- Otherwise there are indices $(i_j)_{j<\omega}$ such that x occurs in each m_{i_j}. Repeating the above argument for each node $(m_0, \ldots, m_{i_j})$ yields $m_g(x) \geq j$ for all j, so $m_g(x) = \infty \geq \widehat{m}_f(x)$. □

Although we do not consider provenance of propositional logic in this thesis, we mention that this may be a motivation to consider finite polynomials. Since each formula is finite, we can interpret it as polynomial in finitely many indeterminates (which then has to be finite). Sets of formulae can use an infinite set of propositional variables, so the conjunction over such a set can be expressed as product over finite polynomials in $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$.

No exponents. Dropping exponents from $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ results in the semiring $\text{PosBool}(\!|\boldsymbol{\mathcal{X}}|\!)$. More formally, we can define the infinite-indeterminate version of $\text{PosBool}(X)$ by the same quotient construction, starting from monomials $\mathsf{Mon}_{\boldsymbol{\mathcal{X}}}(\mathbb{B})$ (i.e., sets of indeterminates) instead of $\mathsf{Mon}_{\boldsymbol{\mathcal{X}}}(\mathbb{N}^\infty)$. All of our proofs for $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$ also go through for $\text{PosBool}(\!|\boldsymbol{\mathcal{X}}|\!)$. The counter-examples, on the other hand, all rely on differences in the exponents and are thus no longer applicable (an exception is Example 3.6.15, which relies on the non-idempotent multiplication of the Viterbi semiring).

This is not a coincidence: in fact, $\text{PosBool}(\!|\boldsymbol{\mathcal{X}}|\!)$ is the free completely-distributive lattice generated by $\boldsymbol{\mathcal{X}}$, and this is already implied by our results on $\mathbb{S}^\infty(\!|\boldsymbol{\mathcal{X}}|\!)$. To see this, we first note that multiplication is now the same as infimum, so the missing part of full continuity in Theorem 3.6.12 becomes trivial. Moreover, the definitions of infinite products via choice functions (Definition 3.6.17) and via finite subproducts (Theorem 3.4.2) both coincide with infima as well. We can thus obtain the following corollaries from Proposition 3.6.18 and Theorem 3.6.19.

Corollary 3.6.27. $\mathrm{PosBool}(\!(\boldsymbol{\mathcal{X}})\!)$ *is a completely-distributive lattice. In particular, it is an absorptive, fully-continuous infinitary semiring with strong distributivity.*

Corollary 3.6.28. *Let K be a completely-distributive lattice (viewed as an absorptive semiring). Every mapping $h\colon \boldsymbol{\mathcal{X}} \to K$ extends uniquely to an infinitary homomorphism $h\colon \mathrm{PosBool}(\!(\boldsymbol{\mathcal{X}})\!) \to K$.*

The free completely-distributive lattice has been studied before. For instance, Markowsky [Mar79] denotes it as $\mathrm{FD}(\!(\boldsymbol{\mathcal{X}})\!)$ and defines its elements to be the downwards-closed subsets $T \subseteq \mathcal{P}(\boldsymbol{\mathcal{X}})$ of the powerset lattice on $\boldsymbol{\mathcal{X}}$, such as $T = \{\{x,y\},\{x\},\{y\},\{z\},\emptyset\}$. We briefly show how this relates to our definition. To simplify notation, we view monomials $m \in \mathsf{Mon}(\mathbb{B})$ as sets $m \subseteq \boldsymbol{\mathcal{X}}$ of indeterminates. Absorption on monomials is the inverse of the subset relation, so polynomials are implicitly upwards-closed (w.r.t. set inclusion) by the quotient construction. For example, $xy + z = xy + z + xyz$. Writing $\overline{m} = \boldsymbol{\mathcal{X}} \setminus m$ for the *complement* of a monomial m, we can define an isomorphism by

$$\pi\colon \mathrm{PosBool}(\!(\boldsymbol{\mathcal{X}})\!) \to \mathrm{FD}(\!(\boldsymbol{\mathcal{X}})\!), \qquad P \mapsto \{\overline{m} \mid m \subseteq \boldsymbol{\mathcal{X}}, m \leq P\}.$$

For the universal property, Markowsky specifies an embedding $I\colon \boldsymbol{\mathcal{X}} \to \mathrm{FD}(\!(\boldsymbol{\mathcal{X}})\!)$ and shows that for any assignment $\Theta\colon \boldsymbol{\mathcal{X}} \to K$ into a completely-distributive lattice K, there is a unique lattice homomorphism Θ^* that preserves suprema and infima and so that $\Theta^* I = \Theta$. In our setting, the embedding $\iota\colon \boldsymbol{\mathcal{X}} \to \mathrm{PosBool}(\!(\boldsymbol{\mathcal{X}})\!)$ is simply $x \mapsto \{x\}$. If h is obtained by applying Corollary 3.6.28 to Θ, then it is easy to verify that the following diagram commutes, so Corollary 3.6.28 indeed reproves Markowsky's result.

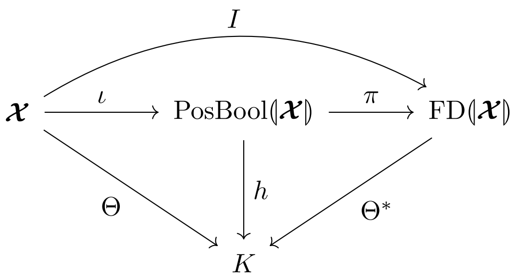

3.7 Absorptive Polynomials with Coefficients

The semirings $\mathbb{S}(X)$ and $\mathbb{S}^\infty(X)$ of (generalised) absorptive polynomials are the universal absorptive (fully-continuous) semirings, which makes them a fruitful tool for proofs about absorptive semirings and also for provenance analysis where they represent the most general information (among the absorptive semirings). However, in some cases it is desirable to have polynomials with coefficients from some semiring K, and for our purposes we are interested in absorptive coefficient semirings. We can of course use the standard polynomial semirings $K[X]$, which inherit some of the properties, such as idempotence, from K. But they do not inherit absorption, since $x + xy \neq x$ (independent of K).

This situation is unsatisfactory for several reasons. First, the theory we develop for absorptive semirings does not apply to $K[X]$, including our main results for fixed-point computation, fixed-point logics, and games. Second, inductive arguments by means of $K[x, y] \cong K[x][y]$ become more difficult: if our argument requires the coefficient semiring to be absorptive, then it does not apply to $K[x][y]$ (which has coefficients in the non-absorptive semiring $K[x]$). Lastly, if we only wish to evaluate polynomials in the absorptive semiring K (or in a larger absorptive semiring that K embeds to), then the polynomials in $K[X]$ may be unnecessarily large and thus inconvenient for applications, as simplifications due to absorption are not applied.

Instead, we propose a general construction of absorptive polynomials $\mathrm{AbsPoly}(K, E, X)$ with coefficients in an absorptive semiring K, exponents in $E \in \{\mathbb{B}, \mathbb{N}, \mathbb{N}^\infty\}$, and a finite set of indeterminates X. It is not immediately obvious how to define such semirings. For example, consider the sum $ax + ax^2 + bx^2$ with $x \in X$ and incomparable coefficients $a, b \in K$. Since x absorbs x^2, we would expect that $(ax + ax^2) + bx^2 = ax + bx^2$. On the other hand, $ax + (ax^2 + bx^2) = ax + (a + b)x^2$, and both must clearly be equal by associativity. To find a suitable representation of this polynomial, we follow an idea suggested by Mrkonjić and think of a monomial ax as also contributing the coefficient a to all smaller monomials, thus counting as ax, ax^2, ax^3, ... (and potentially also as ax^∞). For the monomial x^2 we then take the supremum over all contributed coefficients, which is $\bigsqcup\{a, b\} = a + b$ in both cases so that we indeed have equality. For an elegant implementation of this idea, we define $\mathrm{AbsPoly}(K, E, X)$ as subsemiring of $\mathrm{Poly}(K, E, X)$, consisting only of those polynomials whose coefficients equals the supremum of the contributed values.

The resulting construction works uniformly for all absorptive K and $E \in \{\mathbb{B}, \mathbb{N}, \mathbb{N}^\infty\}$. In particular, choosing $K = \mathbb{B}$ yields an isomorphic construction of the semirings $\mathrm{PosBool}(X)$, $\mathbb{S}(X)$, and $\mathbb{S}^\infty(X)$. Our main results are universal properties for both finite and infinite coefficient semirings, subsuming the universal properties of $\mathrm{PosBool}(X)$ and $\mathbb{S}^\infty(X)$.

3.7.1 Construction

Let K be an absorptive semiring, E a semiring in $\{\mathbb{B}, \mathbb{N}, \mathbb{N}^\infty\}$, and X a finite set of indeterminates. Recall that $\mathsf{Mon}_E(X)$ denotes the set of monomials $m\colon X \to E$. If E and X are clear from the context, we simply write Mon. Further recall that $\mathrm{Poly}(K, E, X)$ is the power series semiring consisting of all functions $P\colon \mathsf{Mon}_E(X) \to K$ with pointwise addition and the usual polynomial multiplication $(P \cdot Q)\colon m \mapsto \widehat{\sum}_{m=m_1\cdot m_2} P(m_1) \cdot Q(m_2)$. Since the natural order is defined by addition, suprema and infima are also evaluated pointwise. That is, for a set $S \subseteq \mathrm{Poly}(K, E, X)$ and assuming that suprema and infima exist in K, we have

$$\Big(\bigsqcup_{P \in S} P\Big)\colon\ m \mapsto \bigsqcup_{P \in S} P(m), \qquad \Big(\bigsqcap_{P \in S} P\Big)\colon\ m \mapsto \bigsqcap_{P \in S} P(m).$$

We define $\mathrm{AbsPoly}(K, E, X)$ as a subsemiring of $\mathrm{Poly}(K, E, X)$. More precisely, as the semiring induced by those power series with the following property, together with polynomial addition and multiplication (but with a different neutral element for multiplication).

Definition 3.7.1. An element $P \in \mathrm{Poly}(K, E, X)$ is *absorption-compatible* if for all monomials $m_1, m_2 \in \mathsf{Mon}_E(X)$, the following implication holds:

$$m_1 \succeq m_2 \quad \Longrightarrow \quad P(m_1) \le P(m_2).$$

Lemma 3.7.2. *Let K be an absorptive, fully-complete semiring. For $P \in \mathrm{Poly}(K, E, X)$, the following are equivalent.*

(1) *P is absorption-compatible,*

(2) *$P(m) = \bigsqcup_{v \succeq m} P(v)$, for all $m \in \mathsf{Mon}$,*

(3) *$P \in \{Q \cdot \top \mid Q \in \mathrm{Poly}(K, E, X)\}$, where $\top \colon m \mapsto 1_K$ is the greatest element.*

Proof. For (3), we observe $\big(Q \cdot \top\big)(m) = \bigsqcup_{m_1 m_2 = m} Q(m_1) \cdot 1_K = \bigsqcup_{v \succeq m} Q(m)$ for all $m \in \mathsf{Mon}$, so (2) and (3) are equivalent. The equivalence of (1) and (2) is clear. □

Lemma 3.7.3. *Let $P, Q \in \mathrm{Poly}(K, E, X)$ be absorption-compatible. Then $P + Q$ and $P \cdot Q$ are absorption-compatible as well.*

Proof. Let $m_1 \succeq m_2$. We have $(P + Q)(m_1) = P(m_1) + Q(m_1) \leq P(m_2) + Q(m_2) = (P + Q)(m_2)$ by monotonicity of K.

For multiplication, we first note that every split $m_1 = v_1 \cdot v_2$ can be extended to a split $m_2 = v_1 \cdot v_2'$ with $v_2 = v_2 \cdot u$ for an appropriate monomial u. Indeed, since E is naturally ordered, $m_1(x) \leq m_2(x)$ implies that there is $e \in E$ with $m_1(x) + e = m_2(x)$, so we can set $u(x) = e$ (and repeat this argument for each $x \in X$). Since Q is absorption-compatible, we have $Q(v_2') \geq Q(v_2)$, and it follows that

$$(P \cdot Q)(m_1) = \bigsqcup_{v_1 v_2 = m_1} P(v_1) Q(v_2) \leq \bigsqcup_{v_1 v_2 = m_2} P(v_1) Q(v_2) = (P \cdot Q)(m_2).$$ □

Since absorption-compatible power series are closed under the semiring operations, they induce a semiring. For the neutral elements, notice that 0 is absorption-compatible, but 1 is not. Instead, $\top \colon m \to 1_K$ becomes the neutral element of the new semiring.

Definition 3.7.4. Let K be an absorptive semiring, $E \in \{\mathbb{B}, \mathbb{N}, \mathbb{N}^\infty\}$ and X a finite set of indeterminates. Let $A \subseteq \mathrm{Poly}(K, E, X)$ be the set of absorption-compatible power series. The semiring of *absorptive polynomials with coefficients in K* (over X and exponents E) is defined as

$$\mathrm{AbsPoly}(K, E, X) \coloneqq (A, +, \cdot, 0, 1_A), \text{ with } 1_A \coloneqq \top.$$

Proposition 3.7.5. $\mathrm{AbsPoly}(K, E, X)$ *is an absorptive semiring.*

Proof. Addition and multiplication are well-defined by Lemma 3.7.3. Clearly, 0 is absorption-compatible, neutral for addition and annihilating for multiplication. Moreover, $(1_A \cdot P)(m) = \bigsqcup_{v \geq m} 1_K \cdot P(v) = P(m)$ by Lemma 3.7.2, so 1_A is neutral and $\mathrm{AbsPoly}(K, E, X)$ is a well-defined semiring.

It remains to show absorption. Since $\mathrm{AbsPoly}(K, E, X)$ inherits idempotence from K and is thereby naturally ordered, it suffices to prove that 1_A is the greatest element (by Proposition 3.1.7). This follows from the fact that K is absorptive so that 1_K is the greatest coefficient. □

Proposition 3.7.6. *If K is a lattice semiring, then so is* $\mathrm{AbsPoly}(K, \mathbb{B}, X)$.

Proof. Due to absorption, it suffices to show that multiplication is idempotent. For $P \in \mathrm{AbsPoly}(K, \mathbb{B}, X)$ and each monomial m, we have

$$(P \cdot P)\,(m) = \bigsqcup_{m_1 m_2 = m} P(m_1) \sqcup P(m_2) = P(m) \sqcup P(m) = P(m).$$

To see this, note that $m \cdot m = m$ for exponents in $\mathbb{B}$, so the supremum is assumed for $m_1 = m_2 = m$. □

This construction generalises the semirings $\mathbb{S}^\infty(X)$, $\mathbb{S}(X)$, and $\mathrm{PosBool}(X)$ without coefficients. Indeed, we obtain these in the special case $K = \mathbb{B}$.

Proposition 3.7.7. *The following semirings are isomorphic:*

$$\begin{aligned} \mathbb{S}^\infty(X) &\cong \mathrm{AbsPoly}(\mathbb{B}, \mathbb{N}^\infty, X), \\ \mathbb{S}(X) &\cong \mathrm{AbsPoly}(\mathbb{B}, \mathbb{N}, X), \\ \mathrm{PosBool}(X) &\cong \mathrm{AbsPoly}(\mathbb{B}, \mathbb{B}, X). \end{aligned}$$

Proof. In each case, the isomorphism from $\mathrm{AbsPoly}(\mathbb{B}, E, X)$ to the respective semiring is given by

$$f\colon P \mapsto \mathsf{Maximals}\{m \in \mathsf{Mon}_E(X) \mid P(m) = 1\}.$$ □

3.7.2 Continuity

We next consider the completeness and continuity properties of the natural order.

Lemma 3.7.8. *Let $S \subseteq \mathrm{Poly}(K, E, X)$ be a set of absorption-compatible power series. Then $\bigsqcup S$ and $\bigsqcap S$ are absorption-compatible, if they exist.*

Proof. Suprema and infima are defined pointwise, so for $m_1 \succeq m_2$ we have

$$\left(\bigsqcup S\right)(m_1) = \bigsqcup_{P \in S} P(m_1) \leq \bigsqcup_{P \in S} P(m_2) = \left(\bigsqcup S\right)(m_2),$$

and analogously for infima. □

It follows that suprema (and also infima) in $\mathrm{AbsPoly}(K, E, X)$ and $\mathrm{Poly}(K, E, X)$ coincide. Moreover, full continuity of K transfers to $\mathrm{AbsPoly}(K, E, X)$. We first prove this for the most involved case $E = \mathbb{N}^\infty$.

Proposition 3.7.9. *If K is fully continuous, then* $\mathrm{AbsPoly}(K, \mathbb{N}^\infty, X)$ *is fully-continuous as well.*

Proof. Since $\mathrm{AbsPoly}(K, \mathbb{N}^\infty, X)$ is idempotent, addition trivially commutes with suprema. Let $C \subseteq \mathrm{AbsPoly}(K, \mathbb{N}^\infty, X)$ be a chain. Full continuity of K directly implies that addition commutes with infima of chains, and multiplication with suprema:

$$\left(Q + \prod_{P\in C} P\right)(m) = Q(m) + \prod_{P\in C} P(m) = \prod_{P\in C} Q(m) + P(m) = \left(\prod_{P\in C} P + Q\right)(m),$$

and

$$\begin{aligned}\left(Q \cdot \coprod_{P\in C} P\right)(m) &= \coprod_{m_1 m_2 = m} \left(Q(m_1) \cdot \coprod_{P\in C} P(m_2)\right) \\ &= \coprod_{P\in C} \coprod_{m_1 m_2 = m} Q(m_1) \cdot P(m_2) = \left(\coprod_{P\in C} Q \cdot P\right)(m).\end{aligned}$$

The only interesting case is multiplication with infima. The issue is that for monomials m with infinite exponents, we have a supremum over infinitely many splits $m = m_1 m_2$ (and it is not clear a priori whether this supremum commutes with the infimum over C).

We resolve this as follows. For a split $m = m_1 m_2$, define $\widehat{m}_1$ and $\widehat{m}_2$ by setting $\widehat{m}_1(x) = \widehat{m}_2(x) = \infty$ for all x with $m(x) = \infty$. Then clearly $m = \widehat{m}_1 \widehat{m}_2$ and also $P(\widehat{m}_1) \geq P(m_1)$ and $P(\widehat{m}_2) \geq P(m_2)$ for all $P \in \mathrm{AbsPoly}(K, \mathbb{N}^\infty, X)$. We say that a split $m = m_1 m_2$ is ∞-*saturated* if $\widehat{m}_1 = m_1$ and $\widehat{m}_2 = m_2$. Crucially, there are only finitely many ∞-saturated splits for each monomial m (there are only finitely many indeterminates and each exponent of m_1, m_2 must either be ∞ or smaller than the one of m). With this insight, we obtain:

$$\begin{aligned}\left(Q \cdot \prod_{P\in C} P\right)(m) &= \coprod_{m_1 m_2 = m} Q(m_1) \cdot \prod_{P\in C} P(m_2) \\ &= \coprod_{m_1 m_2 = m} \prod_{P\in C} Q(m_1) \cdot P(m_2) \\ &\overset{(1)}{=} \coprod_{\substack{m_1 m_2 = m \\ \infty\text{-sat.}}} \prod_{P\in C} Q(m_1) \cdot P(m_2) \\ &\overset{(2)}{=} \prod_{P\in C} \coprod_{\substack{m_1 m_2 = m \\ \infty\text{-sat.}}} Q(m_1) \cdot P(m_2) \\ &\overset{(1)}{=} \prod_{P\in C} \coprod_{m_1 m_2 = m} Q(m_1) \cdot P(m_2) = \left(\prod_{P\in C} Q \cdot P\right)(m).\end{aligned}$$

For the steps labelled (1), recall that $P(\widehat{m}_i) \geq P(m_i)$, so (by monotonicity of all involved operations) the supremum assumes its largest values for the ∞-saturated splits. For (2), note that the coefficients $Q(m_1) \cdot P(m_2)$ form a chain in P. The supremum is a finite sum and we can thus apply the Chain splitting lemma (Lemma 3.2.8). □

Corollary 3.7.10. *If K is fully continuous, then* $\mathrm{AbsPoly}(K, E, X)$ *is fully-continuous as well.*

Proof. Recall that $E \in \{\mathbb{B}, \mathbb{N}, \mathbb{N}^\infty\}$. The argument in the proof of Proposition 3.7.9 also applies to $\mathbb{B}$ and $\mathbb{N}$, as the number of splits is always finite. □

3.7.3 Universal Properties

We prove two universal properties for fully-continuous target semirings, depending on the coefficient semiring K. The ideas for the universality proof of $\mathbb{S}^\infty(X)$ can be adapted to the more general setting with finite K, but new arguments are needed for infinite K.

We should first clarify what kind of universal property we consider. As usual, every assignment $f\colon X \to L$ into an appropriate target semiring L should lift in a unique way to a fully-continuous homomorphism $h_f\colon \mathrm{AbsPoly}(K,E,X) \to L$. Requiring that $L = K$ would be quite restrictive, so instead we assume that K can be embedded into L by means of a fully-continuous homomorphism $h\colon K \to L$, which has to be respected by h_f. (For $K = \mathbb{B}$, this is the trivial homomorphism $h(0) = 0$ and $h(1) = 1$.) As a commutative diagram:

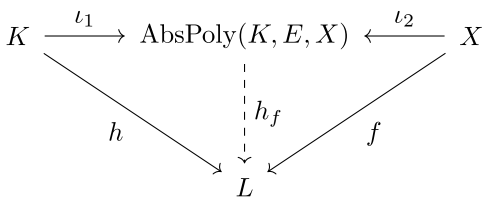

Here, ι_1, ι_2 are the natural embeddings. That is, $\iota_1(k) = P_k$ for $P_k\colon m \mapsto k$, and $\iota_2(x) = x + x^2 + x^3 + \dots + x^\infty$ (for $E = \mathbb{N}^\infty$). Alternatively, we can view $\mathrm{AbsPoly}(K,E,X)$ as coproduct of the free absorptive, fully-continuous semiring (i.e., $\mathbb{S}^\infty(X)$, $\mathbb{S}(X)$, or $\mathrm{PosBool}(X)$) with the coefficient semiring K.

Finite Coefficient Semirings

In order to adapt the universality proof for $\mathbb{S}^\infty(X)$ of [Naa19, DGNT21] to the more general setting with coefficients, we first introduce a new representation of $\mathrm{AbsPoly}(K,E,X)$ that is more convenient when K is finite.

Antichain representation. Consider the set $K \times \mathsf{Mon}_E(X)$ and order it by the product order, i.e., $(c,m) \leq (c',m')$ if $c \leq c'$ (for $c, c' \in K$) and $m \preceq m'$. We represent an absorptive polynomial $P \in \mathrm{AbsPoly}(K,E,X)$ by the set $\mathsf{Maximals}(P) := \mathsf{Maximals}\{(P(m),m) \mid m \in \mathsf{Mon}\}$ of maximal coefficient-monomial pairs occurring in P. For example, we have $\mathsf{Maximals}(0) = (0,1)$ for the zero polynomial.

The original proof for $\mathbb{S}^\infty(X)$ exploits the finiteness of antichains of monomials. This property also holds in the setting with finite K.

Lemma 3.7.11. *Let $P \in \mathrm{AbsPoly}(K,E,X)$.*

(1) *For each $m \in \mathsf{Mon}$ there is $m' \in \mathsf{Mon}$ so that $(P(m),m) \leq (P(m'),m') \in \mathsf{Maximals}(P)$. In particular, $\mathsf{Maximals}(P) \neq \emptyset$.*

(2) *If all ascending chains and antichains in K are finite, then $\mathsf{Maximals}(P)$ is finite.*

Proof. For (1), consider a pair $(P(m),m)$ and suppose no such m' exists. Then $(P(m),m)$ cannot be maximal and by repeating this argument we obtain an infinite sequence $(P(m),m) < (P(m_2),m_2) < (P(m_3),m_3) < \cdots$. Then $m \prec m_2 \prec m_3 \prec \cdots$ is an infinite ascending

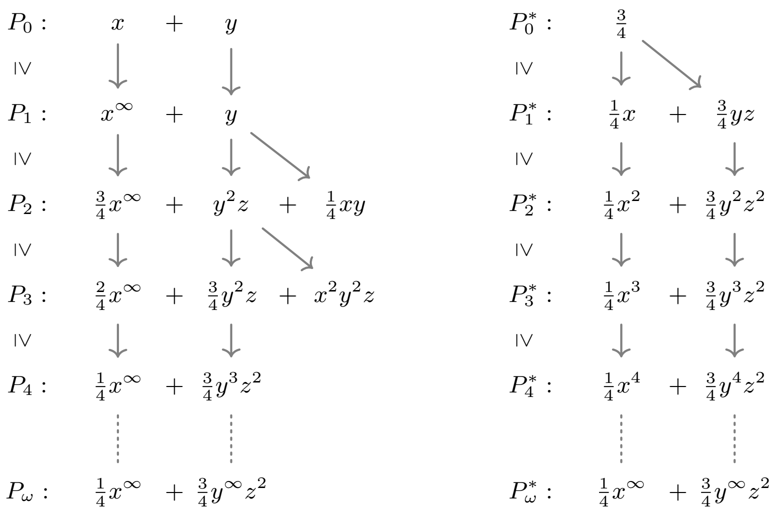

Figure 3.4: A polynomial ω-chain (left) with coefficients in $(\{0, \frac{1}{4}, \frac{2}{4}, \frac{3}{4}, 1\}, \max, \min, 0, 1)$ and the corresponding canonical chain (right) for the proof of Theorem 3.5.7. The arrows indicate absorption between coefficient-monomial pairs of consecutive polynomials and induce a directed graph on which we apply Kőnig's lemma. (Adapted from [DGNT21].)

sequence, but such a sequence does not exist in Mon, contradiction.

For (2), notice that the natural order on K and absorption on $\mathsf{Mon}_E(X)$, and hence also their product, are (duals of) well-partial orders. In particular, they do not admit infinite antichains, and hence the antichain $\mathsf{Maximals}(P)$ must be finite. ☐

The natural order is the pointwise order in the original representation of $\mathrm{AbsPoly}(K, E, X)$. In the antichain representation, it can be characterised in the same way as for $\mathbb{S}^\infty(X)$.

Lemma 3.7.12. *Let $P, Q \in \mathrm{AbsPoly}(K, E, X)$. Then $P \leq Q$ holds if, and only if, for each $(c, m) \in \mathsf{Maximals}(P)$ there is $(c', m') \in \mathsf{Maximals}(Q)$ with $(c, m) \leq (c', m')$.*

Proof. Let $(c, m) \in \mathsf{Maximals}(P)$. Since $P \leq Q$, we have $(c, m) \leq (Q(m), m)$, and by Lemma 3.7.11 (1) we obtain the desired pair $(Q(m), m) \leq (Q(m'), m') \in \mathsf{Maximals}(Q)$. ☐

In the following, we identify P and $\mathsf{Maximals}(P)$ and simply write $(c, m) \in P$ (where always $c = P(m)$). If $\mathsf{Maximals}(P) = \{(c_1, m_1), \dots, (c_k, m_k)\}$, we use the notation $P = c_1 m_1 + \dots + c_k m_k$ as finite sum. In particular, we write cm for the polynomial P with $\mathsf{Maximals}(P) = \{(c, m)\}$.

Universality proof. We now adapt the proof of the universal property of $\mathbb{S}^\infty(X)$ (Theorem 3.5.7, see [DGNT21] for the original proof) to the more general setting with coefficients.

Theorem 3.7.13 (Universality I)**.** *Let K be a finite absorptive semiring and L an absorptive, fully-continuous semiring. For each homomorphism $h \colon K \to L$ and assignment $f \colon X \to L$,* *there is a unique fully-continuous semiring homomorphism* $h_f\colon \mathrm{AbsPoly}(K, \mathbb{N}^\infty, X) \to L$ *that simultaneously extends both* h *and* f*.*

Proof. We first lift f to monomials by $f(m) = \prod_{x\in X} f(x)^{m(x)}$, where $f(x)^\infty$ is the infinitary power operation in L. It follows by comparing exponents that $f(m_1 m_2) = f(m_1)\cdot f(m_2)$ for all monomials $m_1, m_2 \in \mathsf{Mon}$. We then define h_f for $P \in \mathrm{AbsPoly}(K, \mathbb{N}^\infty, X)$ as follows:

$$h_f(P) = \bigsqcup_{m\in\mathsf{Mon}} h(P(m))\cdot f(m) = \bigsqcup_{(c,m)\in P} h(c)\cdot f(m).$$

The two suprema are equal by absorption of L. To see why this is unique, first note that f is uniquely defined on monomials, since h_f must commute with multiplication and $h_f(x^\infty) = \prod_{n<\omega} h_f(x^n)$ by full continuity. Moreover, we must have $h_f(c\cdot m) = h_f(c\cdot 1)\cdot h_f(1\cdot m) = h(c)\cdot f(m)$. Uniqueness then follows since h must commute with suprema (of chains, and hence also of arbitrary sets).

We next show that h commutes with the semiring operations. For multiplication, we have

$$\begin{aligned} h_f(P\cdot Q) &= \bigsqcup_{m\in\mathsf{Mon}} f(m)\cdot h\Big(\bigsqcup_{m_1 m_2 = m} P(m_1)\cdot Q(m_2)\Big) \\ &= \bigsqcup_{m\in\mathsf{Mon}}\ \bigsqcup_{m_1 m_2 = m} f(m_1)\cdot f(m_2)\cdot h(P(m_1))\cdot h(Q(m_2)) \\ &= \bigsqcup_{m_1\in\mathsf{Mon}} f(m_1)\cdot h(P(m_1))\cdot \bigsqcup_{m_2\in\mathsf{Mon}} f(m_2)\cdot h(Q(m_2)) = h_f(P)\cdot h_f(Q). \end{aligned}$$

For suprema over arbitrary sets $S \subseteq \mathrm{AbsPoly}(K, \mathbb{N}^\infty, X)$ (which includes addition for finite S), this follows directly from full continuity of L and h:

$$h_f\Big(\bigsqcup_{P\in S} P\Big) = \bigsqcup_{m\in\mathsf{Mon}} f(m)\cdot h\Big(\bigsqcup_{P\in S} P(m)\Big) = \bigsqcup_{m\in\mathsf{Mon}}\ \bigsqcup_{P\in S} f(m)\cdot h(P(m)) = \bigsqcup_{P\in S} h(P).$$

It remains to show that h commutes with infima of chains. Since $\mathrm{AbsPoly}(K, \mathbb{N}^\infty, X)$ is countable (for finite K and X), it suffices to prove this for descending ω-chains $(P_i)_{i<\omega}$ (by Lemma 3.2.7). One direction trivially follows from monotonicity of h_f, so we only need to prove that

$$\prod_{i<\omega} h_f(P_i) \le h_f\Big(\prod_{i<\omega} P_i\Big).$$

We first consider the case of single monomials. Let $(m_i)_{i<\omega}$ be a descending ω-chain of monomials. Recall that X is finite, so we can write $m_i = \prod_{x\in X} x^{m_i(x)}$. As the m_i form a descending chain, the exponents $(m_i(x))_{i<\omega}$ form an ascending chain for each $x \in X$. By Lemma 3.2.8 (Chain splitting) and the definition of f,

$$\prod_{i<\omega} f(m_i) = \prod_{x\in X}\prod_{i<\omega} f(x)^{m_i(x)} \overset{(*)}{=} \prod_{x\in X} f(x)^{\bigsqcup_i m_i(x)} = f\Big(\prod_{i<\omega} m_i\Big),$$

where $(*)$ can easily be seen by case distinction on the exponent $\bigsqcup_{i<\omega} m_i(x) \in \mathbb{N}^\infty$.

For the general case, let $P_\omega = \prod_{i<\omega} P_i$ and write this as $P_\omega = c_1 m_1 + \cdots + c_n m_n$. The idea for the proof is to define a second canonical ω-chain $(P_i^*)_{i<\omega}$ with the same infimum and a simple structure that makes it possible to apply the Chain splitting lemma. For a monomial $m \in \mathsf{Mon}$, we define the canonical monomial chain $(m_j^*)_{j<\omega}$ as follows:

$$m_j^*(x) = \min(j, m(x)), \qquad \text{for all } x \in X.$$

For a coefficient-monomial pair (c, m), we consider the canonical chain $(c, m_j^*)_{j<\omega}$ in $K \times \mathsf{Mon}$, using the same coefficient in each step. We need the following properties of these chains for the proof:

(A) If $(c, m) \leq (d, v)$, then $(c, m_j^*) \preceq (d, v_j^*)$ for all $j < \omega$.

(B) If $(c, m) = \prod_{i<\omega}(c_i, m_i)$ for an ω-chain $(c_i, m_i)_{i<\omega}$, then $\forall j\, \exists i$: $(c, m_j^*) \succeq (c_i, m_i)$.

(C) In particular, $\prod_{j<\omega} m_j^* = m$.

Properties (1) and (3) are straightforward. For (2), the finiteness of K implies that $c = c_i$ for some i. Analogously, for each indeterminate x the exponent $m_j^*(x)$ must be reached (or exceeded) in m_i for some i, and taking the maximum of all these indices proves the claim.

The canonical polynomial chain $(P_j^*)_{j<\omega}$ is then defined by $P_j^* = c_1(m_1)_j^* + \cdots + c_n(m_n)_j^*$ for each $j < \omega$ (see Fig. 3.4 for an example). We claim that the following holds:

$$\forall j\, \exists i\colon\ P_j^* \geq P_i.$$

We first show that the claim implies the theorem:

$$\begin{aligned}
\prod_{i<\omega} h(P_i) \overset{(1)}{\leq} \prod_{j<\omega} h(P_j^*) &= \prod_{j<\omega} \Big(h(c_1)f((m_1)_j^*) + \cdots + h(c_n)f((m_n)_j^*)\Big) \\
&\overset{(2)}{=} h(c_1) \prod_{j<\omega} f((m_1)_j^*) + \cdots + h(c_n) \prod_{j<\omega} f((m_n)_j^*) \\
&\overset{(3)}{=} h(c_1) f\Big(\prod_{j<\omega} (m_1)_j^*\Big) + \cdots + h(c_n) f\Big(\prod_{j<\omega} (m_n)_j^*\Big) \\
&\overset{(4)}{=} h(c_1)f(m_1) + \cdots + h(c_n)f(m_n) = h_f(P_\omega),
\end{aligned}$$

where (1) follows from the claim, (2) holds by Lemma 3.2.8 (Chain splitting) and full continuity of L, (3) was shown above, and (4) holds due to property (C) above.

To prove the claim, assume towards a contradiction that there is a j such that $P_j^* \not\geq P_i$ for all $i < \omega$. Let us fix an $i < \omega$ for the moment. Because of $P_j^* \not\geq P_i$, there is a pair $(c_i, m_i) \in P_i$ with $P_j^* \not\geq c_i m_i$. Because of $P_{i-1} \geq P_i$, there is further $(c_{i-1}, m_{i-1}) \in P_{i-1}$ with $(c_{i-1}, m_{i-1}) \geq (c_i, m_i)$. But then also $P_j^* \not\geq c_{i-1}m_{i-1}$ (as otherwise $P_j^* \geq c_{i-1}m_{i-1} \geq c_i m_i$). By repeating this argument, we obtain a finite chain $(c_0, m_0) \geq (c_1, m_1) \geq \cdots \geq (c_i, m_i)$ with the property that $(c_k, m_k) \in P_k$ and $P_j^* \not\geq c_k m_k$ for all $0 \leq k \leq i$.

This argument applies to all $i < \omega$, so we obtain arbitrarily long finite chains with this property. By Kőnig's lemma (recall that $\mathsf{Maximals}(P_i)$ is finite for all i), there must be an infinite chain $(c_i, m_i)_{i<\omega}$ with $(c_i, m_i) \in P_i$ and $P_j^* \not\geq c_i m_i$ for all $i < \omega$. Let $(c_\omega, m_\omega) = \prod_{i<\omega}(c_i, m_i)$. Because of $c_i m_i \leq P_i$ for all i, we have $c_\omega m_\omega \leq P_\omega$, so by Lemma 3.7.12 there is a pair $(d, v) \in P_\omega$ with $(c_\omega, m_\omega) \leq (d, v)$. By considering the corresponding canonical monomial chains $(d, v_k^*)_{k<\omega}$ and $(c_\omega, (m_\omega)_k^*)_{k<\omega}$ at $k = j$, we obtain a contradiction: By property (B) there is an i with $(c_\omega, (m_\omega)_j^*) \geq (c_i, m_i)$, and by property (A) we have $(d, v_j^*) \geq (c_\omega, (m_\omega)_j^*)$. Because of $(d, v_j^*) \in P_j^*$, we obtain $P_j^* \geq dv_j^* \geq c_\omega(m_\omega)_j^* \geq c_i m_i$, contradicting our assumption. This proves the claim and ends the overall proof. □

We also note the following result that applies in particular to PosBool(X). Since AbsPoly($K, \mathbb{B}, X$) is finite (for finite K), the continuity part becomes trivial.

Corollary 3.7.14. *Let K be a finite lattice semiring and L a fully-continuous lattice semiring. For each homomorphism $h\colon K \to L$ and assignment $f\colon X \to L$, there is a unique semiring homomorphism $h_f\colon$ AbsPoly($K, \mathbb{B}, X$) $\to L$ that commutes with arbitrary suprema and infima, and simultaneously extends both h and f.*

Proof. Since K is a lattice semiring and we consider exponents in $\mathbb{B}$, also AbsPoly($K, \mathbb{B}, X$) is a lattice semiring. Being finite, it is (trivially) a complete lattice, and due to distributivity (and the dual property) it is moreover fully continuous.

Define $f(m) = \prod_{x \in X} f(x)^{m(x)}$ and $h_f(P) = \bigsqcup_{m \in \mathsf{Mon}} h(P(m)) \cdot f(m)$ as before. By comparing exponents and using multiplicative idempotence of L, we again have $f(m_1 m_2) = f(m_1) f(m_2)$ for all $m_1, m_2 \in \mathsf{Mon}$. The same arguments as in the proof of Theorem 3.7.13 show that h_f is uniquely defined and commutes with the semiring operations. Since AbsPoly($K, \mathbb{B}, X$) is a finite lattice semiring, this already implies that h_f commutes with arbitrary suprema and infima. ☐

Remark 3.7.15. The semirings AbsPoly($K, \mathbb{N}, X$), for finite K, also satisfy a universal property: every ω-continuous homomorphism $h\colon K \to L$ and assignment $f\colon X \to L$ into an absorptive, ω-continuous semiring L uniquely extend to an ω-continuous homomorphism $h_f\colon$ AbsPoly($K, \mathbb{N}, X$) $\to L$. Indeed, we can again set $h_f(P) = \bigsqcup_m h(P(m)) \cdot f(m)$ and the arguments for multiplication and suprema also apply here (notice that the set $\mathsf{Mon}_{\mathbb{N}}(X)$ is countable, so that the supremum is well-defined).

However, we are interested mostly in *fully*-continuous homomorphism, as these commute with the evaluation of fixed-point logic, and these are not guaranteed for AbsPoly($K, \mathbb{N}, X$). A counterexample for the case $K = \mathbb{B}$ (i.e., in $\mathbb{S}(X)$), is given by the chain $1 \geq x \geq x^2 \geq x^3 \geq \dots$ with infimum $\bigsqcap_{i<\omega} x_i = 0$ in $\mathbb{S}(x)$. For the homomorphism h_f induced by the assignment $f(x) = 1$, we get $\bigsqcap_i h_f(x^i) = \bigsqcap_i 1 = 1$ and $h_f(\bigsqcap_i x_i) = 0$.

Infinite Coefficient Semirings

The arguments in the above universality proof are no longer applicable for infinite K, as the antichain representations need not be finite. Indeed, the following example uses an infinite ascending chain in the Viterbi semiring to obtain an infinite absorptive polynomial in AbsPoly($\mathbb{V}, \mathbb{N}^\infty, \{x\}$):

$$\frac{1}{2}x + \frac{2}{3}x^2 + \frac{3}{4}x^3 + \frac{4}{5}x^4 + \cdots + 1x^\infty.$$

Instead, we require that the natural order of the target semiring L is completely distributive (as complete lattice). We can then prove full continuity of the homomorphism $h_f\colon$ AbsPoly(K, E, X) $\to L$ by reasoning about choice functions on monomials and exploiting the finite number of indeterminates. The counterexample in Remark 3.7.15 also applies here, so we prove the universal property for exponents in either $\mathbb{B}$ or $\mathbb{N}^\infty$. Since both are well-founded (so that minimal exponents exist) and closed under suprema, the reasoning via choice functions works for both.

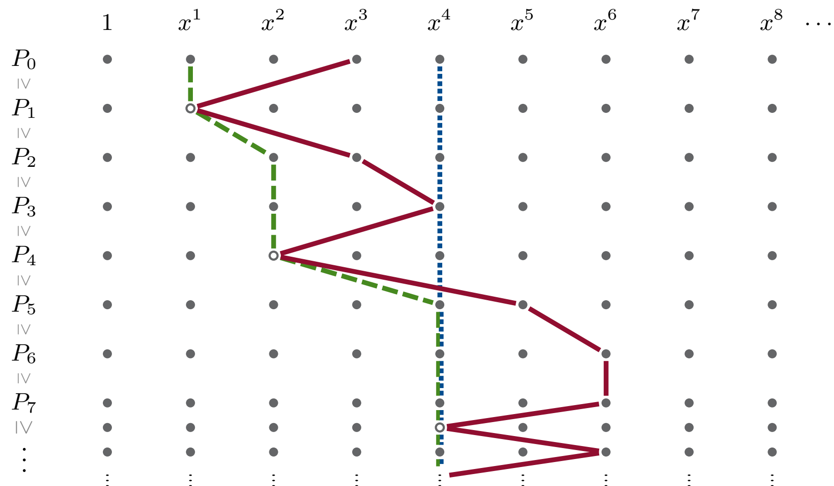

Figure 3.5: Schematic illustration for the proof of Theorem 3.7.16 (with $X = \{x\}$). The choice function g is shown in red, the monotone function g' in green (dashed), and the constant function m in blue (dotted), assuming that g continues to alternate between x^4 and x^6. Empty dots mark the monomials $g(Q^*)$ used to define g'.

Theorem 3.7.16 (Universality II)**.** *Let K be an absorptive, fully-continuous semiring, and $E \in \{\mathbb{B}, \mathbb{N}^\infty\}$. Let further L be an absorptive, fully-continuous semiring whose natural order is completely distributive.*

Then for each fully-continuous homomorphism $h\colon K \to L$ and assignment $f\colon X \to L$, there is a unique fully-continuous semiring homomorphism $h_f\colon \mathrm{AbsPoly}(K, E, X) \to L$ that simultaneously extends both h and f.

Proof. We define h_f as in the proof of Theorem 3.7.13, by lifting f to monomials and setting

$$h_f(P) = \bigsqcup_{m \in \mathsf{Mon}} f(m) \cdot h(P(m)).$$

The proof that h_f is unique and commutes with the semiring operations and suprema did not make use of the finiteness of K, so it still applies here (for $E = \mathbb{B}$ and $E = \mathbb{N}^\infty$).

Let $C \subseteq \mathrm{AbsPoly}(K, E, X)$ be a chain. It remains to show that $h_f(\bigsqcap C) = \bigsqcap h_f(C)$. One one side, we have

$$h_f(\bigsqcap_{P \in C} P) = \bigsqcup_{m \in \mathsf{Mon}} f(m) \cdot h\Big(\bigsqcap_{P \in C} P(m)\Big) = \bigsqcup_{m \in \mathsf{Mon}} \bigsqcap_{P \in C} f(m) \cdot h(P(m)),$$

and for the other side we apply complete distributivity:

$$\bigsqcap_{P \in C} h_f(P) = \bigsqcap_{P \in C} \bigsqcup_{m \in \mathsf{Mon}} f(m) \cdot h(P(m)) = \bigsqcup_{g \in \mathcal{G}} \underbrace{\bigsqcap_{P \in C} f(g(P)) \cdot h(P(g(P)))}_{=:\mathrm{val}(g)},$$

where $\mathcal{G}$ is the set of all functions $g\colon C \to \mathsf{Mon}$. To prove equality of the two suprema, we have to show that the supremum over $\mathcal{G}$ can be restricted to constant functions $g = m$

without affecting its value. We first show that monotone functions suffice, and then restrict these further to constant functions.

Monotone choice functions. We say that $g\colon C \to \mathsf{Mon}$ is *monotone in an indeterminate* $x \in X$, if $P \geq P'$ implies $g(P)(x) \leq g(P')(x)$. (Notice the inverse order on exponents.) If g is monotone in all $x \in X$, then g is monotone as a function. Let $X = \{x_1, \dots, x_n\}$. We prove by induction on $k \leq n$ that the following claim holds: for each $g \in \mathcal{G}$, there is $g' \in \mathcal{G}$ such that g' is monotone in $x_1, \dots, x_k$ and $\mathrm{val}(g) \leq \mathrm{val}(g')$.

Assume that the claim holds for $k < n$ (initially $k = 0$ and $g' = g$). We can then assume that g is monotone in $x_1, \dots, x_k$. We now define $g' \in \mathcal{G}$ so that g' is monotone in $x_1, \dots, x_k$ and also x_{k+1}. Intuitively, for each P we look at the behaviour of g for all smaller $Q \leq P$ (i.e., coming after P in the chain) and choose a Q with minimal exponent of x_{k+1}. We then set g' to skip all the monomials g chooses between P and Q, and immediately jump to $g(Q)$. See Fig. 3.5 for a schematic illustration.

Formally, let $e_{\min} = \min\{g(Q)(x_{k+1}) \mid Q \in C_{\leq P}\}$ be the minimal exponent of x_{k+1} that g chooses in the remaining chain $C_{\leq P}$. Let further $\mathbf{Q}_{\min}^{\leq P} = \{Q \in C_{\leq P} \mid g(Q)(x) = e_{\min}\}$ be the set of polynomials for which this minimal exponent is assumed. Now choose $Q^* \in \mathbf{Q}_{\min}^{\leq P}$ for which all exponents $g(Q)(x_i)$ for $1 \leq i \leq k$ are minimal. This minimum exists since g is monotone in $x_1, \dots, x_k$. (In fact, we can choose the maximum of $\mathbf{Q}_{\min}^{\leq P}$, i.e. the first element of $\mathbf{Q}_{\min}^{\leq P}$ appearing in the chain after P, if such a maximum exists.) This choice may not be unique, but this does not affect our argument. We then define $g'(P) \coloneqq g(Q^*)$.

Then g' is monotone in x_{k+1} by construction. We show that g' is also monotone in x_i for all $i \leq k$. Let $P_1, P_2 \in C$ with $P_1 \geq P_2$. Let further $Q_1^* \in \mathbf{Q}_{\min}^{\leq P_1}$ and $Q_2^* \in \mathbf{Q}_{\min}^{\leq P_2}$ so that $g'(P_1) = g(Q_1^*)$ and $g'(P_2) = g(Q_2^*)$. We proceed by case distinction.

- Case 1: $g'(P_1)(x_{k+1}) < g'(P_2)(x_{k+1})$. Then $Q_1^* > P_2$, as otherwise Q_1^* would be a better choice than Q_2^* for $g(P_2)$. Hence $Q_1^* > Q_2^*$ and thus $g'(P_1)(x_i) \leq g'(P_2)(x_i)$ by monotonicity of g in x_i.
- Case 2: $g'(P_1)(x_{k+1}) = g'(P_2)(x_{k+1})$. Then we have both $Q_1^*, Q_2^* \in \mathbf{Q}_{\min}^{\leq P_1}$, and since Q_1^* minimizes the exponents of $x_1, \dots, x_k$, we again have $g'(P_1)(x_i) \leq g'(P_2)(x_i)$.

Hence g' is monotone in $x_1, \dots, x_{k+1}$. It remains to show that $\mathrm{val}(g') \geq \mathrm{val}(g)$. For each $P \in C$ and the associated Q^* with $g'(P) = g(Q^*)$, we have

$$\begin{aligned} f(g'(P)) \cdot h(P(g'(P))) &= f(g(Q^*)) \cdot h(P(g(Q^*))) \\ &\geq f(g(Q^*)) \cdot h(Q^*(g(Q^*))) \geq \mathrm{val}(g). \end{aligned}$$

Since this holds for each P, we have $\mathrm{val}(g') \geq \mathrm{val}(g)$. We can thus restrict the supremum over $\mathcal{G}$ to monotone functions without affecting its value.

Constant choice functions. Now let $g \in G$ be monotone and consider the family $(g(P))_{P \in C}$ of chosen monomials. Since g is monotone, this is a chain. Let m be the monomial with $m(x) = \bigsqcup_{P \in C} g(P)(x)$ for all $x \in X$. Then $g(P) \succeq m$ for all $P \in C$ and hence, since P is absorption-compatible, also $P(g(P)) \leq P(m)$. Using fully continuity of L and the Chain splitting lemma (Lemma 3.2.8), it is now easy to prove that $\mathrm{val}(g) \leq \mathrm{val}(m)$:

$$\begin{aligned}
\prod_{P \in C} f(g(P)) \cdot h(P(g(P))) &\leq \prod_{P \in C} f(g(P)) \cdot h(P(m)) \\
&= \prod_{P \in C} f(g(P)) \;\cdot\; \prod_{P \in C} h(P(m)) \\
&= \prod_{x \in X} f(x)^{\bigsqcup_{P \in C} g(P)(x)} \;\cdot\; \prod_{P \in C} h(P(m)) \\
&= f(m) \cdot \prod_{P \in C} h(P(m)) \\
&= \prod_{P \in C} f(m) \cdot h(P(m)) = \mathrm{val}(m).
\end{aligned}$$

This proves equality of the two suprema, hence h_f is fully continuous. □

Chapter 4

Fixed-Point Computation

In this chapter, we show how one can efficiently compute least and, most importantly, *greatest* solutions of polynomial equation systems over absorptive, fully-continuous semirings. This is motivated by semiring semantics for fixed-point logics, discussed in Chapter 5, which replaces Boolean evaluation of formulae by computations in absorptive semirings. From this point of view, a formula is essentially a polynomial expression over a semiring, and fixed-point formulae evaluate to least or greatest solutions of polynomial equation systems. Since we are interested in infinite semirings, such as the tropical semiring and absorptive polynomials, it is not clear a priori how these solutions can be (efficiently) computed. This question may also be of independent interest, so we formulate the results in this chapter in terms of polynomial equation systems; the connection to logic is made explicit in Chapter 5.

The textbook approach to compute fixed points is the fixed-point iteration (also called Kleene iteration): start by setting all indeterminates to the smallest (or greatest) semiring value, then repeatedly evaluate the equations to obtain new values for all indeterminates. In the Boolean setting, this terminates in at most n steps on n indeterminates (due to monotonicity), but we are also interested in larger and especially infinite semirings such as the tropical semiring[1] $\mathbb{T} = (\mathbb{R}^{\infty}_{\geq 0}, \min, +_{\mathbb{R}}, \infty, 0)$. Several techniques have been developed to compute least solutions. Hopkins and Kozen [HK99] have considered systems of inequalities over commutative Kleene algebras, which include all ω-continuous idempotent semirings. They define a faster iteration scheme based on (formal) differentials and show that it converges to the least solution in $\mathcal{O}(3^n)$ steps. More recently, Esparza, Kiefer, and Luttenberger [EKL10] built on this idea to generalise Newton's method, a well-known approximation scheme for zeroes of real-valued functions, to ω-continuous semirings. This works surprisingly well and yields an iteration scheme for least solutions over a wide variety of semirings. In the idempotent case, this *Newton iteration* is in fact identical to the iteration of Hopkins and Kozen, and a more detailed analysis using derivation trees shows that it converges in just n steps [EKL10]. Further techniques have been developed for linear systems (see, e.g., [GM08]).

Our goal is to complement the results in [HK99, EKL10] by also computing *greatest* solutions, as both least and greatest fixed points are required to define a semiring semantics for fixed-point logics. Since this semantics works best in absorptive semirings (cf. Section 5.2), we focus on absorptive, fully-continuous semirings and present two approaches to compute greatest fixed points.

[1] We use $+_{\mathbb{R}}$ for the addition on $\mathbb{R}$ to distinguish it from the semiring operation $+$.

Example. Consider the following graph whose edges are annotated by cost values in the tropical semiring. A natural example of a greatest fixed point is the minimal cost of an infinite path. This corresponds to the greatest solution of the equation system given on the right, where each node is represented by an indeterminate and costs appear as coefficients (notice that the right-hand sides are indeed polynomial expressions in terms of the semiring operations).

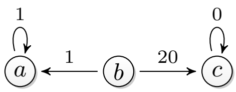

$$X_a = 1 +_{\mathbb{R}} X_a$$
$$X_b = \min(1 +_{\mathbb{R}} X_a, 20 +_{\mathbb{R}} X_c)$$
$$X_c = 0 +_{\mathbb{R}} X_c$$

When we speak of *least* or *greatest* solutions, we always refer to the natural order of the semiring. In the case of the tropical semiring, this is the inverse of the standard order, so $\infty <_{\mathbb{T}} 20 <_{\mathbb{T}} 1 <_{\mathbb{T}} 0$. While the least solution of the above system is trivially $X_a = X_b = X_c = \infty$, the fixed-point iteration for the greatest solution is infinite:

$$\begin{pmatrix}0\\0\\0\end{pmatrix} \mapsto \begin{pmatrix}1\\1\\0\end{pmatrix} \mapsto \begin{pmatrix}2\\2\\0\end{pmatrix} \mapsto \begin{pmatrix}3\\3\\0\end{pmatrix} \mapsto \cdots \mapsto \begin{pmatrix}20\\20\\0\end{pmatrix} \mapsto \begin{pmatrix}21\\20\\0\end{pmatrix} \mapsto \begin{pmatrix}22\\20\\0\end{pmatrix} \mapsto \begin{pmatrix}23\\20\\0\end{pmatrix} \mapsto \cdots$$

and converges to the greatest solution $X_a = \infty$, $X_b = 20$, and $X_c = 0$.

Main result. The essential idea to compute such solutions is that greatest fixed points are composed of two parts: a cyclic part that is repeated indefinitely (the loop at a or c) and a reachability part to get to the cycle (the edges from b). As both parts can consist of at most n nodes, all information we need is already present after n steps of the fixed-point iteration; we can use this information to abbreviate the iteration. The formal proof of this observation is based on (infinite) derivation trees, inspired by the derivation trees in the analysis of the Newton iteration [EKL10] and infinite strategy trees in [DGNT21]. We show that these trees provide an alternative description of the fixed-point iteration; a careful analysis of the shape of the derivation trees then leads to our main result:

Theorem 4.3.1 (informal)**.** *Let F be the operator induced by a polynomial equation system in n indeterminates over an absorptive, fully-continuous, commutative semiring. We can compute in a polynomial number of semiring operations:*

- *the least solution: $F^n(\mathbf{0})$,*
- *the greatest solution: $F^n(\,F^n(\mathbf{1})^\infty\,)$.*

The result for least solutions has already been shown in [EKL11] as part of the derivation tree analysis of the Newton iteration. In our context, it is a byproduct of the more involved proof for greatest solutions.

Symbolic computation. Our second approach is a technique to eliminate indeterminates one by one, based on the ideas of Hopkins and Kozen [HK99]. We show how their approach can be extended to compute greatest fixed points in the semiring $\mathbb{S}^\infty(X)$ of generalised absorptive polynomials. For other semirings, we can work symbolically by first abstracting coefficients to fresh indeterminates, computing the solution in $\mathbb{S}^\infty(X)$, and then undoing the abstraction again.

Both approaches presented in this chapter have previously been published in [Naa21a, Naa21b], with the exception of the contents in Sections 4.3.4, 4.3.5, and 4.5.

4.1 Polynomial Equation Systems

We begin with some notation. Throughout this chapter, we fix a finite set $\boldsymbol{X} = \{X_1, \dots, X_\ell\}$ of ℓ pairwise different indeterminates. We write polynomials in $K[\boldsymbol{X}]$ (or $\mathrm{Poly}^\omega(K, \mathbb{N}, \boldsymbol{X})$ in the notation of Chapter 3) as sums $P = \sum_{i=1}^{k} c_i \cdot m_i$ with indeterminates in $\boldsymbol{X}$ and coefficients c_i in $K \setminus \{0\}$. Recall that we represent monomials as mappings $m_i \colon \boldsymbol{X} \to \mathbb{N}$. Slightly abusing notation, we write $m \in P$ if there is an i with $m = m_i$, and $c \cdot m \in P$ if additionally $c = c_i$. We may write $P(X_1, \dots, X_\ell)$ to make the indeterminates explicit. Then $P(a_1, \dots, a_\ell) \in K$ is the semiring value obtained by instantiating each indeterminate X_i by $a_i \in K$ and evaluating the resulting expression in K.

Tuples are denoted by bold symbols: $\boldsymbol{a} = (a_1, \dots, a_\ell)$. In particular, $\mathbf{0} = (0, \dots, 0)$ and $\mathbf{1} = (1, \dots, 1)$. To ease the presentation, we often avoid numbered indices and instead index tuples by indeterminates. That is, for a tuple $\boldsymbol{a} = (a_1, \dots, a_\ell)$ and an indeterminate $X \in \boldsymbol{X}$, we write $\boldsymbol{a}_X$ for the entry a_i such that $X_i = X$. We lift homomorphisms and the infinitary power to tuples by applying them to all entries, i.e., $(\boldsymbol{a}^\infty)_X = (\boldsymbol{a}_X)^\infty$ for $X \in \boldsymbol{X}$.

Definition 4.1.1. A *polynomial equation system* $\mathcal{E}$ over a semiring K and indeterminates $\boldsymbol{X} = \{X_1, \dots, X_\ell\}$ is a family of equations $\mathcal{E} \colon \big(X_i = P_i(X_1, \dots, X_\ell)\big)_{1 \le i \le \ell}$ with $P_i \in K[\boldsymbol{X}]$. We associate with $\mathcal{E}$ the operator

$$F_\mathcal{E} \colon K^\ell \to K^\ell, \qquad F_\mathcal{E}(a_1, \dots, a_\ell)_X = P_X(a_1, \dots, a_\ell), \text{ for each } X \in \boldsymbol{X}.$$

The least (greatest) solution to $\mathcal{E}$ is the least (greatest) fixed point of $F_\mathcal{E}$. We drop the index and simply write F if $\mathcal{E}$ is clear from the context.

Notice that these are square systems, with the number of equations equal to the number of indeterminates. We recall the introductory example in the tropical semiring (where semiring addition is min and semiring multiplication is $+_\mathbb{R}$). Using $\boldsymbol{X} = \{X_a, X_b, X_c\}$, we refer to the polynomial equation system as $(X = P_X)_{X \in \boldsymbol{X}}$. For example, $P_{X_b} = \min(1 +_\mathbb{R} X_a, 20 +_\mathbb{R} X_c)$ consists of the two coefficient-monomial pairs $1 +_\mathbb{R} X_a$ and $20 +_\mathbb{R} X_c$.

Whenever addition and multiplication are fully continuous, then so are the operators induced by polynomials. It follows that least and greatest solutions of polynomial equation systems are preserved by fully-continuous homomorphisms.

Lemma 4.1.2. *Let $h \colon K_1 \to K_2$ be a fully-continuous homomorphism on absorptive, fully-continuous semirings. Let $\mathcal{E} \colon (X_i = P_i)_{1 \le i \le \ell}$ be a polynomial equation system over K_1.*

Let $h(\mathcal{E}) \colon (X_i = h(P_i))_{1 \le i \le \ell}$ be the polynomial equation system resulting from $\mathcal{E}$ by applying h to all coefficients. Then $\mathbf{lfp}(F_{h(\mathcal{E})}) = h(\mathbf{lfp}(F_\mathcal{E}))$ and $\mathbf{gfp}(F_{h(\mathcal{E})}) = h(\mathbf{gfp}(F_\mathcal{E}))$.

Proof. Since h commutes with the semiring operations, we have $h(F_{\mathcal{E}}(\boldsymbol{a})) = F_{h(\mathcal{E})}(h(\boldsymbol{a}))$ for all $\boldsymbol{a} \in K^\ell$, and by induction also $h(F^n_{\mathcal{E}}(\boldsymbol{a})) = F^n_{h(\mathcal{E})}(h(\boldsymbol{a}))$ for all $n < \omega$. Recall that the semiring operations are fully continuous in K_1 and K_2. Hence $F_{\mathcal{E}}$ and $F_{h(\mathcal{E})}$ are fully continuous as well, and by Kleene's fixed-point theorem and the continuity of h, we get:

$$h(\mathbf{lfp}(F_{\mathcal{E}})) = h\Big(\bigsqcup_{n<\omega} F^n_{\mathcal{E}}(\mathbf{0})\Big) = \bigsqcup_{n<\omega} h\big(F^n_{\mathcal{E}}(\mathbf{0})\big) = \bigsqcup_{n<\omega} \big(F^n_{h(\mathcal{E})}(\mathbf{0})\big) = \mathbf{lfp}(F_{h(\mathcal{E})}).$$

The proof for greatest solutions is symmetric. □

4.2 Derivation Trees

Inspired by the analysis of the Newton iteration [EKL10], we use derivation trees to describe the behaviour of polynomial equation systems. For the intuition behind this notion, think of a polynomial system as a formal grammar: indeterminates are nonterminal symbols, coefficients are terminal symbols, and each monomial in P_X gives rise to a production rule for X. We then consider derivation trees of this grammar in the usual sense, except that we ignore the order of children (we use commutative semirings) and allow infinite derivations.

Example 4.2.1. Consider the following system with coefficients $a, b, c, d \in K$:

$$X = aXY + b$$
$$Y = cY^3 + dX$$

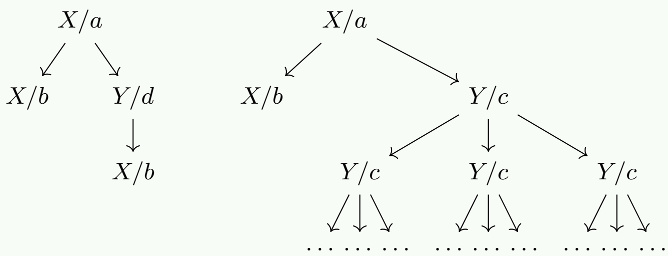

We view this system as a formal grammar with nonterminals X, Y. One possible derivation is then $X \to aXY \to abY \to abdX \to abdb$, and the left tree is the corresponding derivation tree. Each step corresponds to a node, labelled by the chosen nonterminal and the terminal symbol (i.e., the coefficient) of the chosen production rule. We associate with each derivation tree its yield, which is the product over all coefficients appearing in this tree, and hence the commutative image of the generated word. In this case, the yield is ab^2d. The tree on the right indicates an infinite derivation tree with yield abc^∞.

4.2.1 Definition

We formally define derivation trees as follows. We note that our definition is conceptually identical to [EKL10], but our yield labelling is slightly different.

Definition 4.2.2. A *derivation tree* $T = (V, E, \mathsf{var}, \mathsf{yd})$ over a semiring K and indeterminates $\boldsymbol{X}$ is a (possibly infinite) finitely-branching tree (V, E) with node labellings $\mathsf{var}\colon V \to \boldsymbol{X}$ and $\mathsf{yd}\colon V \to K$, the *yield* of v. We denote the root node as ε and refer to trees with $\mathsf{var}(\varepsilon) = X$ as derivation trees *from* X. For convenience, we often write $v \in T$ instead of $v \in V$ and refer to v with $\mathsf{var}(v) = X$ as an *occurrence* of X in T.

We associate with each node the monomial $\mathsf{mon}(v) = \prod_{w \in vE} \mathsf{var}(w)$ composed of its children's indeterminates. We say that T is *compatible with the system* $(X = P_X)_{X \in \boldsymbol{X}}$ if for each node, $\mathsf{yd}(v) \cdot \mathsf{mon}(v) \in P_{\mathsf{var}(v)}$. The set of all derivation trees from X that are compatible with the system $\mathcal{E}$ is denoted $\mathcal{T}(\mathcal{E}, X)$.

Given an equation system $\mathcal{E}\colon (X = P_X)_{X \in \boldsymbol{X}}$ and an indeterminate X, a derivation tree $T \in \mathcal{T}(\mathcal{E}, X)$ first chooses from the equation $X = P_X$ a monomial $\mathsf{mon}(\varepsilon)$ together with its coefficient $\mathsf{yd}(\varepsilon)$. On the next level, it then makes analogous choices for all indeterminates occurring in $\mathsf{mon}(\varepsilon)$, where the exponent specifies how often an indeterminate occurs. The leaves v of such a derivation tree (if they exist) have $\mathsf{mon}(v) = 1$ and correspond to absolute coefficients in one of the equations.

We define the yield of an entire tree as the combined yield of all nodes. Since trees can be infinite, we use infinite products $\widehat{\prod}$ (and later also sums $\widehat{\sum}$) of Theorem 3.4.2 that can be defined in all absorptive, fully-continuous semirings.

Definition 4.2.3. The *yield* of a derivation tree $T = (V, E, \mathsf{var}, \mathsf{yd})$ over K is the (possibly infinite) product $\mathsf{yd}(T) = \widehat{\prod}_{v \in V} \mathsf{yd}(v)$.

We compare the yields of two trees by counting occurrences of coefficients or, equivalently, of monomials from the polynomial equation system.

Definition 4.2.4. Let $\mathcal{E}$ be an equation system over $\boldsymbol{X}$, let $Y \in \boldsymbol{X}$ and $m \in P_Y$. For derivation trees $T \in \mathcal{T}(\mathcal{E}, X)$, we define

$$|T|_{m,Y} \coloneqq \big|\{v \in T \mid \mathsf{mon}(v) = m,\ \mathsf{var}(v) = Y\}\big| \in \mathbb{N}^\infty$$

as the number of occurrences of $m \in P_Y$ in T, setting $|T|_{m,Y} \coloneqq \infty$ for infinitely many occurrences. Notice that we use pairs (m, Y) to unambiguously refer to $m \in P_Y$, as the same monomial m may occur in several polynomials of $\mathcal{E}$.

Lemma 4.2.5 (Yield comparison)**.** *Given a polynomial system $\mathcal{E}\colon (X = P_X)_{X \in \boldsymbol{X}}$ over an absorptive, fully-continuous semiring, and trees $T, T' \in \mathcal{T}(\mathcal{E}, X)$,*

(1) *if $|T|_{m,Y} \geq |T'|_{m,Y}$ for all $m \in P_Y$, $Y \in \boldsymbol{X}$, then $\mathsf{yd}(T) \leq \mathsf{yd}(T')$;*

(2) *if $|T|_{m,Y} = 0$ implies $|T'|_{m,Y} = 0$ for all $m \in P_Y$, $Y \in \boldsymbol{X}$, then $\mathsf{yd}(T)^\infty \leq \mathsf{yd}(T')^\infty$.*

Proof. We use associativity of infinite products to group the yields by monomials. For each $m \in P_Y$, let $c_{m,Y} \in K$ be its coefficient so that $c \cdot m \in P_Y$. Then,

$$\mathsf{yd}(T) = \widehat{\prod_{v \in T}} \mathsf{yd}(v) = \prod_{\substack{Y \in \boldsymbol{X},\\ m \in P_Y}} \widehat{\prod_{\substack{v \in T,\\ \mathsf{var}(v) = Y,\\ \mathsf{mon}(v) = m}}} \mathsf{yd}(v) = \prod_{\substack{Y \in \boldsymbol{X},\\ m \in P_Y}} (c_{m,Y})^{|T|_{m,Y}},$$

and the same applies to T'. By monotonicity and absorption, larger exponents lead to smaller values of the product, hence (1) holds. Claim (2) follows by applying the infinitary power: $\mathsf{yd}(T)^\infty = \prod_{m,Y} (c_{m,Y})^{\infty \cdot |T|_{m,Y}}$, where $\infty \cdot |T|_{m,Y}$ is either ∞ (if $|T|_{m,Y} > 0$) or 0. □

4.2.2 Fixed-Point Iteration via Trees

As a first step towards our main result, this section shows that we can express least and greatest solutions in terms of the yields of derivation trees. Notice that a single derivation tree does not correspond to a solution of the equation system, but only to (the derivation of) a single term in the solution. We thus consider the sum over all derivation trees.

For least solutions, this was already shown in [EKL10]. Here we are mostly concerned with the proof for greatest solutions, as this is much more involved due to the trees being infinite.

Theorem 4.2.6. *Let K be an absorptive, fully-continuous semiring. Let $\mathcal{E}\colon (X = P_X)_{X\in\boldsymbol{X}}$ be a polynomial equation system over K. Then for each $X \in \boldsymbol{X}$,*

$$\mathbf{lfp}(F_{\mathcal{E}})_X = \widehat{\sum_{\substack{T\in\mathcal{T}(\mathcal{E},X),\\ T \text{ is finite}}}} \mathsf{yd}(T), \qquad \mathbf{gfp}(F_{\mathcal{E}})_X = \widehat{\sum_{T\in\mathcal{T}(\mathcal{E},X)}} \mathsf{yd}(T).$$

We recall that summation is equivalent to supremum (cf. Section 3.4.2). Here and in the following, we use summation in reminiscence of the general, non-idempotent case (cf. [EKL10]). Towards a proof, we first observe that it suffices to prove Theorem 4.2.6 for the universal semiring. More precisely, the coefficients of the system $\mathcal{E}$ are from $K = \mathbb{S}^\infty(\boldsymbol{A})$, for some finite indeterminate set $\boldsymbol{A}$. We use lowercase letters $\boldsymbol{A} = \{a, b, c, \dots\}$ to clearly distinguish between the indeterminates $\boldsymbol{X}$ of the polynomial system and the indeterminates of $\mathbb{S}^\infty(\boldsymbol{A})$ that we use as coefficients.

Lemma 4.2.7. *If Theorem 4.2.6 holds for $K = \mathbb{S}^\infty(\boldsymbol{A})$ for all finite indeterminate sets $\boldsymbol{A}$, then it also holds for any absorptive, fully-continuous semiring K.*

Proof. Let K be an absorptive, fully-continuous semiring. Given $\mathcal{E}\colon (X = P_X)_{X\in\boldsymbol{X}}$ over K, we construct a symbolic abstraction $\mathcal{E}'\colon (X = P'_X)_{X\in\boldsymbol{X}}$ over $\mathbb{S}^\infty(\boldsymbol{A})$. To this end, we choose $\boldsymbol{A}$ sufficiently large and let P'_X result from P_X by replacing all coefficients with pairwise different indeterminates from $\boldsymbol{A}$. Let $h\colon \boldsymbol{A} \to K$ be the corresponding instantiation of these indeterminates that reverses this process, so that $h(\mathcal{E}') = \mathcal{E}$. This mapping extends to a fully-continuous homomorphism $h\colon \mathbb{S}^\infty(\boldsymbol{A}) \to K$, so by Lemma 4.1.2, we have for each X,

$$\mathbf{gfp}(F_{\mathcal{E}})_X \;=\; \mathbf{gfp}(F_{h(\mathcal{E}')})_X \overset{(4.1.2)}{=} h(\mathbf{gfp}(F_{\mathcal{E}'}))_X \;=\; h\Big(\widehat{\sum_{T'\in\mathcal{T}(\mathcal{E}',X)}} \mathsf{yd}(T')\Big).$$

Notice that the structure of derivation trees $\mathcal{T}(\mathcal{E}, X)$ only depends on the monomials occurring in $\mathcal{E}$, but not on the coefficients. Thus, the derivation trees for $\mathcal{E}$ and $\mathcal{E}'$ are identical up to the labelling yd. Given a tree $T \in \mathcal{T}(\mathcal{E}, X)$, it holds in particular that

$$\mathsf{yd}(T) \;=\; \widehat{\prod_{v\in T}} \mathsf{yd}(v) \;=\; \widehat{\prod_{v\in T'}} h(\mathsf{yd}(v)) \overset{(3.4.3)}{=} h\Big(\widehat{\prod_{v\in T'}} \mathsf{yd}(v)\Big) \;=\; h\big(\mathsf{yd}(T')\big),$$

where $T' \in \mathcal{T}(\mathcal{E}', X)$ is the tree corresponding to T (so that only yd is changed according to the coefficients in $\mathcal{E}'$). By using the one-to-one correspondence between trees $T \in \mathcal{T}(\mathcal{E}, X)$

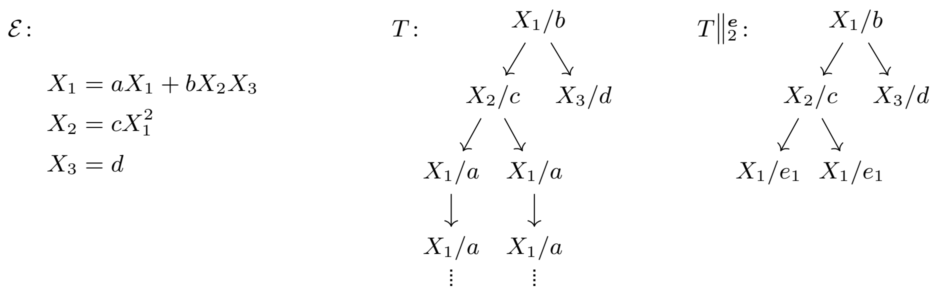

Figure 4.1: A derivation tree $T \in \mathcal{T}(\mathcal{E}, X_1)$ and its $(2, \boldsymbol{e})$-truncation, with nodes represented by their labels $\mathsf{var}(v)/\mathsf{yd}(v)$. The trees have yield $\mathsf{yd}(T) = a^\infty bcd$ and $\mathsf{yd}(T\|^{\boldsymbol{e}}_2) = e_1^2 bcd$.

and $T' \in \mathcal{T}(\mathcal{E}', X)$, we can conclude

$$\mathbf{gfp}(F_{\mathcal{E}})_X = h\Big(\widehat{\sum_{T' \in \mathcal{T}(\mathcal{E}',X)}} \mathsf{yd}(T') \Big) \overset{(3.4.3)}{=} \widehat{\sum_{T' \in \mathcal{T}(\mathcal{E}',X)}} h\big(\mathsf{yd}(T')\big) = \widehat{\sum_{T \in \mathcal{T}(\mathcal{E},X)}} \mathsf{yd}(T).$$

The proof for $\mathbf{lfp}(F_{\mathcal{E}})$ is symmetric. □

For the remaining section, we fix a polynomial equation system $\mathcal{E} \colon (X = P_X)_{X \in \boldsymbol{X}}$ over an absorptive, fully-continuous semiring K, with induced operator F. The proof of Theorem 4.2.6 proceeds by induction on the fixed-point iterations $\big(F^n(\mathbf{0})\big)_{n<\omega}$ and $\big(F^n(\mathbf{1})\big)_{n<\omega}$, but requires some preparation. The idea is that if complete derivation trees correspond to the fixed points, their prefixes should correspond to the steps of the iteration. These prefixes are defined by simply cutting off the derivation trees at a certain depth and assigning a specific yield to the nodes at the cut-off depth (eventually, we will simply assign 0 for the least and 1 for the greatest fixed point).

Definition 4.2.8. Let $T = (V, E, \mathsf{var}, \mathsf{yd}) \in \mathcal{T}(\mathcal{E}, X)$, $n < \omega$ and $\boldsymbol{b} \in K^\ell$. Let $V_{\leq n} \subseteq V$ be the nodes at depth $\leq n$. We define the $(n, \boldsymbol{b})$-*truncation* of T as

$$T\|^{\boldsymbol{b}}_n \coloneqq (V_{\leq n},\, E \cap V_{\leq n}^2,\, \mathsf{var},\, \mathsf{yd}'), \quad \mathsf{yd}'(v) = \begin{cases} \boldsymbol{b}_{\mathsf{var}(v)}, & v \text{ at depth } n, \\ \mathsf{yd}(v), & \text{otherwise.} \end{cases}$$

This defines a derivation tree (compatible with $\mathcal{E}$ except for its leaves) and we define $\mathsf{mon}(v)$ and $\mathsf{yd}(T\|^{\boldsymbol{b}}_n)$ as in Definitions 4.2.2 and 4.2.3 (cf. Fig. 4.1).

The following, mostly technical, lemma establishes the general connection between truncations of derivation trees and the fixed-point iteration.

Lemma 4.2.9 (Tree iteration)**.** *Given $\boldsymbol{b} \in K^\ell$, we have $F^n(\boldsymbol{b})_X = \widehat{\sum}_{T \in \mathcal{T}(\mathcal{E},X)} \mathsf{yd}(T\|^{\boldsymbol{b}}_n)$ for all $n < \omega$ and $X \in \boldsymbol{X}$.*

Proof. Induction on n. For $n = 0$, we trivially have $F^n(\boldsymbol{b})_X = \mathsf{yd}(T\|^{\boldsymbol{b}}_0) = \boldsymbol{b}_X$ for all derivation trees $T \in \mathcal{T}(\mathcal{E}, X)$. For the induction step, assume that $F^n(\boldsymbol{b})_X = \widehat{\sum}_{T\in\mathcal{T}(\mathcal{E},X)} \mathsf{yd}(T\|^{\boldsymbol{b}}_n)$ for all $X \in \boldsymbol{X}$. We have to show that

$$F^{n+1}(\boldsymbol{b})_X = P_X(F^n(\boldsymbol{b})) \overset{!}{=} \widehat{\sum_{T\in\mathcal{T}(\mathcal{E},X)}} \mathsf{yd}(T\|^{\boldsymbol{b}}_{n+1}).$$

To simplify notation, let $\boldsymbol{a}$ be the tuple with $\boldsymbol{a}_X = \widehat{\sum}_{T\in\mathcal{T}(\mathcal{E},X)} \mathsf{yd}(T\|^{\boldsymbol{b}}_n)$. We can rewrite the left-hand side as follows (recall that $m(X)$ denotes the exponent of X in m):

$$\begin{aligned} P_X(\boldsymbol{a}) &= \sum_{c\cdot m\in P_X} (c\cdot m)(\boldsymbol{a}) = \sum_{c\cdot m\in P_X} \Big(c \cdot \prod_{X\in\boldsymbol{X}} \boldsymbol{a}_X^{m(X)}\Big) \\ &= \sum_{c\cdot m\in P_X} c \cdot \big(\underbrace{\boldsymbol{a}_{X_1}\cdot\boldsymbol{a}_{X_1}\cdots\boldsymbol{a}_{X_1}}_{m(X_1)\text{ times}} \cdot \ \ldots \ \cdot \underbrace{\boldsymbol{a}_{X_\ell}\cdots\boldsymbol{a}_{X_\ell}}_{m(X_\ell)\text{ times}}\big) \end{aligned}$$

Notice that the unfolded product is finite, since m only has finite exponents. Since finite multiplication distributes over the infinite sums $\boldsymbol{a}_{X_1}, \boldsymbol{a}_{X_2}, \ldots$ (by Theorem 3.4.2), the product $\boldsymbol{a}_{X_1}\cdot\boldsymbol{a}_{X_1}\cdots\boldsymbol{a}_{X_\ell}$ can be rewritten as infinite sum:

$$\begin{aligned} &= \widehat{\sum}\left\{ c\cdot \prod_{1\le i\le \ell} \mathsf{yd}(T_{i,1}\|^{\boldsymbol{b}}_n)\cdots\mathsf{yd}(T_{i,m(X_i)}\|^{\boldsymbol{b}}_n) \;\middle|\; \begin{array}{l} c\cdot m\in P_X \text{ and all possible} \\ \text{choices of trees } T_{i,j}\in\mathcal{T}(\mathcal{E},X_i) \end{array} \right\} \\ &= \widehat{\sum}\left\{ \mathsf{yd}(T\|^{\boldsymbol{b}}_{n+1}) \mid c\cdot m\in P_X,\, T\in\mathcal{T}(\mathcal{E},X) \text{ with } \mathsf{mon}(\varepsilon)=m,\, \mathsf{yd}(\varepsilon)=c \right\} \\ &= \widehat{\sum}\left\{ \mathsf{yd}(T\|^{\boldsymbol{b}}_{n+1}) \mid T\in\mathcal{T}(\mathcal{E},X) \right\}. \end{aligned}$$

For these last three steps, recall that derivation trees from X first choose a monomial (and corresponding coefficient) $c\cdot m \in P_X$. The root ε then has yield c and the children are derivation trees from the indeterminates occurring in m. By associativity of the infinite product $\mathsf{yd}(T\|^{\boldsymbol{b}}_{n+1})$, we can group together the yields of the child subtrees $T_{i,j}$. □

To prove Theorem 4.2.6, all that is left to do is to consider the supremum of the iteration $\big(F^n(\mathbf{0})\big)_{n<\omega}$ and the corresponding tree truncations, and dually the infimum of $\big(F^n(\mathbf{1})\big)_{n<\omega}$. For the infimum, one last obstacle needs to be resolved: we must show that whenever we pick for each n some n-truncation, their infimum can still be realised as yield of an actual (infinite) tree, even if we pick a different tree to truncate for each n. A similar, but more involved, argument is needed for strategy trees of model-checking games (cf. Puzzle lemma 6.2.8), so we only sketch the proof of the simpler observation we need here. Essentially, the finite number of indeterminates in $\boldsymbol{A}$ allows us to choose a sufficiently large n such that the truncation $T_n\|^{\mathbf{1}}_n$ contains a "nice" part that we can repeat to obtain the desired infinite tree.

Lemma 4.2.10 (Simple puzzle lemma). *Let $X \in \boldsymbol{X}$ and let $(T_n)_{n<\omega}$ be a family of trees $T_n \in \mathcal{T}(\mathcal{E}, X)$ such that their yields $\big(\mathsf{yd}(T_n\|^{\mathbf{1}}_n)\big)_{n<\omega}$ form a descending chain in K. Then there is a tree $T' \in \mathcal{T}(\mathcal{E}, X)$ with $\mathsf{yd}(T') \geq \prod_{n<\omega} \mathsf{yd}(T_n\|^{\mathbf{1}}_n)$.*

Proof sketch. The proof proceeds in the following steps.

(1) *Chain splitting*

Let $e_n(m, Y) = |T_n\|^{\mathbf{1}}_n|_{m,Y} \in \mathbb{N}^\infty$ be the number of occurrences of $m \in P_Y$ in the tree $T\|^{\mathbf{1}}_n$. As usual, let $c_{m,Y}$ be the coefficient of m in P_Y so that

$$y_n \coloneqq \mathsf{yd}(T_n\|^{\mathbf{1}}_n) = \prod_{\substack{Y \in \boldsymbol{X},\\ m \in P_Y}} (c_{m,Y})^{e_n(m,Y)}.$$

Using Lemma 3.2.8 (Chain splitting), we can write the infimum as

$$\bigsqcap_{n<\omega} y_n = \prod_{\substack{Y \in \boldsymbol{X},\\ m \in P_Y}} (c_{m,Y})^{e(m,Y)}, \qquad e(m, Y) \coloneqq \bigsqcup_{n<\omega} e_n(m, Y).$$

(2) *Problematic monomials*

We say that a monomial $m \in P_Y$ is *problematic* if $e(m, Y)$ is finite. Unproblematic monomials do not impose any restrictions, as they may appear arbitrarily often (finite or infinite) in the tree T' we construct (and T' still satisfies the desired inequality $\mathsf{yd}(T') \geq \bigsqcap_n y_n$). As the polynomial equation system is finite, there are only finitely many problematic monomials.

(3) *Decomposition into k-layers*

We partition each of the trees $T_n\|^{\mathbf{1}}_n$ into k-layers, each of which simply consists of k consecutive levels of the tree. We choose the layers such that they cover the entire tree and do not overlap. Given a constant k (to be determined later), we choose a large enough n such that there is a k-layer in $T_n\|^{\mathbf{1}}_n$ that does not contain any problematic monomials (as there are only finitely many).

(4) *Repetition of a k-layer*

We now construct T' by first following the tree T_n (for the chosen n), but upon reaching the k-layer without problematic monomials, we continue T' by repeating this layer over and over, so that no further problematic monomials occur in T'.

An easy way to describe the repetition is by the notion of deterministic trees introduced in the next section. First note that the k-layer is a forest consisting of several trees; then determinise each such tree S by Lemma 4.3.2 and obtain a (possibly infinite) deterministic tree S' that we use in T' to replace S. To ensure that Lemma 4.3.2 can be applied, we choose $k = \ell + 1$, so that each path trough the k-layer must contain a repetition of indeterminates (cf. Corollary 4.3.4).

(5) *Conclusion*

This construction ensures that the number of occurrences of problematic monomials in T' is bounded by their occurrences in $T_n\|^{\mathbf{1}}_n$. In other words, $|T'|_{m,Y} \leq e_n(m, Y) \leq e(m, Y)$ for all problematic $m \in P_Y$ and thus $|T'|_{m,Y} \leq e(m, Y)$ for all (problematic and unproblematic) m, Y. Hence

$$\mathsf{yd}(T') = \prod_{\substack{Y \in \boldsymbol{X},\\ m \in P_Y}} (c_{m,Y})^{|T'|_{m,Y}} \geq \prod_{\substack{Y \in \boldsymbol{X},\\ m \in P_Y}} (c_{m,Y})^{e(m,Y)} = \bigsqcap_{n<\omega} y_n. \qquad \square$$

With this taken care of, we can prove that the sum of all (finite) derivation trees gives the least and greatest solutions.

Proof of Theorem 4.2.6. Recall that F is fully continuous, as it is defined by polynomial expressions over a fully-continuous semiring. By Kleene's fixed-point theorem, we can thus express its least (or greatest) fixed point as supremum of $F^n(\mathbf{0})$ (or infimum of $F^n(\mathbf{1})$) over $n < \omega$. By idempotence, sums coincide with suprema, so for the least solution we immediately obtain:

$$\mathbf{lfp}(F)_X = \bigsqcup_{n<\omega} F^n(\mathbf{0})_X \overset{(4.2.9)}{=} \bigsqcup_{n<\omega} \Big(\widehat{\sum_{T \in \mathcal{T}(\mathcal{E},X)}} \mathsf{yd}(T\|^{\mathbf{0}}_n) \Big) = \widehat{\sum_{T \in \mathcal{T}(\mathcal{E},X)}} \Big(\bigsqcup_{n<\omega} \mathsf{yd}(T\|^{\mathbf{0}}_n) \Big).$$

Now observe that $\mathsf{yd}(T\|^{\mathbf{0}}_n) = \mathsf{yd}(T)$ if T has height $< n$, otherwise $\mathsf{yd}(T\|^{\mathbf{0}}_n) = 0$. Hence $\bigsqcup_n \mathsf{yd}(T\|^{\mathbf{0}}_n) = \mathsf{yd}(T)$ if T is finite and 0 otherwise.

For the greatest solution, it suffices to consider $K = \mathbb{S}^\infty(\boldsymbol{A})$ by Lemma 4.2.7. We can thus apply Proposition 3.5.8 to express the infimum in $\mathbb{S}^\infty(\boldsymbol{A})$ as a supremum over monomial chains:

$$\begin{aligned}
\mathbf{gfp}(F)_X &= \prod_{n<\omega} F^n(\mathbf{1})_X \overset{(4.2.9)}{=} \prod_{n<\omega} \Big(\widehat{\sum_{T \in \mathcal{T}(\mathcal{E},X)}} \mathsf{yd}(T\|^{\mathbf{1}}_n) \Big) \\
&\overset{(3.5.8)}{=} \bigsqcup \Big\{ \prod_{n\in\mathbb{N}} y_n \;\Big|\; \begin{array}{l} (y_n)_{n<\omega} \text{ is a descending chain of monomials} \\ \text{with } y_n = \mathsf{yd}(T_n\|^{\mathbf{1}}_n) \text{ for some } T_n \in \mathcal{T}(\mathcal{E},X) \end{array} \Big\} \\
&\overset{(4.2.10)}{=} \bigsqcup \Big\{ \mathsf{yd}(T') \;\Big|\; T' \in \mathcal{T}(\mathcal{E},X) \Big\} = \widehat{\sum_{T \in \mathcal{T}(\mathcal{E},X)}} \mathsf{yd}(T).
\end{aligned}$$

In the last line, we apply the Simple puzzle lemma. This gives us for each monomial chain $(y_n)_{n<\omega}$ an infinite tree T' with $\mathsf{yd}(T') \geq \prod_n y_n$. Conversely, each tree T' induces the monomial chain defined by $y_n = \mathsf{yd}(T'\|^{\mathbf{1}}_n)$. It is easy to see that this chain has infimum $\mathsf{yd}(T')$, so we have equality. □

4.3 Closed-Form Solution

This section is devoted to the proof of this chapter's main result:

Theorem 4.3.1. *Let K be an absorptive, fully-continuous semiring. Let $\mathcal{E}\colon (X = P_X)_{X\in\boldsymbol{X}}$ be a polynomial equation system over K and $\boldsymbol{X} = \{X_1, \dots, X_\ell\}$ with induced operator $F_\mathcal{E}\colon K^\ell \to K^\ell$. Then,*

$$\mathbf{lfp}(F_\mathcal{E}) = F^\ell_\mathcal{E}(\mathbf{0}), \qquad \mathbf{gfp}(F_\mathcal{E}) = F^\ell_\mathcal{E}(\, F^\ell_\mathcal{E}(\mathbf{1})^\infty \,).$$

Towards the proof, we again fix a polynomial equation system $\mathcal{E}\colon (X = P_X)_{X\in\boldsymbol{X}}$ over an absorptive, fully-continuous semiring K with induced operator F. Recall that ℓ denotes the number of equations (and indeterminates) of $\mathcal{E}$. Our strategy is to prove that we can always find derivation trees of a certain shape, and that the yield of all other derivation trees is absorbed by these trees.

4.3.1 Deterministic Derivation Trees

A derivation tree is deterministic if all its subtrees rooted in an indeterminate X are identical, so that it makes a unique choice how to continue the derivation from X (for all $X \in \boldsymbol{X}$).

Definition 4.3.1. A derivation tree $T = (V, E, \mathsf{var}, \mathsf{yd})$ is said to be *deterministic* if $\mathsf{mon}(v)$ depends only on $\mathsf{var}(v)$. That is, the relation $\{(\mathsf{var}(v), \mathsf{mon}(v)) \mid v \in V\}$ is a function.

To reason about $\mathbf{gfp}(F)$, we must reason about and construct infinite derivation trees. This is straightforward for deterministic trees.

Lemma 4.3.2 (Deterministic construction)**.** *Let $\boldsymbol{X}_0 \subseteq \boldsymbol{X}$. For each $X \in \boldsymbol{X}_0$, let m_X be a monomial with $m_X \in P_X$ such that all indeterminates occurring in m_X are contained in $\boldsymbol{X}_0$. Then for each $X \in \boldsymbol{X}_0$, there is a deterministic tree $T \in \mathcal{T}(\mathcal{E}, X)$ with $\mathsf{var}(v) \in \boldsymbol{X}_0$ and $\mathsf{mon}(v) = m_{\mathsf{var}(v)}$ for all nodes $v \in T$.*

Proof sketch. Starting from the root $\mathsf{var}(\varepsilon) = X$, define the (possibly infinite) tree T inductively by repeatedly adding to each leaf v with label $\mathsf{var}(v) = Y$ child nodes according to m_Y. The invariant $\mathsf{var}(v) \in \boldsymbol{X}_0$ for all nodes v is always maintained. □

It is easy to see that deterministic trees are uniquely defined by their prefix up to depth $\ell - 1$, as at depth ℓ each path must either end or start to repeat (recall that there are only ℓ different indeterminates). Moreover, once we consider the infinitary power $\mathsf{yd}(T)^\infty$, it does not matter how often a particular coefficient c occurs in T, since $(c^n)^\infty = c^\infty$ for all $n > 0$. This leads to the following simple but essential observations.

Lemma 4.3.3. *If $T \in \mathcal{T}(\mathcal{E}, X)$ is deterministic, every indeterminate that occurs in T also occurs in the truncation $T\|^{\mathbf{1}}_{\ell-1}$. It follows that $\mathsf{yd}(T)^\infty = \mathsf{yd}(T\|^{\mathbf{1}}_{\ell})^\infty$.*

Corollary 4.3.4. *For each $T \in \mathcal{T}(\mathcal{E}, X)$, there is a deterministic tree $T' \in \mathcal{T}(\mathcal{E}, X)$ such that $\mathsf{yd}(T\|^{\mathbf{1}}_{\ell})^\infty \leq \mathsf{yd}(T')^\infty$.*

Proof. Choose any way to determinise $T\|^{\mathbf{1}}_{\ell}$ (by Lemma 4.3.2) using only monomials appearing in $T\|^{\mathbf{1}}_{\ell}$. This is always possible, as $T\|^{\mathbf{1}}_{\ell}$ contains at most ℓ indeterminates and hence every path must contain a repetition or end in a leaf.

More precisely, we apply Lemma 4.3.2 as follows. We define $\boldsymbol{X}_0$ and monomials m_X for $X \in \boldsymbol{X}_0$ inductively by traversing the tree T level by level, starting with the root. We always maintain the following invariant: after traversing level k, all indeterminates occurring in m_X for $X \in \boldsymbol{X}_0^{(k)}$ are contained in $\boldsymbol{X}_0^{(k)}$ or occur at level $k+1$ in T.

For level 0 (consisting only of the root with $\mathsf{var}(\varepsilon) = X$), we set $\boldsymbol{X}_0^{(0)} := \{X\}$ and $m_X := \mathsf{mon}(\varepsilon)$. The invariant holds due to the children of ε. Assume we have processed level k. Let $\boldsymbol{X}_{\text{todo}}^{(k)}$ be the set of indeterminates that occur in m_X for $X \in \boldsymbol{X}_0^{(k)}$, but are not contained in $\boldsymbol{X}_0^{(k)}$. If $\boldsymbol{X}_{\text{todo}}^{(k)} = \emptyset$, we are done. Otherwise, for each $Y \in \boldsymbol{X}_{\text{todo}}^{(k)}$, there is a node v at level $k+1$ with $\mathsf{var}(v) = Y$ by the invariant. Choose any such node and set $m_Y := \mathsf{mon}(v)$. By definition of derivation trees, all indeterminates of m_Y occur as children of v on level $k+2$, so the invariant holds. Then set $\boldsymbol{X}_0^{(k+1)} := \boldsymbol{X}_0^{(k)} \cup \boldsymbol{X}_{\text{todo}}^{(k)}$.

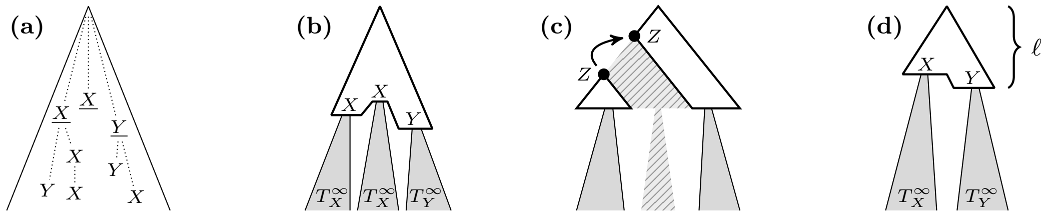

Figure 4.2: Illustration of the construction steps in the proof of Lemma 4.3.5.

This process never enters level ℓ: On each level, at least one indeterminate is added to $\boldsymbol{X}_0^{(k)}$ (or we stop) and there are only ℓ different indeterminates. Hence all monomials m_X we choose occur in the truncation $T\|_{\ell}^{\mathbf{1}}$. Now obtain T' by applying Lemma 4.3.2 to $\bigcup_k \boldsymbol{X}_0^{(k)}$ and the chosen monomials m_X.

The inequality $\mathsf{yd}(T\|_{\ell}^{\mathbf{1}})^{\infty} \leq \mathsf{yd}(T')^{\infty}$ follows by Lemma 4.2.5 (Yield comparison). □

4.3.2 Constructing Simple Trees

The main insight behind Theorem 4.3.1 is that when we sum over the yields of all derivation trees, it suffices to consider trees of a particular shape corresponding to our intuition of a reachability and a cyclic part in the introductory example. That is, these trees consist of an arbitrary prefix up to (at most) depth ℓ (the reachability part), followed by deterministic trees (the cyclic part). See Fig. 4.2d for an illustration.

Lemma 4.3.5 (Simple tree)**.** *For each $T \in \mathcal{T}(\mathcal{E}, X)$, there is a derivation tree $T' \in \mathcal{T}(\mathcal{E}, X)$ such that all subtrees rooted at depth ℓ in T' are deterministic and use only monomials $m \in P_Y$ (with $Y \in \boldsymbol{X}$) that occur infinitely often in T. Moreover, $|T'|_{m,Y} \leq |T|_{m,Y}$ for all $m \in P_Y, Y \in \boldsymbol{X}$.*

Proof. Let $\boldsymbol{X}_\infty \subseteq \boldsymbol{X}$ be the (possibly empty) set of indeterminates that occur infinitely often in T. We write $V_X = \{v \in T \mid \mathsf{var}(v) = X\}$ for the set of nodes labeled X. For each $X \in \boldsymbol{X}_\infty$, the set V_X is infinite. As the polynomial P_X is finite, there must thus be infinitely many $v \in V_X$ with the same monomial $\mathsf{mon}(v)$. For each $X \in \boldsymbol{X}_\infty$, choose such an infinitely often occurring monomial $m_X \in P_X$. Using Lemma 4.3.2, we obtain for each $X \in \boldsymbol{X}_\infty$ a deterministic tree $T_X^\infty \in \mathcal{T}(\mathcal{E}, X)$ such that for all $v \in T_X^\infty$, we have $\mathsf{var}(v) \in \boldsymbol{X}_\infty$ and $\mathsf{mon}(v)$ occurs infinitely often in T.

Let W be the set of earliest occurrences of $\boldsymbol{X}_\infty$ in T (Fig. 4.2a). Formally, $W = \{v \in T \mid \mathsf{var}(v) \in \boldsymbol{X}_\infty$, there is no ancestor $v' \sqsubset v$ with $\mathsf{var}(v') \in \boldsymbol{X}_\infty\}$. Now let S be the tree that results from T by replacing the subtree at each $v \in W$ with the tree $T_{\mathsf{var}(v)}^\infty$ (Fig. 4.2b). The tree S is almost of the desired shape, but the trees T_X^∞ may be rooted at depth $> \ell$. To fix this, we consider the prefix up to the subtrees T_X^∞ and eliminate all repetitions of indeterminates within the prefix. As all indeterminates in the prefix occur only finitely often, we can eliminate repetitions by replacing each first occurrence of an indeterminate Z by a last occurrence of Z within the prefix (Fig. 4.2c).

More formally, call a path $v_0 v_1 v_2 \dots v_k$ from the root of S *unresolved* if $k \geq \ell$ and no node on the path is contained in one of the deterministic subtrees T_X^∞. Given an unresolved path, there must be an indeterminate Z that occurs twice on the path. Further, $Z \notin \boldsymbol{X}_\infty$, as otherwise the nodes labeled Z would lie within one of the deterministic subtrees by construction of S and W. Let $i < j$ be indices such that v_i is the first and v_j the last occurrence of Z on the path, so $\mathsf{var}(v_i) = \mathsf{var}(v_j) = Z$. Now let S' result from S by replacing the subtree S_{v_i} rooted at v_i with the subtree S_{v_j} rooted at v_j, thereby removing at least one occurrence of Z from the tree (Fig. 4.2c).

Apply this elimination step exhaustively, until there are no more unresolved paths. As all indeterminates in the prefix of S occur only finitely often, this process terminates. Let T' be the resulting tree (notice that T' may not be uniquely determined, but this does not affect our argument). Then T' has the desired shape: when no unresolved path exists, then all nodes at depth ℓ (if any) must be contained in one of deterministic subtrees T_X^∞.

Moreover, the elimination step only removes nodes of S, but neither adds nodes nor modifies any node labels. It follows that $|T'|_{m,Y} \leq |S|_{m,Y}$ for all $m \in P_Y, Y \in \boldsymbol{X}$. As the trees T_X^∞ only use monomials that occur infinitely often in T, we further have $|S|_{m,Y} \leq |T|_{m,Y}$ for all m, Y, closing the proof. □

4.3.3 Proof of the Main Result

We relate infinite trees of this shape to the expression $F^\ell(\,F^\ell(\mathbf{1})^\infty\,)$. The deterministic trees rooted at depth ℓ correspond to the inner term $F^\ell(\mathbf{1})^\infty$, relying on Lemma 4.3.3 to ensure that ℓ applications of F suffice. The outer applications of F correspond to the prefix on which we impose no further restrictions (except that it has height at most ℓ). The following lemma formalizes this intuition.

Lemma 4.3.6. *Let $\boldsymbol{b}$ be the tuple with $\boldsymbol{b}_X = \widehat{\sum}_{T \in \mathcal{T}(\mathcal{E},X)} \mathsf{yd}(T\|^{\mathbf{1}}_{\ell})^\infty$ for $X \in \boldsymbol{X}$. For each $T \in \mathcal{T}(\mathcal{E}, X)$, there is a tree $T' \in \mathcal{T}(\mathcal{E}, X)$ such that $\mathsf{yd}(T) \leq \mathsf{yd}(T'\|^{\boldsymbol{b}}_{\ell})$.*

Proof. Let $T \in \mathcal{T}(\mathcal{E}, X)$. Using Lemma 4.3.5, we obtain a tree T' of a certain shape: Let $S_1, \ldots, S_k$ be the subtrees of T' rooted at depth ℓ. These subtrees are deterministic and all monomials occurring in $S_1, \ldots, S_k$ occur infinitely often in T and moreover, $|T'|_{m,Y} \leq |T|_{m,Y}$ for all $m \in P_Y, Y \in \boldsymbol{X}$. We claim that

$$\mathsf{yd}(T) \overset{!}{\leq} \mathsf{yd}(T'\|^{\mathbf{1}}_{\ell}) \cdot \prod_{i=1}^{k} \mathsf{yd}(S_i)^\infty.$$

To see this, we expand the definition of yd and rearrange terms. Borrowing the notation $c_{m,Y}$ for the coefficient of $m \in P_Y$ from the proof of Lemma 4.2.5, we obtain

$$\prod_{\substack{Y \in \boldsymbol{X},\\ m \in P_Y}} (c_{m,Y})^{|T|_{m,Y}} \overset{!}{\leq} \prod_{\substack{Y \in \boldsymbol{X},\\ m \in P_Y}} (c_{m,Y})^{|T'\|^{\mathbf{1}}_{\ell}|_{m,Y} + \sum_{i=1}^{k} \infty \cdot |S_i|_{m,Y}}.$$

By monotonicity, it suffices to show that $|T|_{m,Y} \geq |T'\|^{\mathbf{1}}_{\ell}|_{m,Y} + \sum_{i=1}^{k} \infty \cdot |S_i|_{m,Y}$ for all m, Y. This holds by construction of T': If $|S_i|_{m,Y} > 0$ for some i, then m occurs infinitely often in T and hence $|T|_{m,Y} = \infty$. Otherwise, the right-hand side is equal to $|T'\|^{\mathbf{1}}_{\ell}|_{m,Y} \leq |T'|_{m,Y} \leq |T|_{m,Y}$. This proves our claim.

Now let $v_1, \dots, v_k \in T'$ be the root nodes of the deterministic subtrees $S_1, \dots, S_k$. By Lemma 4.3.3, $\mathsf{yd}(S_i)^\infty = \mathsf{yd}(S_i \|^{\mathbf{1}}_{\ell})^\infty \leq \boldsymbol{b}_{\mathsf{var}(v_i)}$, and thus

$$\mathsf{yd}(T) \;\leq\; \mathsf{yd}(T' \|^{\mathbf{1}}_{\ell}) \cdot \prod_{i=1}^{k} \mathsf{yd}(S_i)^\infty \;\leq\; \mathsf{yd}(T' \|^{\mathbf{1}}_{\ell}) \cdot \prod_{i=1}^{k} \boldsymbol{b}_{\mathsf{var}(v_i)} \;=\; \mathsf{yd}(T' \|^{\boldsymbol{b}}_{\ell}). \qquad \square$$

We are now ready to prove our main result. The statement on the least solution follows rather directly from our earlier considerations. For greatest fixed points, the previous lemma already proves the difficult direction.

Proof of Theorem 4.3.1. It suffices to consider the case $K = \mathbb{S}^\infty(\boldsymbol{A})$, as the general statement follows with Lemma 4.1.2. We first consider the least solution. It is clear by monotonicity of F that $F^\ell(\mathbf{0})_X \leq \mathbf{lfp}(F)_X$. By Theorem 4.2.6 and Lemma 4.2.9, it thus suffices to prove

$$\mathbf{lfp}(F)_X \;=\; \widehat{\sum_{\substack{T \in \mathcal{T}(\mathcal{E},X) \\ T \text{ finite}}}} \mathsf{yd}(T) \;\overset{!}{\leq}\; \widehat{\sum_{T \in \mathcal{T}(\mathcal{E},X)}} \mathsf{yd}(T \|^{\mathbf{0}}_{\ell}) \;=\; F^\ell(\mathbf{0})_X.$$

To this end, let $T \in \mathcal{T}(\mathcal{E}, X)$ be finite and obtain a simple tree T' by Lemma 4.3.5. As T is finite, no monomials can occur infinitely often. Hence T' has no subtrees rooted at depth ℓ and is thus of height $< \ell$. But then, $\mathsf{yd}(T) \leq \mathsf{yd}(T') = \mathsf{yd}(T' \|^{\mathbf{0}}_{\ell}) \leq F^\ell(\mathbf{0})_X$.

For the greatest solution, we know that $\mathbf{gfp}(F)_X = \widehat{\sum}_{T \in \mathcal{T}(\mathcal{E},X)} \mathsf{yd}(T)$. On the other hand, we apply Lemma 4.2.9 (Tree iteration) to $F^\ell(\mathbf{1})$. Since we work in $\mathbb{S}^\infty(\boldsymbol{A})$, we know that infinitary power commutes with infinite summation (Lemma 3.5.6), thus

$$(F^\ell(\mathbf{1})_X)^\infty = \left(\widehat{\sum_{S \in \mathcal{T}(\mathcal{E},X)}} \mathsf{yd}(S \|^{\mathbf{1}}_{\ell}) \right)^\infty = \widehat{\sum_{S \in \mathcal{T}(\mathcal{E},X)}} \mathsf{yd}(S \|^{\mathbf{1}}_{\ell})^\infty.$$

Now let $\boldsymbol{b} = F^\ell(\mathbf{1})^\infty$. Applying Lemma 4.2.9 again gives

$$F^\ell(F^\ell(\mathbf{1})^\infty)_X = F^\ell(\boldsymbol{b})_X = \widehat{\sum_{T \in \mathcal{T}(\mathcal{E},X)}} \mathsf{yd}(T \|^{\boldsymbol{b}}_{\ell}).$$

The direction $\mathbf{gfp}(F) \leq F^\ell(F^\ell(\mathbf{1})^\infty)$ follows immediately from Lemma 4.3.6. For the other direction, let $T \in \mathcal{T}(\mathcal{E}, X)$ and let $v_1, \dots, v_k$ be the nodes at depth ℓ in T. By distributivity, we get

$$\begin{aligned}
\mathsf{yd}(T \|^{\boldsymbol{b}}_{\ell}) &= \mathsf{yd}(T \|^{\mathbf{1}}_{\ell}) \cdot \prod_{1 \leq i \leq k} \boldsymbol{b}_{\mathsf{var}(v_i)} \\
&\overset{\text{dist.}}{=} \mathsf{yd}(T \|^{\mathbf{1}}_{\ell}) \cdot \widehat{\sum} \Big\{ \prod_{1 \leq i \leq k} \mathsf{yd}(S_i \|^{\mathbf{1}}_{\ell})^\infty \;\Big|\; S_i \in \mathcal{T}(\mathcal{E}, \mathsf{var}(v_i)) \text{ for all } i \Big\} \\
&\overset{(4.3.4)}{\leq} \mathsf{yd}(T \|^{\mathbf{1}}_{\ell}) \cdot \widehat{\sum} \Big\{ \prod_{1 \leq i \leq k} \mathsf{yd}(S'_i)^\infty \;\Big|\; S'_i \in \mathcal{T}(\mathcal{E}, \mathsf{var}(v_i)) \text{ for all } i \Big\} \\
&\overset{\text{abs.}}{\leq} \mathsf{yd}(T \|^{\mathbf{1}}_{\ell}) \cdot \widehat{\sum} \Big\{ \prod_{1 \leq i \leq k} \mathsf{yd}(S'_i) \;\Big|\; S'_i \in \mathcal{T}(\mathcal{E}, \mathsf{var}(v_i)) \text{ for all } i \Big\} \\
&\overset{\text{dist.}}{=} \widehat{\sum} \Big\{ \mathsf{yd}(T) \;\Big|\; T \in \mathcal{T}(\mathcal{E}, X) \Big\} = \mathbf{gfp}(F)_X. \qquad \square
\end{aligned}$$

Using this result, we can compute least and, most importantly, greatest solutions of polynomial equation systems in a polynomial number of semiring operations (including the infinitary power). Notice that no details of the polynomial system are needed for the computation, we only need to evaluate F. Moreover, although the proof relied on $\mathbb{S}^\infty(\boldsymbol{A})$, the computation happens only in the target semiring. For instance, recall the introductory example in the tropical semiring.

Example 4.3.7. We recall $X_a = 1 +_\mathbb{R} X_a$, $X_b = \min(1 +_\mathbb{R} X_a, 20 +_\mathbb{R} X_c)$ and $X_c = 0 +_\mathbb{R} X_c$. Notice that the one-element of the tropical semiring is the real value 0. Using Theorem 4.3.1, we can collapse the infinite fixed-point iteration to

$$\begin{pmatrix}0\\0\\0\end{pmatrix} \overset{F}{\mapsto} \begin{pmatrix}1\\1\\0\end{pmatrix} \overset{F}{\mapsto} \begin{pmatrix}2\\2\\0\end{pmatrix} \overset{F}{\mapsto} \begin{pmatrix}3\\3\\0\end{pmatrix} \overset{\infty}{\longmapsto} \begin{pmatrix}\infty\\\infty\\0\end{pmatrix} \overset{F}{\mapsto} \begin{pmatrix}\infty\\20\\0\end{pmatrix} \circlearrowleft^F$$

and obtain the expected solution. In this example, one iteration of F would actually suffice (instead of $\ell = 3$ iterations), since reachability paths and cycles have length one (see the graph in the introduction). In general, all ℓ steps are required (see Example 4.3.8 below).

This result can be used to compute semiring semantics of fixed-point logics (see Section 5.2), using a polynomial number of semiring operations per fixed-point computation. Whether these operations can be computed efficiently depends on the semiring. This is not the case for the most general semiring $\mathbb{S}^\infty(\boldsymbol{X})$, as each multiplication can in the worst case double the number of monomials (which is not avoidable, as one can easily construct equation systems whose solution is a polynomial describing all exponentially many paths in a graph). Even if the solution consists of few monomials, we may get an intermediary blowup when computing $F^n(\mathbf{1})$, as seen in the following example. In such cases, the symbolic approach presented in Section 4.4 may be preferable.

Example 4.3.8. Consider the polynomial equation system over $\mathbb{S}^\infty(a, b, c)$ shown on the left. (This results from evaluating the L_μ-formula $\nu X.\ \Box X \wedge \Diamond P$ (via its translation into LFP) in a semiring interpretation that uses the indeterminates a, b, c to track whether the atom P holds at vertices v_1, v_2, v_3, cf. Section 5.2.)

$$\begin{aligned} X_1 &= b \cdot X_2 \\ X_2 &= (b + c) \cdot X_2 X_3 \\ X_3 &= a \cdot X_1 \end{aligned}$$

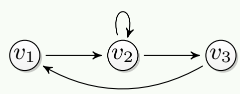

We apply Theorem 4.3.1 to compute the greatest solution:

$$\mathbf{1} \overset{F}{\mapsto} \begin{pmatrix} b \\ b+c \\ a \end{pmatrix} \overset{F}{\mapsto} \begin{pmatrix} b^2 + bc \\ ab^2 + abc + ac^2 \\ ab \end{pmatrix} \overset{F}{\mapsto} \begin{pmatrix} ab^3 + ab^2c + abc^2 \\ a^2b^4 + a^2b^3c + a^2b^2c^2 + a^2bc^3 \\ ab^2 + abc \end{pmatrix} \overset{\infty}{\longmapsto} \begin{pmatrix} a^\infty b^\infty \\ a^\infty b^\infty \\ a^\infty b^\infty \end{pmatrix} \circlearrowleft^F$$

Here we need all $\ell = 3$ inner applications of F until a appears in the first entry. We also see that the intermediate polynomials can become much longer than the solution.

Remark 4.3.9. The intermediary blowup in the size of the polynomials can be reduced. We see in the previous example that the monomials cancel each other by absorption once the infinitary power is applied. We can thus reduce the number of monomials by always applying the infinitary power after each step.

More precisely, let F be an operator induced by a polynomial equation system over $\boldsymbol{X}$. We define the operator $\tilde{F}\colon K^\ell \to K^\ell$ by $\tilde{F}(\boldsymbol{b}) = F(\boldsymbol{b})^\infty$ and observe that

$$\mathbf{gfp}(F) = F^\ell(F^\ell(\mathbf{1})^\infty) = F^\ell(\tilde{F}^\ell(\mathbf{1})).$$

To see this, we can show by induction that $\tilde{F}^n(\boldsymbol{b}) = F^n(\boldsymbol{b})^\infty$. For $n = 1$ this is the definition. For the induction step, we note that for a polynomial $P(\boldsymbol{X})$ and $\boldsymbol{b} \in K^\ell$, we have $P^\infty(\boldsymbol{b}^\infty) = P^\infty(\boldsymbol{b})$. This is easy to prove by induction on P using the properties of the infinitary power in Lemma 3.3.2. Then $\tilde{F}^{n+1}(\boldsymbol{b}) = F(\tilde{F}^n(\boldsymbol{b}))^\infty = F(F^n(\boldsymbol{b})^\infty)^\infty = F(F^n(\boldsymbol{b}))^\infty$.

The iteration of the previous example then becomes:

$$\mathbf{1} \overset{\tilde{F}}{\mapsto} \begin{pmatrix} b^\infty \\ b^\infty + c^\infty \\ a^\infty \end{pmatrix} \overset{\tilde{F}}{\mapsto} \begin{pmatrix} b^\infty \\ a^\infty b^\infty + a^\infty c^\infty \\ a^\infty b^\infty \end{pmatrix} \overset{\tilde{F}}{\mapsto} \begin{pmatrix} a^\infty b^\infty \\ a^\infty b^\infty \\ a^\infty b^\infty \end{pmatrix} \circlearrowleft^{F}$$

4.3.4 Generalised Polynomial Equation Systems

We discuss how to extend the closed-form computation to polynomial equation systems with exponents in $\mathbb{N}^\infty$ instead of $\mathbb{N}$ (i.e., to $\mathrm{Poly}^\omega(K, \mathbb{N}^\infty, \boldsymbol{X})$). Derivation trees are then no longer finitely branching, as a monomial X^∞ corresponds to a node with infinitely many children. The key arguments of the proof remain the same; we sketch the few modifications below.

Example 4.3.10. Consider this generalised polynomial equation system over $\mathbb{S}^\infty(\boldsymbol{A})$:

$$X = aXY^\infty + bY^\infty$$
$$Y = cY^\infty + d$$

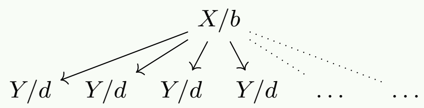

The least solution is given by $X = bd^\infty$, $Y = d$. This corresponds to (possibly infinite) derivation trees of finite height, as indicated by the infinite tree above with yield bd^∞.

For the greatest solution, we compute

$$\mathbf{1} \overset{F}{\mapsto} \begin{pmatrix} a+b \\ c+d \end{pmatrix} \overset{F}{\mapsto} \begin{pmatrix} a^2c^\infty + a^2d^\infty + bc^\infty + bd^\infty \\ c^\infty + d \end{pmatrix}$$

and then $F^2(F^2(\mathbf{1})^\infty)$ gives the result:

$$\begin{pmatrix} a^\infty c^\infty + a^\infty d^\infty + b^\infty c^\infty + b^\infty d^\infty \\ c^\infty + d^\infty \end{pmatrix} \overset{F}{\mapsto} \begin{pmatrix} a^\infty c^\infty + a^\infty d^\infty + bc^\infty + bd^\infty \\ c^\infty + d \end{pmatrix} \circlearrowleft^{F}$$

Homomorphisms. We first note that least and greatest solutions are still preserved by fully-continuous homomorphisms. The proof of Lemma 4.1.2 needs to be modified, as the operators $F_\mathcal{E}$, $F_{h(\mathcal{E})}$ may not be fully-continuous (consider $\mathcal{E}\colon X = X^\infty$ and Example 3.3.4), but this is easily resolved by considering the transfinite fixed-point iteration instead of stopping at ω. It follows that Lemma 4.2.7 (restriction to $K = \mathbb{S}^\infty(\boldsymbol{A})$) remains valid.

Iteration via trees. Derivation trees can be defined as before, but now a node v may have countably infinitely many children if the monomial $\mathsf{mon}(v)$ has infinite exponents. We remark that the infinite number of children may appear more powerful than it actually is: there are only finitely many choices for each children's indeterminate, so one choice must be repeated infinitely often; it is then always better (in terms of a larger yield) to make the same choice for all children.

The number of indeterminates remains finite, and hence Lemma 4.2.5 (Yield comparison), Lemma 4.3.2 (Deterministic construction) and Lemma 4.2.10 (Simple puzzle lemma) still hold. Modifications are needed in the proof of Lemma 4.2.9, we sketch these below. The proof of Theorem 4.2.6 then goes through, we only have to replace finite trees with trees of finite height for least solutions:

$$\mathbf{lfp}(F_{\mathcal{E}})_X = \widehat{\sum_{\substack{T \in \mathcal{T}(\mathcal{E},X),\\ T \text{ has finite height}}}} \mathsf{yd}(T).$$

For Lemma 4.2.9, we first note that we only need the lemma for $K = \mathbb{S}^\infty(\boldsymbol{A})$. We adapt the proof to account for the infinite products induced by infinite exponents.

Lemma 4.2.9′ (Tree iteration, modified)**.** *Let $K = \mathbb{S}^\infty(\boldsymbol{A})$ and $\boldsymbol{b} \in K^\ell$. Then $F^n(\boldsymbol{b})_X = \widehat{\sum}_{T \in \mathcal{T}(\mathcal{E},X)} \mathsf{yd}(T\|^{\boldsymbol{b}}_n)$ for all $n < \omega$ and $X \in \boldsymbol{X}$.*

Modified proof sketch. Recall that we need to prove $P_X(\boldsymbol{a}) = \widehat{\sum}_{T \in \mathcal{T}(\mathcal{E},X)} \mathsf{yd}(T\|^{\boldsymbol{b}}_{n+1})$ for the induction step, where $\boldsymbol{a}$ is defined by $\boldsymbol{a}_X = \widehat{\sum}_{T \in \mathcal{T}(\mathcal{E},X)} \mathsf{yd}(T\|^{\boldsymbol{b}}_n)$. For a monomial m, let I_m be the set of pairs (i,j) with $1 \le i \le \ell$ and $1 \le j \le m(X_i)$ if $m(X_i)$ is finite, or $1 \le j < \omega$ otherwise. With this notation, we can write

$$P_X(\boldsymbol{a}) = \sum_{c \cdot m \in P_X} \Big(c \cdot \prod_{X \in \boldsymbol{X}} \boldsymbol{a}_X^{m(X)} \Big) = \sum_{c \cdot m \in P_X} \Big(c \cdot \widehat{\prod_{(i,j) \in I_m}} \boldsymbol{a}_{X_i} \Big)$$

We now apply strong distributivity of $\mathbb{S}^\infty(\boldsymbol{A})$ (Corollary 3.5.10):

$$= \widehat{\sum} \left\{ c \cdot \widehat{\prod_{(i,j) \in I_m}} \mathsf{yd}(T_{i,j}\|^{\boldsymbol{b}}_n) \;\middle|\; \begin{array}{l} c \cdot m \in P_X \text{ and all possible choices} \\ \text{of trees } T_{i,j} \in \mathcal{T}(\mathcal{E}, X_i), \text{ for } (i,j) \in I_m \end{array} \right\}$$

$$= \widehat{\sum} \Big\{ \mathsf{yd}(T\|^{\boldsymbol{b}}_{n+1}) \mid T \in \mathcal{T}(\mathcal{E}, X) \Big\}.$$

The argument for the last step is the same as before: a derivation tree from X first chooses some $c \cdot m \in P_X$ and then has subtrees according to m. These subtrees are in bijection with I_m, so the yield of the (truncated) entire tree is the same as the product over the yields of the (truncated) subtrees $T_{i,j}$ by associativity of infinite products. □

Closed-form solution. Determinstic trees are defined as before. Their construction in Lemma 4.3.2 and the observations in Lemma 4.3.3 and Corollary 4.3.4 are not affected, as they only rely on the finite number of indeterminates. Since trees can be infinite in both height and width, we need a slightly stronger version of Lemma 4.3.5:

Lemma 4.3.5′ (Simple tree, modified)**.** *For each $T \in \mathcal{T}(\mathcal{E}, X)$, there is a derivation tree $T' \in \mathcal{T}(\mathcal{E}, X)$ such that all subtrees rooted at depth ℓ in T' are deterministic and use only monomials $m \in P_Y$ (with $Y \in \boldsymbol{X}$) that occur* on infinitely many levels *of T. Moreover, $|T'|_{m,Y} \leq |T|_{m,Y}$ for all $m \in P_Y, Y \in \boldsymbol{X}$.*

The proof is as before, but $\boldsymbol{X}_\infty$ now contains only those indeterminates that appear at infinitely many levels of T. The prefix up to the first occurrences of $\boldsymbol{X}_\infty$ is still of finite height, but might be of infinite width. Consequently, the elimination step may have to be applied to infinitely many subtrees in parallel, but this does not affect the correctness.

Finally, the proof of Theorem 4.3.1 (Section 4.3.3) goes through with only minor modifications. The argument for least solution still works, by using Lemma 4.3.5′. For greatest solutions, we note that Lemma 4.3.6 still holds (there can be infinitely many subtrees $S_1, S_2, \ldots$ at depth ℓ, but this is not an issue). The proof then works as before, using strong distributivity of $\mathbb{S}^\infty(\boldsymbol{A})$ as there can now be infinitely many nodes $v_1, v_2, \ldots$ at depth ℓ.

4.3.5 Comparison to Newton Iteration

The computation of least fixed points over semirings has already been studied in a more general setting by Esparza, Kiefer, and Luttenberger [EKL10], who consider Newton's method for approximating zeroes of differentiable functions and generalise it to ω-continuous semirings. This may appear surprising at first, given that semirings lack division and subtraction operators to define differentials, but the resulting Newton iteration turns out to be even more robust than Newton's method, as it always converges to the least fixed point. This works for vectors of power series over any ω-continuous semiring, for instance the language semiring $(2^{\Sigma^*}, \cup, \cdot, \emptyset, \{\varepsilon\})$ over a finite alphabet Σ, which is neither absorptive nor commutative. Greatest fixed points are not considered in [EKL10], and these are also the reason for our restriction to absorptive semirings.

Looking at least fixed points of systems of n equations, it has been shown in [EKL10] that the Newton iteration always converges in n steps in idempotent commutative semirings. This applies to absorptive semirings, and in this case also the standard (Kleene) fixed-point iteration converges in n steps, as shown in [EKL11] and as part of Theorem 4.3.1. This raises the question whether the two iterations coincide in absorptive semirings. A negative answer is implicit in [EKL10], where it was observed that the steps of the Newton iteration correspond to derivation trees of increasing *dimension*, whereas Kleene iteration corresponds to trees of increasing *height* (cf. Lemma 4.2.9 (Tree iteration)). The notion of dimension has been introduced in [EKL10], but was later found to be identical to the *Strahler number* of a tree (see [ELS14] for a fascinating survey). These different characterisations suggest that the Newton and Kleene iterations are different also for absorptive semirings, and we give a concrete example below.

Example 4.3.11. Consider the following equation system over $\mathbb{S}^\infty(a, c)$.

$$X = Y^2 + aY \qquad Y = Z^2 + aZ \qquad Z = U^2 + aU \qquad U = c$$

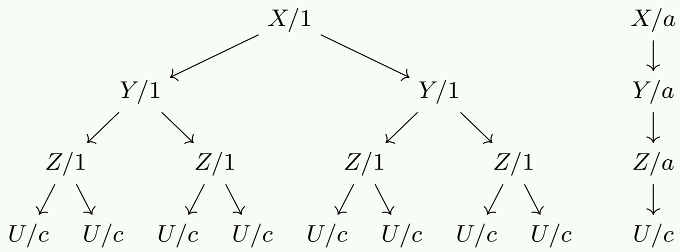

This system has been chosen with the intention of having derivation trees with large dimension (and thus also large height; left tree), but also trees with large height and small dimension (right tree). Due to the large dimension, Newton iteration takes as long as Kleene iteration, but the yield of the tree with small dimension is computed much earlier. We verify this intuition by computing the Kleene iteration:

$$\mathbf{0} \overset{F}{\mapsto} \begin{pmatrix} 0 \\ 0 \\ 0 \\ c \end{pmatrix} \overset{F}{\mapsto} \begin{pmatrix} 0 \\ 0 \\ c^2 + ac \\ c \end{pmatrix} \overset{F}{\mapsto} \begin{pmatrix} 0 \\ c^4 + ac^2 + a^2c \\ c^2 + ac \\ c \end{pmatrix} \overset{F}{\mapsto} \begin{pmatrix} c^8 + ac^4 + a^2c^2 + a^3c \\ c^4 + ac^2 + a^2c \\ c^2 + ac \\ c \end{pmatrix} \circlearrowleft F$$

We have not defined the Newton iteration here, but applied to this example it gives:

$$F(\mathbf{0}) = \begin{pmatrix} 0 \\ 0 \\ 0 \\ c \end{pmatrix} \mapsto \begin{pmatrix} a^2c^2 + a^3c \\ ac^2 + a^2c \\ c^2 + ac \\ c \end{pmatrix} \mapsto \begin{pmatrix} ac^4 + a^2c^2 + a^3c \\ c^4 + ac^2 + a^2c \\ c^2 + ac \\ c \end{pmatrix} \mapsto \begin{pmatrix} c^8 + ac^4 + a^2c^2 + a^3c \\ c^4 + ac^2 + a^2c \\ c^2 + ac \\ c \end{pmatrix} \circlearrowleft$$

Notice that the Newton iteration starts from $F(\mathbf{0})$, so for the comparison we do not count the first step of the Kleene iteration. We see that both iterations take 3 steps to generate the monomial c^8 (corresponding to the left tree), but a^3c (right tree) appears earlier in the Newton iteration and so the two iterations are indeed different.

4.4 Symbolic Computation

This section complements our main result by a second approach, focused specifically on polynomial equation systems over the semiring $\mathbb{S}^\infty(\boldsymbol{A})$, which works by eliminating indeterminates one by one. This may be preferable in cases where the closed-form approach for $\mathbb{S}^\infty(\boldsymbol{A})$ leads to a large blowup in the size of the intermediary values during the computation (cf. Example 4.3.8). For this second approach, we adapt results of Hopkins and Kozen on Kleene algebras [HK99], since we can view absorptive semirings as a special case of Kleene algebras (by setting $a^* = 1$). These results are based on symbolic derivatives of polynomials (which is also the basis for the Newton iteration [EKL10]) to express least solutions. We show that the notion of derivatives can also be used to express greatest solutions, with the help of the infinitary power operation.

4.4.1 Setting and Derivatives

It is convenient to slightly reformulate our problem setting: instead of a system $\mathcal{E}\colon (X = P_X)_{X\in\boldsymbol{X}}$ of polynomials P_X with indeterminates in $\boldsymbol{X}$ and coefficients in $\mathbb{S}^\infty(\boldsymbol{A})$, we now regard P_X as an absorptive polynomial $P_X \in \mathbb{S}^\infty(\boldsymbol{A}\cup\boldsymbol{X})$ (so we no longer distinguish between indeterminates of the polynomial system and indeterminates occurring in coefficients). This allows a more uniform treatment when we eliminate indeterminates one by one, and it is straightforward to see that this does not affect the solutions. This also means that we allow the exponent ∞, so we always work with generalised polynomial equation systems.

To simplify notation, we write $\mathbb{S}^\infty(\boldsymbol{A}, X)$ for $\mathbb{S}^\infty(\boldsymbol{A} \cup \{X\})$. Recall that we write $P(X) \in \mathbb{S}^\infty(\boldsymbol{A}, X)$ to make explicit that X may occur in P; then $P(a) \in \mathbb{S}^\infty(\boldsymbol{A})$ denotes the polynomial that results from $P(X)$ by replacing X with $a \in \mathbb{S}^\infty(\boldsymbol{A})$. In the following, we let $\boldsymbol{V}$ be an arbitrary finite indeterminate set (in particular, any subset of $\boldsymbol{X} \cup \boldsymbol{A}$).

Definition 4.4.1. Let $X \in \boldsymbol{V}$ and $P(X) \in \mathbb{S}^\infty(\boldsymbol{V})$. We denote the *partial derivative* of P with respect to X as P' (leaving X implicit) and define it inductively by

$$X' = 1, \quad Y' = 0 \text{ for } X \neq Y \in \boldsymbol{V},$$
$$(PQ)' = P' \cdot Q + P \cdot Q', \quad (P+Q)' = P' + Q', \quad \big(P^\infty\big)' = P^\infty \cdot P',$$

where $P(X), Q(X) \in \mathbb{S}^\infty(\boldsymbol{V})$.

Notice that the case P^∞ can only apply to $P = X$, as we are working with polynomials in $\mathbb{S}^\infty(\boldsymbol{V})$. However, the following proofs only require the axioms of $^\infty$-algebra (see Section 3.3), so they are not specific to $\mathbb{S}^\infty(\boldsymbol{V})$ and would also hold for arbitrary terms over an $^\infty$-algebra.

4.4.2 Solutions in One Dimension

We first show how least and greatest solutions of a single equation $X = P(X)$ can be expressed using derivatives. Since absorptive semirings can be regarded as Kleene algebras by setting $a^* = 1$ for all elements a, we can follow the proof in [HK99] and only add arguments for the infinitary power.

Lemma 4.4.2. *Let $P(X), Q(X) \in \mathbb{S}^\infty(\boldsymbol{V})$ with $X \in \boldsymbol{V}$, and $a, b, c \in \mathbb{S}^\infty(\boldsymbol{V})$. Then,*

(1) $P(Q)' = P'(Q) \cdot Q'$, *(chain rule, cf. [HK99])*

(2) $P(a+b) = P(a) + P'(a+b) \cdot b$ *(Taylor's theorem, cf. [HK99]),*

(3) $ac \le bc \implies P(a)c \le P(b)c$ *(cf. [HK99]).*

Proof. By structural induction on P. We only consider the case $P = H^\infty$, the proofs for all other cases are identical to [HK99]. For the chain rule, we have

$$\begin{aligned}\big(H^\infty(Q)\big)' &= \big(H(Q)^\infty\big)' = H(Q)^\infty \cdot H(Q)' \\ &= H(Q)^\infty \cdot H'(Q) \cdot H' = \big(H(Q)^\infty\big)' \cdot Q' = \big(H^\infty(Q)\big)' \cdot Q'.\end{aligned}$$

For Taylor's theorem,

$$\begin{aligned} H^\infty(a+b) &= H(a+b)^\infty \\ &= \big(H(a) + H'(a+b)\cdot b\big)^\infty \\ &\overset{(*)}{=} H(a)^\infty + \big(H(a) + H'(a+b)\cdot b\big)^\infty \cdot H'(a+b)\cdot b \\ &= H(a)^\infty + H(a+b)^\infty \cdot H'(a+b)\cdot b \\ &= H^\infty(a) + H^\infty(a+b)\cdot H'(a+b)\cdot b \\ &= H^\infty(a) + (H^\infty)'(a+b)\cdot b. \end{aligned}$$

In $(*)$, we use the fact that $(a+b)^\infty = a^\infty + (a+b)^\infty \cdot b$ is a theorem of $^\infty$-algebra. For (3), the case $P = H^\infty$ is Lemma 3.3.7. □

Using these observations, Hopkins and Kozen prove that the least solution is $P'(P(0))^* \cdot P(0)$. In our setting, this is equal to $P(0)$ and is in fact a direct consequence of absorption (without the need of derivatives). But it serves as inspiration to prove a similar result also for greatest solutions:

Theorem 4.4.3. *Let $P(X) \in \mathbb{S}^\infty(\boldsymbol{V}, X)$. Then $X = P(X)$ has the least solution $P(0)$ and the greatest solution $P(0) + P'(1)^\infty$ in $\mathbb{S}^\infty(\boldsymbol{V})$.*

Proof. For a direct proof of the least solution, let $c_0 \in \mathbb{S}^\infty(\boldsymbol{V})$ be the absolute coefficient of P (i.e., the sum of all monomials not containing X) so that $P(0) = c_0$. For a monomial $m(X)$ containing X, we have $m(c_0) \leq c_0$ by absorption. It follows that $P(c_0) = c_0$ by absorption, so $P(0)$ is the least solution.

For the greatest solution, we first prove that $P(0) + P'(1)^\infty$ is a solution to the inequality $X \leq P(X)$ using (2), (3) of Lemma 4.4.2:

$$\begin{aligned} P\big(P(0) + P'(1)^\infty\big) &\overset{(2)}{=} P(0) + P'\big(P(0) + P'(1)^\infty\big) \cdot \big(P(0) + P'(1)^\infty\big) \\ &\geq P(0) + P'(P'(1)^\infty) \cdot P'(1)^\infty \\ &\overset{(3)}{\geq} P(0) + P'(1) \cdot P'(1)^\infty = P(0) + P'(1)^\infty. \end{aligned}$$

We next show that this is the greatest solution to $X \leq P(X)$. To this end, let $a \in \mathbb{S}^\infty(\boldsymbol{V})$ be a solution, i.e., $a \leq P(a)$. As $1 \geq a$, we get

$$P(0) + P'(1)\cdot a \geq P(0) + P'(a)\cdot a \overset{(2)}{=} P(a) \geq a.$$

Using the axiom ($^\infty$-ind) of $^\infty$-algebra (cf. Definition 3.3.5), we can conclude

$$P(0) + P'(1)^\infty \overset{\text{abs}}{\geq} P(0) + P'(1)^\infty \cdot a \geq a.$$

Finally, note that the greatest solution to $X \leq P(X)$ is also the greatest solution to $X = P(X)$, since the latter is guaranteed to exist in $\mathbb{S}^\infty(\boldsymbol{V})$. □

4.4.3 Solutions of Larger Systems

To solve systems with more than one equation, we eliminate indeterminates one by one, in each step applying Theorem 4.4.3. The main theoretical underpinning is the uniformity of the solutions in one indeterminate, which in our setting follows from the universal property of $\mathbb{S}^\infty(\boldsymbol{V})$.

For the sake of simplicity, we only consider systems of two equations; we can inductively apply the same approach to larger systems. Moreover, we only state the result for greatest solutions, as least solutions are symmetric. We use the notation $\mathbf{gfp}(X \mapsto P(X,Y))$ to refer to the greatest solution of the equation $X = P(X,Y)$ using Theorem 4.4.3 (where we treat the additional indeterminate Y as a coefficient). With this notation, we can formulate the solution of a system in two indeterminates as follows:

Theorem 4.4.4. *Consider the equation system $\mathcal{E}\colon X = P(X,Y), Y = Q(X,Y)$ with $P, Q \in \mathbb{S}^\infty(\boldsymbol{A}, X, Y)$. Let further*

$$\begin{aligned} H(Y) &= \mathbf{gfp}(X \mapsto P(X,Y)) && \in \mathbb{S}^\infty(\boldsymbol{A}, Y),\\ b &= \mathbf{gfp}(Y \mapsto Q(H(Y),Y)) && \in \mathbb{S}^\infty(\boldsymbol{A}). \end{aligned}$$

Then $(H(b), b)$ is the greatest solution of $\mathcal{E}$.

Proof. It is easy to see that $(H(b), b)$ is a solution: By definition of b, we have $Q(H(b), b) = b$. By definition of H, we further have $P(H(Y), Y) = H(Y)$, and by applying the instantiation $Y \mapsto b$ we get $P(H(b), b) = H(b)$.

To prove that $(H(b), b)$ is the greatest solution, we make use of the universal property. Let (c, d) be any solution with $c, d \in \mathbb{S}^\infty(\boldsymbol{A})$. We claim that $H(d) = \mathbf{gfp}(X \mapsto P(X, d))$. To see this, consider the definition of $H(Y)$ and apply the instantiation $Y \mapsto d$. By the universal property, this instantiation is fully continuous and thus preserves greatest fixed points.

Now, since (c, d) is a solution, we have $P(c, d) = c$, and since $H(d)$ is the greatest solution to $X = P(X, d)$, we must have $H(d) \geq c$. Then also $Q(H(d), d) \geq Q(c, d) = d$. Since b is the greatest solution to $Q(H(Y), Y) = Y$ and hence also to $Q(H(Y), Y) \geq Y$, we have $b \geq d$. Finally, $H(b) \geq H(d) \geq c$, so $(H(b), b)$ is indeed the greatest solution. □

We can apply this second computational technique to semirings other than $\mathbb{S}^\infty(\boldsymbol{A})$ by first performing a symbolic abstraction. That is, we replace all coefficients by pairwise different indeterminates from $\boldsymbol{A}$, then compute the solution in $\mathbb{S}^\infty(\boldsymbol{X}, \boldsymbol{A})$ and apply the reverse instantiation (which preserves solutions).

Example 4.4.5. Recall our introductory example in the tropical semiring. By replacing coefficients with indeterminates a, b, c, we obtain the equation system on the right:

$$\begin{aligned} X_a &= 1 +_{\mathbb{R}} X_a & & & X_a &= a \cdot X_a\\ X_b &= \min(1 +_{\mathbb{R}} X_a, 20 +_{\mathbb{R}} X_c) & &\rightsquigarrow & X_b &= a \cdot X_a + b \cdot X_c\\ X_c &= 0 +_{\mathbb{R}} X_c & & & X_c &= c \cdot X_c \end{aligned}$$

We solve the system over $\mathbb{S}^\infty(X_a, X_b, X_c, a, b, c)$ by eliminating indeterminates:

- $\mathbf{gfp}(X_a \mapsto a \cdot X_a) = 0 + a^\infty = a^\infty$ (by Theorem 4.4.3)

- $\mathbf{gfp}(X_b \mapsto a \cdot a^\infty + b \cdot X_c) = a^\infty + b \cdot X_c$ (we first instantiate X_a by a^∞)
- $\mathbf{gfp}(X_c \mapsto c \cdot X_c) = c^\infty$

The greatest solution is thus $X_a = a^\infty$, $X_b = a^\infty + bc^\infty$, $X_c = c^\infty$. Applying the reverse substitution, we get the expected solution $(\infty, 20, 0)$ in the tropical semiring.

Usually, the closed-form solution in Theorem 4.3.1 is preferable, as we can work directly in the target semiring. The symbolic technique is best suited to compute solutions in $\mathbb{S}^\infty(\boldsymbol{A})$, which is of interest for semiring provenance analysis. Compared to the closed-form solution, we need fewer computation steps and can often avoid an intermediate blowup in the size of the polynomials.

Example 4.4.6. Recall the equation system $X_1 = b \cdot X_2$, $X_2 = (b+c) \cdot X_2X_3$, $X_3 = a \cdot X_1$ of Example 4.3.8. Eliminating indeterminates one by one, we get

- $\mathbf{gfp}(X_1 \mapsto b \cdot X_2) = b \cdot X_2$
- $\mathbf{gfp}(X_2 \mapsto (b+c) \cdot X_2X_3) = \big((b+c)X_3\big)^\infty = b^\infty X_3^\infty + c^\infty X_3^\infty$
- $\mathbf{gfp}(X_3 \mapsto ab \cdot (b^\infty X_3^\infty + c^\infty X_3^\infty)) = a^\infty b^\infty + a^\infty b^\infty c^\infty = a^\infty b^\infty$

and by substituting the results backwards, we obtain $X_1 = X_2 = X_3 = a^\infty b^\infty$.

Remark 4.4.7. The results in this section are not specific to $\mathbb{S}^\infty(\boldsymbol{A})$. Indeed, the proof in [HK99] holds for polynomials over any idempotent, commutative Kleene algebra, and one can similarly generalise our results by representing polynomial systems using absorptive polynomials $\mathrm{AbsPoly}(K, \mathbb{N}^\infty, \boldsymbol{X})$ with other coefficient semirings (cf. Section 3.7). However, we already have the closed-form solution for computing greatest fixed points directly in K, and for provenance application $\mathbb{S}^\infty(\boldsymbol{A})$ is sufficient.

It remains an open question whether a purely algebraic proof for greatest solutions is possible, without continuity arguments (see the axiomatisation of the infinitary power in Section 3.3 for a first step in this direction).

4.5 Nested Fixed Points

Polynomial equation systems arise from the evaluation of formulae in fixed-point logics under semiring semantics, so our methods to compute (greatest) fixed points enable us to evaluate formulae with a single fixed-point operator. Fixed-point logics such as LFP or L_μ (see Section 5.2) also permit *nested* fixed-point operators. These do not immediately translate to polynomial systems, but our methods can still be applied without modifications. The reason is that a nested inner fixed point can always be described by a (generalised) polynomial. We illustrate this argument on an example.

Running example. Throughout this section, we consider the following nested system:

$$\mathcal{E}_1\colon \begin{bmatrix} X_1 = cX_2^2 + Y_1^{(\nu)} \\ X_2 = Y_2^{(\nu)} + d \end{bmatrix}, \qquad \mathcal{E}_2^{(\boldsymbol{X})}\colon \begin{bmatrix} Y_1 = X_2Y_1 + aY_2 \\ Y_2 = bY_2 + X_2 \end{bmatrix},$$

with coefficients $a, b, c, d \in K$ in some absorptive, fully-continuous semiring. We interpret $\boldsymbol{Y}^{(\nu)} = (Y_1^{(\nu)}, Y_2^{(\nu)})$ as the greatest solution to $\mathcal{E}_2$. Notice that X_1, X_2 can occur in $\mathcal{E}_2$, so this greatest solution depends on the values of X_1, X_2 (in this example only on X_2). Our goal is to solve the system $\mathcal{E}_1$. We also refer to $\mathcal{E}_1$ as the *outer*, and $\mathcal{E}_2$ as the *inner* system.

Simultaneous fixed points. First assume we are interested in the greatest solution to $\mathcal{E}_1$. While the arguments for dealing with alternations below also apply here, there is a particularly simple solution due to the well-known Bekić principle (see, e.g., [AN01, Section 1.4.2]). This applies to monotone functions on complete lattices, which covers our setting, and intuitively states that nested fixed points of the same kind are equivalent to simultaneous fixed points. In our example, this means that the greatest solution to $\mathcal{E}_1$ is equal to the greatest solution of the combined system (which happens to be $\mathbf{1}$):

$$\mathcal{E}_{1,2}\colon \begin{bmatrix} X_1 = cX_2^2 + Y_1 \\ X_2 = Y_2 + d \\ Y_1 = X_2Y_1 + aY_2 \\ Y_2 = bY_2 + X_2 \end{bmatrix}.$$

Notice that the size of the combined system grows linearly with the nesting depth.

Alternating fixed points. Now assume we are interested in the least solution to $\mathcal{E}_1$, so we are dealing with the more challenging case of alternating fixed points. For a given instantiation $a_1, a_2 \in K$ of X_1, X_2, we can compute the inner solution $\mathbf{gfp}(F_{\mathcal{E}_2^{(a_1,a_2)}})$ as usual. In contrast, the outer system $\mathcal{E}_1$ is not defined by polynomials (since it includes the inner greatest fixed point) and it is thus not clear whether our techniques are applicable. However, we show that we can rewrite $\mathcal{E}_1$ as a polynomial system $\mathcal{E}_1'$, by first solving $\mathcal{E}_2$ symbolically.

We construct $\mathcal{E}_1'$ as follows. Consider the symbolic abstraction of $\mathcal{E}_2$ and treat X_1, X_2 as coefficients. In our case, we let $\boldsymbol{A} = \{a, b, c, d\}$ and view a, b, c, d and X_1, X_2 as indeterminates, so that we have a system $\mathcal{E}_2'$ over $\mathbb{S}^\infty(\boldsymbol{A}, \boldsymbol{X})$. Then compute the greatest solution to $\mathcal{E}_2'$ in $\mathbb{S}^\infty(\boldsymbol{A}, \boldsymbol{X})$ (using either the closed-form or the symbolic approach):

$$\mathcal{E}_2'\colon \begin{bmatrix} Y_1 = X_2Y_1 + aY_2 \\ Y_2 = bY_2 + X_2 \end{bmatrix}, \qquad \mathbf{gfp}(\mathcal{E}_2') = \begin{pmatrix} X_2^\infty + ab^\infty + aX_2 \\ b^\infty + X_2 \end{pmatrix}.$$

Recall that the greatest solution holds uniformly under all instantiations of $\boldsymbol{A}$ and X_2, due to the universal property of $\mathbb{S}^\infty(\boldsymbol{A}, \boldsymbol{X})$ and Lemma 4.1.2. We can thus undo the abstraction (i.e., replace $\boldsymbol{A}$ by their original values $a, b, c, d \in K$) and instantiate $\boldsymbol{Y}^{(\nu)}$ by the result to obtain the desired system $\mathcal{E}_1'$:

$$\mathcal{E}_1'\colon \begin{bmatrix} X_1 = cX_2^2 + X_2^\infty + ab^\infty + aX_2 \\ X_2 = b^\infty + X_2 + d \end{bmatrix}.$$

This is indeed a (generalised) polynomial system over K, which is equivalent to $\mathcal{E}_1$ by our construction, so we can now compute the least solution as usual:

$$\mathbf{lfp}(F_{\mathcal{E}_1}) = \mathbf{lfp}(F_{\mathcal{E}_1'}) = \begin{pmatrix} b^\infty + d^\infty + ad + cd^2 \\ b^\infty + d \end{pmatrix}.$$

This approach requires the computation of a potentially large polynomial $\mathbf{gfp}(\mathcal{E}_2')$ in $\mathbb{S}^\infty(\boldsymbol{A},\boldsymbol{X})$, which may be undesirable when we want to compute the overall solution in some other semiring K. In this case, recall that the closed-form solution does not depend on the actual polynomials of the polynomial system, but only requires us to evaluate the associated operator F. We can thus forget about $\mathcal{E}_1'$ and $\mathbf{gfp}(\mathcal{E}_2')$ altogether (we only need their existence) and evaluate F directly on $\mathcal{E}_1$. More precisely, we use the following operator:

$$F_{\mathcal{E}_1}\colon K^2 \to K^2, \quad \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} \mapsto \begin{pmatrix} c \cdot v_2^2 + \mathbf{gfp}(\mathcal{E}_2^{(v_1,v_2)})_1 \\ \mathbf{gfp}(\mathcal{E}_2^{(v_1,v_2)})_2 + d \end{pmatrix}.$$

That is, for each iteration of the outer system, we have to recompute the solution to the inner system (with new values for X_1, X_2). This is the same as the naive evaluation of nested fixed points in the classical Boolean setting. In other words, the generalisation from the Boolean to absorptive semirings does not cause additional complications for nested fixed points. We remark that we have focused on the running example only for the presentation, our arguments can be repeated inductively to cover nested fixed points of arbitrary alternation depth.

4.6 Outlook

We have presented two methods to compute least and, most importantly, greatest solutions of polynomial equation systems over absorptive, fully-continuous semirings. Both methods require only polynomially many applications of the semiring operations and the infinitary power, in terms of the number of equations. The main result of this chapter is a closed-form solution that works in any absorptive, fully-continuous semiring and is as easy as computing the standard fixed-point iteration with an added application of the infinitary power. The proof is based on an analysis of infinite derivation trees: because of absorption, only trees of a particularly simple shape need to be considered for the computation.

These methods can also be used to compute nested fixed points with the naive recursive computation. In the classical Boolean setting, recent breakthroughs in solving parity games have lead to algorithms computing nested fixed points in quasipolynomial time, roughly $\mathcal{O}(n^{\log d})$ instead of the naive $\mathcal{O}(n^d)$ for systems of n equations and alternation depth d. Hausmann and Schröder [HS21] have described a progress-measure algorithm based on universal graphs that solves systems of equations over finite lattices in quasipolynomial time. At the same time, Arnold, Niwiński, and Parys [ANP21] independently found a quasipolynomial-time algorithm based on universal trees that computes nested fixed points of the form $\mu z.\nu y.\mu x.f(x,y,z)$ for a monotone function f on a powerset lattice. It is an interesting question whether the techniques using universal trees can also be used in our setting. More precisely, consider a nested system of the form (assuming even d):

$$\mu \boldsymbol{X}^{(1)}.\ \nu \boldsymbol{X}^{(2)}.\ \mu \boldsymbol{X}^{(3)} \dots\ \nu \boldsymbol{X}^{(d)}.\ F(\boldsymbol{X}^{(1)}, \dots, \boldsymbol{X}^{(d)}),$$

where each $\boldsymbol{X}^{(i)}$ is a tuple of n indeterminates. The semantics is defined as in [ANP21], but we additionally assume that F is defined by polynomials of the form $P_i(\boldsymbol{X}^{(1)}, \dots, \boldsymbol{X}^{(d)})$, i.e.,

$$F\colon (K^n)^d \to K^n, \qquad F(\boldsymbol{a}^{(1)}, \dots, \boldsymbol{a}^{(d)}) = \begin{pmatrix} P_1(\boldsymbol{a}^{(1)}, \dots, \boldsymbol{a}^{(d)}) \\ P_2(\boldsymbol{a}^{(1)}, \dots, \boldsymbol{a}^{(d)}) \\ \vdots \\ P_n(\boldsymbol{a}^{(1)}, \dots, \boldsymbol{a}^{(d)}) \end{pmatrix}.$$

Open Question. *Given polynomials $P_1, \dots, P_n$ over an absorptive, fully-continuous semiring, can the above expression be computed using only quasipolynomially many evaluations of F (with respect to n and d)?*

Let us briefly discuss why this might be an interesting question. Both [HS21] and [ANP21] rely on the finiteness of the domain. For instance, an important argument in [ANP21] is that an n-dimensional Boolean vector can only be increased n times during the computation, which is used to show that universal trees of width n suffice for the computation. This argument fails in our setting, since absorptive semirings may have infinite ascending or descending chains, so it is not clear how their proof could be extended. On the other hand, we have shown that the fixed-point iteration stagnates after n steps for least fixed points, and greatest fixed points are only slightly more difficult to compute. So perhaps this provides a sufficient bound on the number of iterations to make a quasipolynomial-time algorithm work.

A good starting point might be the (dual version of the) *asymmetric* algorithm in [ANP21], which applies universal trees only to the simpler least fixed-point iteration, resulting in a system for which one could then compute the greatest fixed point by Theorem 4.3.1.

A further question for future work is the extension to more general semirings. Motivated by the problem of evaluating semiring semantics for LFP, we have here focussed on absorptive semirings. However, the Newton iteration for least solutions applies to a much more general setting, including non-absorptive and non-commutative semirings [EKL10], and it would be interesting to see if also greatest solutions can be computed in such semirings (provided they exist). Since absorption plays a crucial role in our simplification of derivation trees, such a generalisation likely requires new arguments.

Chapter 5

Semiring Semantics

Semiring provenance aims at explaining the truth of a formula, or the answers to a database query, by evaluating formulae or queries in a semiring semantics that generalises classical Boolean semantics by permitting multiple truth values from some commutative semiring. These additional truth values, also called *nuances of truth* or *shades of truth* can then be used to encode additional information such as costs, bag semantics, or provenance information.

This idea was first proposed in a seminal paper by Green, Karvounarakis, and Tannen [GKT07], who defined a semiring semantics for positive relational algebra and datalog. The premise of this semantics is simple: starting from annotated tuples in the input database (called K-relations, for a semiring K), these annotations are propagated through the evaluation of the query to produce a semiring annotation for each answer tuple. For this propagation, joint use of information (e.g., conjunction, natural join) is interpreted by semiring multiplication, and alternative use of information (e.g., disjunction, union) is interpreted by semiring addition. This works for any commutative semiring, or for all ω-continuous semirings in case of datalog, which leads to a very general framework. By using different semirings, one obtains bag semantics (in $\mathbb{N}$ or $\mathbb{N}^\infty$), applications such as cost computation (in the tropical semiring), classical semantics (in $\mathbb{B}$), and different forms of provenance information (in polynomial semirings such as $\mathbb{N}[X]$ or $\mathrm{PosBool}(X)$).

Both positive relational algebra and datalog are negation-free languages and this is not a coincidence: semirings only have two operations and when these are used as described above, there remains no general way to define a negation operator in an arbitrary commutative semiring. How can this obstacle be overcome to extend semiring semantics to logics with negation, such as first-order logic and stratified datalog? This has been addressed by Grädel and Tannen [GT17a, GT] for first-order logic based on two ideas. First, semiring semantics is defined only for formulae in *negation normal form*, and semiring annotations (called K-interpretations) are provided not only for positive, but also for negative literals. As a consequence, any form of desired consistency between positive and negative literals has to be imposed on the K-interpretations. One possibility is to work with *model-defining* K-interpretations, where for each literal L, exactly one of the annotations for L and its negation $\neg L$ must be zero. This works for many applications, but is not satisfactory for provenance tasks such as computing repairs, where it is not clear a priori whether a literal or its negation should be true. For this case, Grädel and Tannen propose semirings of *dual-indeterminate* polynomials, with the idea that an indeterminate x and its dual $\overline{x}$ cancel each other and can thus be used as annotations for opposing literals.

Based on these results, semiring semantics has been extended to several logics and the semantics became interesting not only for provenance analysis, but also in its own right by studying its model-theoretic properties. We discuss the model theory of semiring semantics in Section 7.1; this chapter provides an overview of semiring semantics for several logics relevant to this thesis, most importantly first-order logic (FO), least fixed-point logic (LFP), and infinitary logic ($L_{\infty\omega}$). A novel result is the embedding of LFP into the newly defined semiring semantics for infinitary logics (Section 5.3), which makes use of the fixed-point computation in Chapter 4. We close this chapter with a brief survey on semiring semantics and ideas for future research.

5.1 First-Order Logic

We begin with an overview of semiring semantics for first-order logic (FO), following [GT17a, GT]. The syntax of FO-formulae is standard (see Chapter 2 and [EF95]), it is only the models and the semantics that is modified. Throughout this section, τ denotes a finite relational signature and K a commutative semiring.

5.1.1 *K*-Interpretations

We first generalise the models from classical τ-structures to structures annotated with semiring values. For a finite universe A, we write $\mathrm{Atoms}_A(\tau) = \{R\boldsymbol{a} \mid R \in \tau,\ \boldsymbol{a} \in A^{\mathrm{arity}(R)}\}$ for the set of *instantiated* atoms, and $\mathrm{NegAtoms}_A(\tau)$ for their negations. The set of all instantiated literals is denoted as $\mathrm{Lit}_A(\tau) = \mathrm{Atoms}_A(\tau) \cup \mathrm{NegAtoms}_A(\tau)$. For a literal $L \in \mathrm{Lit}_A(\tau)$, we write $\overline{L}$ for the complementary literal. We then annotate all instantiated literals by semiring values, where false literals are usually annotated by 0, and true literals by some positive semiring value.

Definition 5.1.1 (Interpretation)**.** A K*-interpretation* (of signature τ and universe A) is a mapping $\pi\colon \mathrm{Lit}_A(\tau) \to K$ into a commutative semiring K.

We say that π is *model-defining* if exactly one of the values $\pi(L)$, $\pi(\overline{L})$ equals 0 for every pair of complementary literals $L, \overline{L} \in \mathrm{Lit}_A(\tau)$. In this case, π induces a classical τ-structure $\mathfrak{A}_\pi$ with universe A such that $\mathfrak{A}_\pi \models L$ if, and only if, $\pi(L) \neq 0$.

Since we work with formulae in negation normal form, the semantics of negation is pushed to the interpretations $\pi(L)$ and $\pi(\overline{L})$ of complementary literals. This is why we mostly work with model-defining interpretations, which retain the Boolean semantics in the sense that they choose which of L and $\overline{L}$ should be true by assigning a positive value. This is not the only possibility; one can, for instance, only require $\pi(L) \cdot \pi(\overline{L})$ for each literal, or one can use model-compatible interpretations (see Definition 5.1.7 below).

A K-interpretation $\pi\colon \mathrm{Lit}_A(\tau) \to K$ extends in a straightforward way to an interpretation $\pi[\![\varphi(\boldsymbol{a})]\!]$ of instantiated FO-formulae. We first set $\pi[\![L]\!] = \pi(L)$ for literals $L \in \mathrm{Lit}_A(\tau)$ and map equalities and inequalities to their truth values (using the neutral elements of K):

$$\pi[\![a = b]\!] := \begin{cases} 1 & \text{if } a = b, \\ 0 & \text{if } a \neq b, \end{cases} \qquad \text{and} \qquad \pi[\![a \neq b]\!] := \begin{cases} 0 & \text{if } a = b, \\ 1 & \text{if } a \neq b. \end{cases}$$

For connectives and quantifiers, we follow the *joint versus alternative use* paradigm and interpret them by the semiring operations:

$$\pi[\![\psi(\boldsymbol{a}) \vee \vartheta(\boldsymbol{a})]\!] := \pi[\![\psi(\boldsymbol{a})]\!] + \pi[\![\vartheta(\boldsymbol{a})]\!], \qquad \pi[\![\psi(\boldsymbol{a}) \wedge \vartheta(\boldsymbol{a})]\!] := \pi[\![\psi(\boldsymbol{a})]\!] \cdot \pi[\![\vartheta(\boldsymbol{a})]\!],$$
$$\pi[\![\exists x\, \vartheta(\boldsymbol{a}, x)]\!] := \sum_{b \in A} \pi[\![\vartheta(\boldsymbol{a}, b)]\!], \qquad \pi[\![\forall x\, \vartheta(\boldsymbol{a}, x)]\!] := \prod_{b \in A} \pi[\![\vartheta(\boldsymbol{a}, b)]\!].$$

This defines a semiring semantics for formulae in negation normal form (nnf). It is sometimes convenient to write formulae with negation, for which we simply set

$$\pi[\![\neg\varphi]\!] := \pi[\![\mathrm{nnf}(\neg\varphi)]\!].$$

Definition 5.1.2. Semiring semantics of $\mathrm{FO}(\tau)$ is defined by $\pi[\![\cdot]\!]$ for K-interpretations π over finite universe and finite, relational signature τ, assigning to each instantiated formula $\varphi(\boldsymbol{a})$ a semiring value.

We emphasize that only finite universes are considered, so that the semantics of quantifiers is well-defined in any semiring K. This is inherited from the database setting, where provenance for finite databases was studied. More recently, a theory of semirings with infinitary operations has been developed (cf. Section 3.4) that enables an extension to infinite universes (see [BGMN24]), but for the scope of this thesis we stick to the setting of finite models.

Equivalence of formulae now takes into account the semiring values and is thus more fine-grained than Boolean equivalence.

Definition 5.1.3 (Equivalence)**.** Two formulae $\psi(\boldsymbol{x})$, $\varphi(\boldsymbol{x})$ are *K-equivalent*, denoted $\psi \equiv_K \varphi$, if $\pi[\![\psi(\boldsymbol{a})]\!] = \pi[\![\varphi(\boldsymbol{a})]\!]$ for every model-defining K-interpretation π (over finite universe A) and every tuple $\boldsymbol{a} \subseteq A$.

Notice that only model-defining interpretations are considered for equivalence. This means that equivalence $\equiv_{\mathbb{B}}$ over the Boolean semiring coincides with the standard notion of logical equivalence, which serves as a first sanity check for our semantics. Other semirings may distinguish more formulae. For example, $\varphi \vee \varphi \not\equiv_{\mathbb{N}} \varphi$ since addition of natural numbers is not idempotent. Many algebraic properties of the semiring directly translate to equivalences, for instance $\varphi \vee (\varphi \wedge \vartheta) \equiv_K \varphi$ holds precisely in the absorptive semirings (where $a + ab = a$).

Standard concepts such as isomorphism and elementary equivalence naturally extend to K-interpretations, but their relationship is different under semiring semantics (see Section 7.1).

5.1.2 Main Properties

Arguably the most important property of semiring semantics is the interplay with homomorphisms, called *fundamental property* in [GT17a].

Proposition 5.1.4 (Fundamental property)**.** *Let $h\colon K_1 \to K_2$ be a semiring homomorphism. For every K_1-interpretation π, the mapping $h \circ \pi$ is a K_2-interpretation, and for every sentence $\varphi \in \mathrm{FO}$, we have $h(\pi[\![\varphi]\!]) = (h \circ \pi)[\![\varphi]\!]$.*

As diagram:

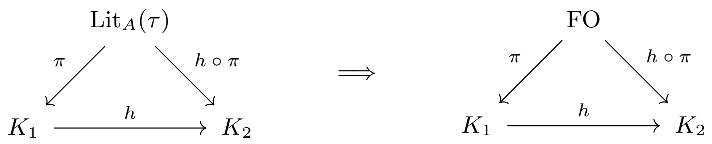

From a provenance perspective, this means that freely generated semirings are the most informative ones, as they permit homomorphisms into any other semiring. For first-order logic, the most general such semirings are the free commutative semirings $\mathbb{N}[X]$.

Example 5.1.5. For a simple example, consider the following graph with a unary relation P. We first consider the $\mathbb{N}[X]$-interpretation π shown on the left, where edges are annotated by indeterminates $X = \{a, b, c, d, e\}$ and the relation P is not tracked (i.e., mapped to 0/1).

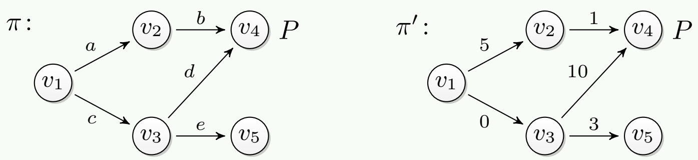

More precisely, π is defined as follows:

$$\pi(Ev_1v_2) = a, \qquad \pi(Ev_3v_4) = d, \qquad \pi(Pv_4) = 1,$$
$$\pi(Ev_2v_4) = b, \qquad \pi(Ev_3v_5) = e, \qquad \pi(Pv_i) = 0 \text{ otherwise},$$
$$\pi(Ev_1v_3) = c, \qquad \pi(Ev_iv_j) = 0 \text{ otherwise},$$

and negative literals are mapped to 0 or 1 so that π is model-defining. For the formula φ evaluated at v_1, we obtain the following polynomial in $\mathbb{N}[X]$:

$$\varphi(x) := \exists y \exists z (Exy \wedge Eyz \wedge Pz), \qquad \pi[\![\varphi(v_1)]\!] = ab + cd.$$

This tells us that there are two ways to satisfy the formula, by using either edges a and b, or edges c and d.

Now consider the $\mathbb{T}$-interpretation π' shown on the right, where edges have been annotated with cost values in the tropical semiring (and P is not tracked, i.e., mapped to $\infty/0$). Without evaluating φ again, we can compute its semiring value by simply applying the homomorphism h induced by the assignment $a \mapsto 5$, $b \mapsto 1$, $c \mapsto 0$, $d \mapsto 10$, $e \mapsto 3$:

$$\pi'[\![\varphi(v_1)]\!] = h\left(\pi[\![\varphi(v_1)]\!]\right) = h(ab + cd) = \min(5 + 1, 0 + 10) = 6.$$

We see that the minimal cost to prove $\varphi(v_1)$ is 6 (assuming we have to pay for the use of each literal). Similarly, we can consider the counting interpretation $\pi_\#$ that maps every true literal to 1 just as the underlying Boolean structure, but interpreted in the semiring $\mathbb{N}$. This amounts to counting the number of proof trees (here: paths from v_1 to P), which can again be obtained from $ab + cd$ by applying the homomorphism $h_\#$ induced by $a, b, c, d, e \mapsto 1$:

$$\pi_\#[\![\varphi(v_1)]\!] = h_\#(ab + cd) = 1 + 1 = 2.$$

A second key property is that model-defining interpretations are *truth-preserving*, in the sense that any formula φ that is true in the induced Boolean interpretation, i.e. $\mathfrak{A}_\pi \models \varphi$, is mapped to a nonzero semiring value $\pi[\![\varphi]\!]$, under the assumption that the semiring is positive. Recall that a positive semiring is one where $a + b = 0$ implies $a, b = 0$, and further $ab = 0$ implies $a = 0$ or $b = 0$, so that the value 0 cannot be produced from nonzero values. An elegant way to prove this property is by considering the *truth projection* into the Boolean semiring:

$$\dagger_K : K \to \mathbb{B}, \qquad a \mapsto \begin{cases} 1, & \text{if } a \neq 0, \\ 0, & \text{if } a = 0. \end{cases}$$

This is a homomorphism if, and only if, K is positive. If π is model-defining, then $\dagger_K \circ \pi$ coincides with $\mathfrak{A}_\pi$ and we obtain the following corollary to Proposition 5.1.4:

Corollary 5.1.6 (Truth preserving). *Let π be a model-defining K-interpretation in a positive semiring K. Then for every sentence $\varphi \in \mathrm{FO}$,*

$$\pi[\![\varphi]\!] \neq 0 \quad \Longleftrightarrow \quad \mathfrak{A}_\pi \models \varphi.$$

Lastly, we mention that semiring semantics can also be characterised by proof trees. Grädel and Tannen show in [GT] that the value $\pi[\![\psi]\!]$ is equal to the sum over the values of all proof trees, where the value of a proof tree is the product over all literals occurring in the tree. This justifies the cost computation and path counting in Example 5.1.5 and also means that, in general, we can always count the number of proof trees in the semiring $\mathbb{N}$.

5.1.3 Dual Indeterminates

Typical provenance applications are the explanation of missing answers and the computation of (minimal) repairs (see, e.g., [CJ09, XZAT18]). For instance, say we have a graph that fails to satisfy a property ψ. We are given a set of edges that may be added and a set of edges that may be removed, and we want to find the minimal sets of edges to modify so that ψ holds. When we model this as a K-interpretation, we choose a suitable polynomial semiring and map edges that may be modified to indeterminates. However, we run into the problem of not knowing upfront whether Euv or $\neg Euv$ should be true, so the only option is to map both to different indeterminates. This has the downside that the resulting interpretation is no longer model-defining, and the constraint that only one of Euv and $\neg Euv$ can be true is lost. To resolve this issue, Grädel and Tannen introduce *dual indeterminates* [GT17a, XZAT18].

Let X be a finite set of indeterminates. We define $\overline{X} = \{\overline{x} \mid x \in X\}$ to be a copy of X, called the *dual* version of X. Given a polynomial semiring such as $\mathbb{N}[X]$, we then define its dual-indeterminate version $\mathbb{N}[X, \overline{X}]$ as the quotient of $\mathbb{N}[X \cup \overline{X}]$ by the congruence generated by equalities $x \cdot \overline{x} = 0$ for all $x \in X$, formally $\mathbb{N}[X, \overline{X}] := \mathbb{N}[X \cup \overline{X}]/_{\langle x\overline{x}=0\rangle}$. Intuitively, the semiring operations work as in $\mathbb{N}[X \cup \overline{X}]$, but after each operation we drop all monomials that contain both an indeterminate x and its dual $\overline{x}$. For example, $x \cdot (\overline{x}y + xz) = x^2 z$. Notice that the resulting semiring is no longer positive, as $x \cdot \overline{x} = 0$ gives a divisor of zero. The quotient semiring inherits a universal property; we simply require that the assignment of the indeterminates factors through the quotient construction: every mapping $h\colon X \cup \overline{X} \to K$ that respects dual-indeterminates, i.e., $h(x) \cdot h(\overline{x}) = 0$ for all $x \in X$, extends uniquely to a semiring homomorphism $h\colon \mathbb{N}[X,\overline{X}] \to K$. The same quotient construction can be applied to other semirings of polynomials, including $\mathrm{PosBool}(X,\overline{X})$ and $\mathbb{S}^\infty(X,\overline{X})$.

Instead of model-defining interpretations, we now allow interpretations that use dual indeterminates for complementary literals:

Definition 5.1.7 (Model-compatible)**.** A $\mathbb{N}[X,\overline{X}]$-interpretation is called *model-compatible* if for each positive atom $L \in \mathrm{Atoms}_A(\tau)$, either

- $\pi(L) = x$ and $\pi(\neg L) = \overline{x}$, for some $x \in X$, or
- one of $\pi(L)$, $\pi(\neg L)$ is 1, the other one 0.

We say that π is *non-overlapping* if no two literals are mapped to the same indeterminate.

We use the same definition also for other semirings of dual-indeterminate polynomials. We have seen that model-defining interpretations π are truth-preserving in positive semirings, so that nonzero values are given to those formulae satisfied by $\mathfrak{A}_\pi$. For model-compatible interpretations, semiring semantics captures satisfiability in an entire class of classical models. Formally, let π be a model-compatible interpretation. A classical structure $\mathfrak{A}$ is *compatible* with π if $\mathfrak{A} \models L$ for all literals with $\pi(L) = 1$. Writing Mod_π for the set of structures compatible with π, semiring semantics in π can be characterised as follows [GT17a].

Proposition 5.1.8. *Let π be a non-overlapping model-compatible $\mathbb{N}[X,\overline{X}]$-interpretation and φ a FO-sentence. Then*

(1) *π is Mod_π-satisfiable if, and only if, $\pi[\![\varphi]\!] \neq 0$,*

(2) *π is Mod_π-valid if, and only if, $\pi[\![\neg\varphi]\!] = 0$.*

This follows from applying the universal property to assignments $h\colon X \cup \overline{X} \to \mathbb{B}$ corresponding to the compatible models. This argument applies to all dual-indeterminate polynomial semirings, in particular to $\mathrm{PosBool}(X,\overline{X})$. Since this is the simplest such semiring, it is the right choice to answer questions about Mod_π-satisfiability. In fact, the argument is much stronger: assume $\pi[\![\psi]\!] = m_1 + \cdots + m_k$ is some polynomial in $\mathrm{PosBool}(X,\overline{X})$. Since an assignment h corresponding to a satisfying structure must evaluate to 1, this polynomial must contain a monomial m_i consisting only of indeterminates with $h(x) = 1$, so that $h(m_i) = 1$. This means that the monomials in $\pi[\![\psi]\!]$ closely correspond to the models within Mod_π. For a more rigorous argument, we refer to Section 6.3.5 where this reasoning is used for repairs in Büchi games. Here we illustrate this approach on a small example.

Example 5.1.9. Consider the graph on the left, which fails to satisfy the FO-sentence ψ.

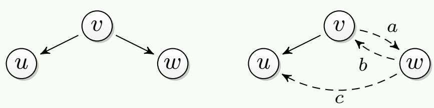

$$\psi = \exists x \exists y (Exy \wedge \forall z (z = y \vee \neg Exz))$$

Say we fix the edge $v \to u$, but are allowed to remove the edge $v \to w$ or to additionally add edges $w \to v$, $w \to u$. We then consider the $\mathrm{PosBool}(X,\overline{X})$-interpretation π shown in the middle, with dual-indeterminates $X = \{a, b, c\}$ used for the edges we may modify. For example, $\pi(Evw) = a$ and $\pi(\neg Evw) = \overline{a}$. The edge $v \to w$ is not tracked and simply mapped to 1. Evaluating ψ results in a polynomial from which we can read off the repairs:

$$\pi[\![\psi]\!] = \overline{a} + \overline{b}c + b\overline{c}.$$

That is, we can either remove edge a, or we can add exactly one of b and c. This indeed covers all graphs satisfying ψ within Mod_π, i.e., all possible repairs under the given constraints.

5.2 Least Fixed-Point Logic

Least fixed point logic (LFP) extends first-order logic by least and greatest fixed points of definable monotone operators. Least fixed points have already been considered in the original proposal of semiring semantics for datalog [GKT07], where ω-continuous semirings were used to obtain a well-defined semantics. Adding greatest fixed points (and also arbitrary interleavings of least and greatest fixed points), leads to a number of complications that we discuss in this section. It turns out that, although the semantics can be defined in any fully-continuous semiring, *absorptive* semirings and particularly generalised absorptive polynomials $\mathbb{S}^\infty(X)$ are the right provenance semirings for LFP.

This section summarises the semiring semantics for LFP that has been developed in [Naa19, DGNT21] and establishes the connection to the fixed-point computation in Chapter 4.

5.2.1 Semantics

We briefly recall the classical semantics of LFP. As for FO, we assume a finite relational signature τ and only consider finite structures. The syntax of LFP extends FO by adding operators for least and greatest fixed points: if $\varphi(R, \boldsymbol{x})$ is a formula of signature $\tau \cup \{R\}$ in which the relational variable R occurs only positively, and the length of $\boldsymbol{x}, \boldsymbol{y}$ matches the arity of R, then $[\mathbf{lfp}\, R\, \boldsymbol{x}.\ \varphi](\boldsymbol{y})$ and $[\mathbf{gfp}\, R\, \boldsymbol{x}.\ \varphi](\boldsymbol{y})$ are formulae of signature τ. The semantics of these formulae is that $\boldsymbol{y}$ is contained in the least (or greatest) fixed point of the update operator $F_\varphi \colon R \mapsto \{\boldsymbol{a} \mid \psi(R, \boldsymbol{a})\}$. Due to the positivity of R in ψ, any such operator F_φ is monotone and has, by the Knaster-Tarski theorem, a least fixed point $\mathbf{lfp}(F_\varphi)$ and a greatest fixed point $\mathbf{gfp}(F_\psi)$. We refer to [GKL+07, EF95] for more background on LFP. The duality between least and greatest fixed points implies that $[\mathbf{gfp}\, R\, \boldsymbol{x}.\ \varphi](\boldsymbol{x}) \equiv \neg[\mathbf{lfp}\, R\, \boldsymbol{x}.\ \neg\varphi[R/\neg R]](\boldsymbol{x})$, where $\varphi[R/\neg R]$ denotes the replacement of each positive R-atom by its negation. (One can thus define LFP by only adding **lfp**-formulae, which explains the name.) By this duality together with de Morgan's laws, every LFP-formula can be brought into negation normal form. The positive fragment of LFP, denoted posLFP, consists of the formulae in negation normal form with only least fixed-point operators.

We note that fixed-point formulae can be nested. A particular source of complexity (or expressive power) are interleavings of least and greatest fixed points, e.g.:

$$\psi(\boldsymbol{a}) = \big[\mathbf{lfp}\, R\, \boldsymbol{x}.\ \vartheta(R, \boldsymbol{x}) \vee [\mathbf{gfp}\, P\, \boldsymbol{y}.\ \varphi(R, P, \boldsymbol{y})](\boldsymbol{x})\big](\boldsymbol{a}).$$

Assuming that R occurs in φ, this is an instance of a *fixed-point alternation*, and the number of such alternations is called the *alternation depth* (see, e.g., [GKL+07, Chapter 3] for a precise definition). If we evaluate ψ using the naive fixed-point iteration, then we have to recompute the inner greatest fixed point for each iteration of the outer least fixed point, due to the dependency on R; in general, the number of required evaluations of the innermost formula is exponential in the alternation depth.

In order to define a semiring semantics, we use the same notation of K-interpretations as for FO, but additionally assume K to be fully continuous. We then use the natural order of K to define *least* and *greatest* fixed points. Because of full continuity, these fixed points are guaranteed to exist (by Proposition 3.2.2) and are compatible with the semiring operations. To define the semantics of the new fixed-point operators, we generalise the update operators F_φ to semiring semantics. If R has arity m, then its interpretations in a semiring K (over universe A) are functions $g\colon A^m \to K$, which form the function semiring K^{A^m} under pointwise operations (cf. Section 3.1.3). In particular, we have $g \le g'$ if $g(\boldsymbol{a}) \le g'(\boldsymbol{a})$ for all $\boldsymbol{a} \in A^m$. Given a K-interpretation $\pi\colon \mathrm{Lit}_A(\tau) \to K$ and $g\colon A^m \to K$, we denote by $\pi_{[R\mapsto g]}$ the K-interpretation of signature $\tau \cup \{R\}$ that extends π by setting $\pi_{[R\mapsto g]}(R\boldsymbol{a}) \coloneqq g(\boldsymbol{a})$ for all $\boldsymbol{a} \in A^m$. (Notice that R appears only positively in φ, so the interpretation of negative literals $\neg R\boldsymbol{a}$ is irrelevant.) The formula $\varphi(R, \boldsymbol{x})$ then induces a monotone update operator $F^\varphi_\pi\colon K^{A^m} \to K^{A^m}$ on the function semiring, that maps $g\colon A^m \to K$ to the function

$$F^\varphi_\pi(g)\colon\ \boldsymbol{a} \mapsto \pi_{[R\mapsto g]}[\![\varphi(R, \boldsymbol{a})]\!].$$

If K is fully continuous, then so is K^{A^m}, and hence F^φ_π has well-defined least and greatest fixed points $\mathbf{lfp}(F^\varphi_\pi)$ and $\mathbf{gfp}(F^\varphi_\pi)$, which are again functions $A^m \to K$. We then define semiring semantics for LFP by extending the FO-semantics with the new rules

$$\pi\Big[\!\!\Big[[\mathbf{lfp}\, R\boldsymbol{x}.\ \varphi(R, \boldsymbol{x})](\boldsymbol{a})\Big]\!\!\Big] \coloneqq \mathbf{lfp}(F^\varphi_\pi)(\boldsymbol{a}),$$

$$\pi\Big[\!\!\Big[[\mathbf{gfp}\, R\boldsymbol{x}.\ \varphi(R, \boldsymbol{x})](\boldsymbol{a})\Big]\!\!\Big] \coloneqq \mathbf{gfp}(F^\varphi_\pi)(\boldsymbol{a}).$$

The monotonicity of F^φ_π deserves further attention. Since semiring addition and multiplication are both monotone, it is clear that F^φ_π is monotone for FO-formulae $\varphi(R, \boldsymbol{x})$. One can further show by a simple induction on the fixed-point iteration that F^φ_π is also monotone for LFP-formulae $\varphi(R, \boldsymbol{x})$ that may include further fixed-point operators. It follows that the semantics is well-defined in any fully-continuous[1] semiring.

Definition 5.2.1. Semiring semantics of LFP(τ) in a fully-continuous semiring K is defined by $\pi[\![\cdot]\!]$ for K-interpretations π over finite universe and finite, relational signature τ, assigning to each instantiated LFP-formula $\varphi(\boldsymbol{a})$ a semiring value.

Example 5.2.2. For a simple example of a greatest fixed point, consider the following formula that asserts, in classical semantics, that there is an infinite path from v:

$$\psi(v) = [\mathbf{gfp}\, R\, x.\ \underbrace{\exists y(Exy \wedge Ry)}_{\varphi(R,x)}](v)$$

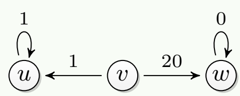

We evaluate $\psi(v)$ in the semiring interpretation π in the tropical semiring shown on the right. Writing a function $g\colon \{u, v, w\} \to \mathbb{T}$ as a tuple $(g(u), g(v), g(w))$, the induced operator

[1]For a well-defined semantics, *fully chain-complete* semirings would suffice, but continuity seems indispensable to prove any interesting results.

$F^{\varphi}_{\pi} \colon g \mapsto g'$ with $g'(z) = \pi_{[R \mapsto g]}[\![\varphi(R, z)]\!]$ for $z \in \{u, v, w\}$ can be written as

$$F^{\varphi}_{\pi} \colon \begin{pmatrix} a \\ b \\ c \end{pmatrix} \mapsto \begin{pmatrix} 1 +_{\mathbb{R}} a \\ \min(1 +_{\mathbb{R}} a,\ 20 +_{\mathbb{R}} c) \\ 0 +_{\mathbb{R}} c \end{pmatrix}.$$

This is exactly the introductory example of Chapter 4, with the infinite iteration

$$\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix} \xmapsto{F^{\varphi}_{\pi}} \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix} \xmapsto{F^{\varphi}_{\pi}} \begin{pmatrix} 2 \\ 2 \\ 0 \end{pmatrix} \xmapsto{F^{\varphi}_{\pi}} \cdots \xmapsto{F^{\varphi}_{\pi}} \begin{pmatrix} 20 \\ 20 \\ 0 \end{pmatrix} \xmapsto{F^{\varphi}_{\pi}} \begin{pmatrix} 21 \\ 20 \\ 0 \end{pmatrix} \xmapsto{F^{\varphi}_{\pi}} \begin{pmatrix} 22 \\ 20 \\ 0 \end{pmatrix} \xmapsto{F^{\varphi}_{\pi}} \cdots$$

with infimum and greatest fixed point $(\infty, 20, c)$, so that $\pi[\![\psi(v)]\!] = 20$. In general, evaluating $\psi(v)$ in the tropical semiring yields the *minimal cost* of an infinite path from v.

For the Boolean semiring, this semantics coincides with the classical semantics. In general, however, there is an important difference between the classical Boolean semantics and semiring semantics concerning the relationship of fixed-point logics with first-order logic. The (Boolean) evaluation of a fixed-point formula on a finite structure is computed by fixed-point iterations that terminate after a polynomial number of stages (with respect to the size of the structure). Hence, on any fixed finite universe, a fixed-point formula can be unraveled to an equivalent first-order formula. This is not the case for the provenance valuations in infinite semirings, such as the tropical semiring or formal power series. These semirings provide more information than just the truth or falsity of a statement, for instance about the number and properties of successful evaluation strategies, and this information may also be infinite.

As for first-order logic, semiring semantics of LFP is compatible with homomorphisms. However, these homomorphisms must now be fully continuous to preserve fixed points.

Proposition 5.2.3 (Fundamental property). *Let K_1, K_2 be fully chain-complete semirings with greatest elements $\top_1$, $\top_2$, respectively. Let $h \colon K_1 \to K_2$ be a fully-continuous semiring homomorphism with $h(\top_1) = \top_2$. Then for every K_1-interpretation π, the mapping $h \circ \pi$ is a K_2-interpretation and for every sentence $\varphi \in$ LFP, we have $h(\pi[\![\varphi]\!]) = (h \circ \pi)[\![\varphi]\!]$.*

Proof outline. The proof is by induction on φ. For a fixed-point formula, one shows that h maps each step of the fixed-point iteration in π to the corresponding step of the iteration in $h \circ \pi$, and thus in particular preserves the (least or greatest) fixed point (cf. [DGNT21]). □

As diagram:

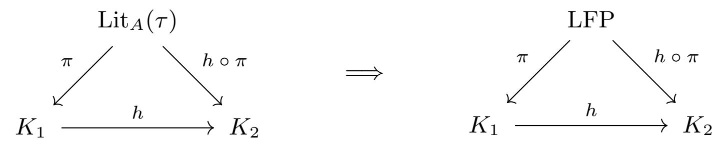

The assumption $h(\top_1) = \top_2$ is needed to ensure that h maps the starting point of the **gfp**-iteration in π to the starting point of the **gfp**-iteration in $h \circ \pi$. It can be dropped in absorptive semirings, where 1 is always the greatest element.

As a sanity check for our semantics, we want to prove that it is truth-preserving. But this is no longer the case in all positive semirings, due to greatest fixed points.

Example 5.2.4. We recall the infinite-path formula from Example 5.2.2:

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

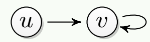

First consider the model-defining Viterbi-interpretation π defined by $\pi(Euv) = 1$ and $\pi(Evv) = 1 - \varepsilon$ for some small $\varepsilon > 0$. We then obtain an overall confidence score of $\pi[\![\varphi(u)]\!] = 0$, due to the fixed-point iteration $1, 1-\varepsilon, (1-\varepsilon)^2, \dots$. So although π induces the graph shown above (as Boolean structure $\mathfrak{A}_\pi$), it evaluates $\varphi(u)$ to 0 which we usually interpret as *false*. Hence the Viterbi semiring is *not truth-preserving*, despite being positive. From a provenance perspective, the confidence score 0 makes sense, as the loop with confidence $\pi(Evv) < 1$ has to be taken infinitely often, so the Viterbi semiring may still be useful for applications.

Consider next the semirings of formal power series $\mathbb{N}^\infty[\![X]\!]$. For the model-defining interpretation π with $\pi(Euv) = x$ and $\pi(Evv) = y$, we get $\pi[\![\varphi(u)]\!] = 0$, as result of the iteration $\top$, $y \cdot \top$, $y^2 \cdot \top$, $y^3 \cdot \top$, $\dots$ with infimum 0 at node v (here, $\top$ is the power series in which all monomials have coefficient ∞). Thus, $\mathbb{N}^\infty[\![X]\!]$ is *not truth-preserving* either.

We notice that both the Viterbi semiring and $\mathbb{N}^\infty[\![X]\!]$ permit decreasing chains of nonzero elements with zero as infimum. Recall that we have used the truth-projection homomorphism to prove that FO-semantics are truth-preserving. In this case, the decreasing chains witness that $\dagger_{\mathbb{V}}$ and $\dagger_{\mathbb{N}^\infty[\![X]\!]}$ are not fully continuous, so that we cannot apply the Fundamental property. If a semiring K is chain-positive (cf. Definition 3.2.3) so that such decreasing chains do not exist, we can use $\dagger_K$ to see that semiring semantics for LFP is truth-preserving. This applies, for instance, to $\mathbb{S}^\infty(X)$.

Corollary 5.2.5 (Truth-preserving)**.** *Let π be a model-defining K-interpretation in a fully-continuous, positive, and chain-positive semiring K. Then for every sentence $\varphi \in$ LFP,*

$$\pi[\![\varphi]\!] \neq 0 \quad \Longleftrightarrow \quad \mathfrak{A}_\pi \models \varphi.$$

5.2.2 Absorption

Semiring semantics for LFP is well-defined in all fully-continuous semirings, so absorption is not a strict requirement. However, from a provenance perspective it turns out that semiring semantics provides particularly useful information when we work with absorptive semirings (cf. Section 6.3), whereas greatest fixed points may not yield useful information in non-absorptive semirings.

Example 5.2.6. Recall the LFP-formula and the graph of Example 5.2.4:

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

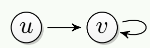

The semiring $\mathbb{N}^\infty$ is used to count proofs of FO- and posLFP-formulae, but this does not work here: While the graph only has one infinite path that we would view as a witness of $\varphi(u)$, setting $\pi(Euv) = \pi(Evv) = 1$ results in $\pi[\![\varphi(u)]\!] = \infty$, as the fixed-point iteration of φ at v is ∞, $1 \cdot \infty$, $1 \cdot \infty$, $\dots$ which stagnates immediately. This hints at a general issue: multiplication with nonzero values in $\mathbb{N}^\infty$ always increases values. The same is true for addition, so fixed-point iterations of **gfp**-formula are likely to result in ∞ and do not give meaningful provenance information (although the semantics in $\mathbb{N}^\infty$ is truth-preserving).

Moreover, the free ω-continuous semiring $\mathbb{N}^\infty[\![X]\!]$ is not the right provenance semiring for LFP. Indeed, recall that evaluating $\varphi(u)$ in $\mathbb{N}^\infty[\![X]\!]$ yields 0 (cf. Example 5.2.4), so we cannot obtain the valuation in $\mathbb{N}^\infty$ from the one in $\mathbb{N}^\infty[\![X]\!]$ by polynomial evaluation.

One argument in favour of absorption is that it increases symmetry. In the Boolean setting, fixed points are computed in the complete lattice of subsets which is inherently symmetric. For instance, a greatest fixed point of a monotone operator is the complement of the least fixed point of the dual operator (which is essential for a negation normal form). Moreover, conjunction and disjunction are symmetric in the sense that one increases values, acting as set union in the lattice of subsets, while the other is decreasing. In the semiring setting, we compute fixed points with respect to the natural order induced by addition. We assume this order to be sufficiently complete and it is clear that addition is increasing in the sense that $a + b \geq a$ for all a, b. The issue is with multiplication: the only constraint that relates addition or the natural order with multiplication is distributivity, but this alone does not suffice to ensure a symmetry similar to the Boolean setting. This is where absorption comes in. Multiplication becomes decreasing ($ab \leq a$ for all a, b) and 1 becomes the greatest element, symmetric to addition and the least element 0. This symmetry helps, for instance, to avoid problems of increasing multiplication as in $\mathbb{N}^\infty$. Fixed-point theory often relies on symmetry and it is thus no surprise that more symmetry leads to more useful provenance information.

A further motivation for absorptive semirings is that they give information about reduced proofs of a formula. The property $a + ab = a$ implies, for example, that a proof containing two literals mapped to a and b, thus having the value ab, is absorbed by a proof only using one literal, with provenance value a. This has the benefit that, unlike formal power series $\mathbb{N}^\infty[\![X]\!]$, provenance information is always finitely representable (in $\mathbb{S}^\infty(X)$).

Example 5.2.7. We recall the setting of Example 5.2.4 and first consider the model-defining $\mathbb{S}^\infty(X)$-interpretation tracking the two edges labelled x and y (left graph):

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

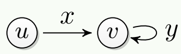

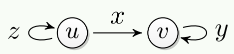

We obtain $\pi[\![\varphi(u)]\!] = xy^\infty$ corresponding to the infinite path $uvvv\ldots$. Notice that, unlike for $\mathbb{N}^\infty[\![X]\!]$, we can obtain the confidence score in the Viterbi semiring by polynomial evaluation: For $h(x) = 1$ and $h(y) = 1 - \varepsilon$, we get $h(xy^\infty) = 1 \cdot (1 - \varepsilon)^\infty = 0$ as expected.

Now set $\pi(Euu) = z$ to obtain the interpretation on the right. There are now infinitely many infinite paths from u to v. However, we obtain only finitely many monomials due to absorption: $\pi[\![\varphi(u)]\!] = xy^\infty + z^\infty$. These correspond to the *simplest* infinite paths, since monomials such as z^2xy^∞ (corresponding to the path $uuuvvv\ldots$) are absorbed by xy^∞.

These arguments suggest that absorptive semirings, and in particular the semirings of generalised absorptive polynomials $\mathbb{S}^\infty(X)$ and their dual-indeterminate versions $\mathbb{S}^\infty(X, \overline{X})$, are the right provenance semirings for LFP. This is further justified by observing that many provenance applications, such as computing repairs (see Section 6.3.5 for an example in LFP), can be solved already in the simple (and absorptive) semiring $\mathrm{PosBool}(X, \overline{X})$. A second justification is the Sum-of-strategies theorem [DGNT21] that we discuss in more detail in Chapter 6. Intuitively, it states that in absorptive, fully-continuous semirings, the value $\pi[\![\psi]\!]$

of an LFP-sentence ψ can be described as the sum over the values of all winning strategies in the associated model-checking game, similar to the sum over all proof trees for FO. In the following, we always evaluate LFP in absorptive semirings.

5.2.3 Fixed-Point Computation

We show how the fixed-point computation of Chapter 4 can be used to evaluate LFP-formulae. To this end, we describe a translation from formulae into polynomial equation systems. Consider a formula $\varphi(R, \boldsymbol{x}) \in \mathrm{FO}(\tau \cup \{R\})$ and let r be the arity of R. Let further π be a K-interpretation over universe A. We introduce indeterminates $\boldsymbol{X} = \{X_{\boldsymbol{a}} \mid \boldsymbol{a} \in A^r\}$ and let $g_{\boldsymbol{X}}$ be the mapping $g_{\boldsymbol{X}} \colon A^r \to K[\boldsymbol{X}]$ with $\boldsymbol{a} \mapsto X_{\boldsymbol{a}}$. Notice that we can view π also as a $K[\boldsymbol{X}]$-interpretation. Then, the polynomial equation system induced by φ in π can be elegantly defined as:

$$\big(X_{\boldsymbol{a}} = \pi_{[R \mapsto g_{\boldsymbol{X}}]}[\![\varphi(\boldsymbol{a})]\!]\big)_{\boldsymbol{a} \in A^r},$$

where $\pi_{[R \mapsto g_{\boldsymbol{X}}]}[\![\varphi(\boldsymbol{a})]\!]$ is a polynomial in $K[\boldsymbol{X}]$. Let $F \colon K^{|A^r|} \to K^{|A^r|}$ be the operator induced by the polynomial equation system. Identifying tuples in $K^{|A^r|}$ with functions $g \colon A^r \to K$, we have $F(g) = F_\pi^\varphi(g)$ by construction, and thus $\mathbf{lfp}(F) = \mathbf{lfp}(F_\pi^\varphi)$ and $\mathbf{gfp}(F) = \mathbf{gfp}(F_\pi^\varphi)$. Since we can represent φ as a polynomial equation system, it follows that $\mathbf{lfp}(F_\pi^\varphi) = (F_\pi^\varphi)^\ell(\mathbf{0})$ and $\mathbf{gfp}(F_\pi^\varphi) = (F_\pi^\varphi)^\ell\big((F_\pi^\varphi)^\ell(\mathbf{1})^\infty\big)$, where $\ell = |A^r|$.

Example 5.2.8. Recall the setting of Example 5.2.7:

$$\psi(u) = [\mathbf{gfp}\, R\, x.\, \underbrace{\exists y(Exy \wedge Ry)}_{\varphi(R,x)}](u)$$

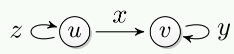

We construct the induced polynomial equation system (notice that this is just for illustration, this is not necessary to evaluate ψ). We use the indeterminates $\boldsymbol{X} = \{X_u, X_v\}$. Evaluating $\varphi(u)$, $\varphi(v)$ in $\pi_{[R \mapsto g_{\boldsymbol{X}}]}$ yields the equations:

$$\begin{aligned} X_u &= zX_u + xX_v, \\ X_v &= yX_v. \end{aligned}$$

We can use Theorem 4.3.1 to evaluate ψ, using $\ell = 2$ iterations:

$$\begin{pmatrix}1\\1\end{pmatrix} \overset{F_\pi^\varphi}{\longmapsto} \begin{pmatrix}z+x\\y\end{pmatrix} \overset{F_\pi^\varphi}{\longmapsto} \begin{pmatrix}z^2+zx+xy\\y^2\end{pmatrix} \overset{\infty}{\longmapsto} \begin{pmatrix}z^\infty + x^\infty y^\infty\\y^\infty\end{pmatrix} \overset{F_\pi^\varphi}{\longmapsto} \begin{pmatrix}z^\infty + xy^\infty\\y^\infty\end{pmatrix} \circlearrowleft F_\pi^\varphi$$

We obtain $\pi[\![\psi(u)]\!] = z^\infty + xy^\infty$ in the first component as expected.

This covers the case of a single fixed point. If $\varphi(R, \boldsymbol{x})$ is an LFP-formula containing further fixed-point formulae, the reasoning in Section 4.5 shows that these inner fixed points can be described by (generalised) polynomials (by evaluating them in $\mathbb{S}^\infty(X)$) and hence also φ gives rise to a generalised polynomial equation system. This means that we can always apply the closed-form solution of Chapter 4 to evaluate LFP-formulae. Notice that computing $F_\pi^\varphi(g)$ may require to first (re-)compute inner fixed points, just as it does in the naive evaluation of LFP in classical Boolean semantics.

Corollary 5.2.9. *Let K be an absorptive, fully-continuous semiring, and let π be a K-interpretation over universe A. Given $\varphi(R, \boldsymbol{x}) \in \mathrm{LFP}(\tau \cup \{R\})$, we have:*

$$\pi[\![[\mathbf{lfp}\, R\,\boldsymbol{x}.\, \varphi](\boldsymbol{a})]\!] = (F_\pi^\varphi)^\ell(\mathbf{0})\,(\boldsymbol{a}),$$

$$\pi[\![[\mathbf{gfp}\, R\,\boldsymbol{x}.\, \varphi](\boldsymbol{a})]\!] = (F_\pi^\varphi)^\ell\big(\,(F_\pi^\varphi)^\ell(\mathbf{1})^\infty\,\big)\,(\boldsymbol{a}),$$

where $\ell = |A^r|$ and r is the arity of R.

5.2.4 Closure Ordinals

Although we consider finite structures, the fixed-point iterations induced by LFP-formulae need not be finite if the semiring is infinite. We show that the iterations are still limited for absorptive semirings, in the sense that the closure ordinal of an operator F_π^φ induced by a LFP-formula is at most ω. That is, $\mathbf{lfp}(F_\pi^\varphi) = \bigsqcup_{n<\omega}(F_\pi^\varphi)^n(\mathbf{0})$ and $\mathbf{gfp}(F_\pi^\varphi) = \bigsqcap_{n<\omega}(F_\pi^\varphi)^n(\mathbf{1})$.

One way to prove this would be to show that the operators F_π^φ are fully-continuous, so that Kleene's fixed-point theorem (and its dual version) can be applied. But this is, in general, not the case for nested fixed points.

Example 5.2.10. Recall the Łukasiwicz semiring $\mathbb{L} = ([0,1], \max, \star, 0, 1)$ with $a \star b = \max(0, a+b-1)$. We consider the (artificial) setting of evaluating the following LFP-formula on the $\mathbb{L}$-interpretation with singleton universe $A = \{a\}$ over signature $\tau = \emptyset$:

$$\psi(z) = [\mathbf{lfp}\, P\, y.\, \underbrace{[\mathbf{gfp}\, R\, x.\, Rx \wedge Px](y)}_{\varphi(P,y)}](z).$$

We claim that the operator F_π^φ for the outer **lfp**-iteration is not fully-continuous. Since the universe only contains a single element, we can simplify notation and view $F_\pi^\varphi \colon K \to K$ as a function on K. Now consider the chain $(x_n)_{n<\omega}$ with $x_n = \frac{n}{n+1}$ and supremum 1. We thus have $F_\pi^\varphi(\bigsqcup_n x_n) = F_\pi^\varphi(1) = \mathbf{gfp}(a \mapsto a \star 1) = 1$. On the other hand,

$$\bigsqcup_{n<\omega} F_\pi^\varphi(x_n) = \bigsqcup_{n<\omega} \mathbf{gfp}(a \mapsto a \star x_n) = \bigsqcup_{n<\omega} 0 = 0,$$

by definition of $\star$ and since $x_n < 1$. Hence F_π^φ does not commute with suprema of chains.

We remark that the chain $(x_n)_{n<\omega}$ used in the example is artificial in the sense that it cannot be generated by LFP-formulae over $\mathbb{L}$. We can avoid such problems by working in $\mathbb{S}^\infty(X)$, where no infinite ascending chains exist. By using the most general $\mathbb{S}^\infty(X)$-interpretation, that maps each literal L to a unique indeterminate $x_L \in X$, one can show that operators F_π^φ are indeed fully continuous in $\mathbb{S}^\infty(X)$, and hence have closure ordinal ω by Kleene's fixed-point theorem [DGNT21]. Since fully-continuous homomorphisms preserve the steps of the fixed-point iteration, this generalises to all absorptive, fully-continuous semirings:

Proposition 5.2.11. *Given a K-interpretation π into an absorptive, fully continuous semiring, all fixed-point iterations for $\mathbf{lfp}(F_\pi^\varphi)$ and $\mathbf{gfp}(F_\pi^\varphi)$ have closure ordinal at most ω.*

5.2.5 Outlook

We have seen that semiring semantics can be defined also for full fixed-point logic including greatest fixed points. We get a beautiful theory for absorptive, fully-continuous semirings: the Fundamental property together with the universality of $\mathbb{S}^\infty(X)$ means that we have a suitable provenance semiring, where provenance information is finitely represented and can be computed in a simple way. The Sum-of-strategies theorem and its applications (see Chapter 6 for both) further show that this information is meaningful.

Modal μ-calculus. The semantics for LFP also induces a semiring semantics for the modal μ-calculus L_μ, as there is a direct translation of L_μ into LFP that we can use to define such a semantics (see [GKL+07, Chapter 3]). A small obstacle is the translation of the modal operator $\Box$, which introduces an implication that first has to be brought into negation normal form. In fact, there are several reasonable ways to define the semantics of $\Box\varphi$, and a detailed analysis of semiring semantics for modal logics is a separate topic (see also [DG21] for a semantics of $\Box\varphi$ that is not based on the LFP-translation). However, for absorptive semirings, the translation of $\psi = \Box\varphi$ into the negation normal form $\psi^*(v) = \forall w(\neg Evw \vee (Evw \wedge \varphi^*(w)))$ (assuming that we evaluate ψ on a transition system with edge relation E) seems to be the most appropriate choice. Using this translation to define a semiring semantics for L_μ, the Fundamental property and the fixed-point computation are inherited from LFP.

This semiring semantics may be of interest also for applications, as the logics CTL and CTL*, which are relevant for formal verification, can both be embedded into L_μ (one should note, however, that CTL embeds directly into the alternation-free fragment, while the embedding of CTL* is more involved and leads to a significant blowup in the size of the formula, see e.g. [BS07] for background on L_μ).

Future research. There is a rich theory of LFP and related fixed-point logics in classical semantics, and it may be worthwhile to study this theory in the context of semiring semantics. We point out two questions that appear interesting. The first concerns the alternation hierarchy of LFP. In classical semantics, it is well-known that LFP collapses to posLFP on finite structures, due to the Stage comparison theorem [GKL+07, Chapter 3]. While it is not clear whether some form of the Stage comparison theorem also holds for semiring semantics, and this may be an interesting question on its own, a simple argument shows that LFP does *not* collapse to posLFP over absorptive, fully-continuous semirings.

Example 5.2.12. We again consider $\varphi(v) = [\mathbf{gfp}\, R\, x.\ \exists y(Exy \wedge Ry)](v)$. We claim that there is no posLFP-formula $\varphi'(v)$ such that $\varphi \equiv_K \varphi'$ for all absorptive, fully-continuous semirings K. To see this, consider a $\mathbb{S}^\infty(x)$-interpretation π consisting of a single node v with a self-loop labeled x. Since FO-formulae correspond to polynomials and least fixed points can be computed in a polynomial number of steps by Corollary 5.2.9, it follows that the value $\pi[\![\varphi'(v)]\!]$ of any posLFP-formula φ' is a polynomial expression in x, without the exponent ∞. On the other hand, $\pi[\![\varphi(v)]\!] = x^\infty$, so φ and φ' cannot be equivalent.

The situation in semiring semantics is thus different from classical semantics, but this simple argument only seems useful for separating LFP and posLFP. New techniques may be required to look at higher levels of the alternation hierarchy.

Open Question. *Is the alternation hierarchy of LFP over finite universes under semiring semantics (in absorptive, fully-continuous semirings) strict?*

A related topic are inflationary fixed points and the corresponding logic IFP. In classical semantics, LFP is known to be equally expressive as IFP, again due to the Stage comparison theorem (cf. [GKL+07, Chapter 3]). Inflationary fixed points can easily be defined in semiring semantics by using semiring addition (which is equal to supremum in absorptive semirings), and it may be an interesting direction for future work to see if this can be simulated by LFP-formulae as in the classical case, or if it leads to new expressive power.

5.3 Infinitary Logic

Semiring semantics has traditionally been restricted to interpretations over finite universes, but this limitation can be lifted with the infinite operations defined in Section 3.4. Besides infinite universes, we can also use these operations to define a semiring semantics for infinitary logic $L_{\infty\omega}$ which extends first-order logic by infinite conjunctions and disjunctions. In this section, we extend the semirings semantics of FO to infinitary logic. Our main contribution is the embedding of LFP (in absorptive, fully-continuous semirings over finite structures) into infinitary logic $L^{\omega}_{\omega_1\omega}$ with finitely many variables, by means of the fixed-point computation in Theorem 4.3.1.

5.3.1 Semantics

As usual, let τ be a finite relational signature. The syntax of infinitary logic $L_{\infty\omega}(\tau)$ extends FO-syntax by the following rule: for every set Φ of $L_{\infty\omega}(\tau)$-formulae, $\bigvee\Phi$ and $\bigwedge\Phi$ are $L_{\infty\omega}(\tau)$-formulae as well. For the purpose of semiring semantics, it makes sense to slightly modify the syntax. Instead of sets Φ, we consider families $(\varphi_i)_{i\in I}$ of $L_{\infty\omega}(\tau)$-formulae and then permit the expressions $\bigvee_{i\in I}\varphi_i$ and $\bigwedge_{i\in I}\varphi_i$. This accounts for the fact that addition and multiplication need not be idempotent. For a more fine-grained analysis, $L_{\kappa\omega}(\tau)$ permits $\bigvee_{i\in I}\varphi_i$ and $\bigwedge_{i\in I}\varphi_i$ only for families of size $|I|<\kappa$. In our case, the logic $L_{\omega_1\omega}(\tau)$ with countable conjunctions and disjunctions will be most relevant. We further write $L^k_{\omega_1\omega}(\tau)$ for its k-variable fragment, and $L^{\omega}_{\omega_1\omega}(\tau)=\bigcup_{k<\omega}L^k_{\omega_1\omega}(\tau)$ for the finite-variable fragment. We refer to [EF95] for more background on infinitary logic. In the following, we omit τ when it is clear from the context.

Semiring semantics of FO extends in a straightforward way to $L_{\infty\omega}$. Let π be a K-interpretation (over finite universe A and signature τ). We assume that K is an infinitary semiring with operations $\widehat{\sum}_{i\in I} a_i$ and $\widehat{\prod}_{i\in I} a_i$. We lift π to $L_{\infty\omega}$-formulae as for FO, with the additional rules:

$$\pi[\![\bigvee_{i\in I}\varphi_i(\boldsymbol{a})]\!]=\widehat{\sum_{i\in I}}\pi[\![\varphi_i(\boldsymbol{a})]\!],\qquad \pi[\![\bigwedge_{i\in I}\varphi_i(\boldsymbol{a})]\!]=\widehat{\prod_{i\in I}}\pi[\![\varphi_i(\boldsymbol{a})]\!].$$

Definition 5.3.1. Semiring semantics of $L_{\infty\omega}(\tau)$ in an infinitary semiring K is defined by $\pi[\![\cdot]\!]$ for K-interpretations π over finite universe and finite, relational signature τ, assigning to each instantiated $L_{\infty\omega}$-formula $\varphi(\boldsymbol{a})$ a semiring value.

Remark 5.3.2. In classical semantics, $L_{\omega_1\omega}$ can axiomatise any class of finite structures (up to isomorphism) with a single conjunction $\bigwedge \Phi$ for $\Phi \subseteq \mathrm{FO}$ (see [EF95, Chapter 3]), relying on characteristic sentences for finite structures. This construction does not work for semiring semantics, since finite K-interpretations cannot be described (up to isomorphism) by (sets of) first-order sentences in many semirings (see [GM21] and Section 7.1), and it remains yet to be studied to what extent classes of semiring interpretations can be axiomatised in infinitary logic under semiring semantics.

Remark 5.3.3. We note that the semantics is well-defined in all lattice semirings (which are not necessarily complete lattices). The reason is that we consider finite universes. Indeed, let K be a (possibly infinite) lattice semiring and π a K-interpretation over finite universe A. Let $K_0 \subseteq K$ be the lattice subsemiring that results from closing the subset $\{\pi(L) \mid L \in \mathrm{Lit}_A(\tau)\}$ under suprema and infima. Since this subset is finite, so is K_0. It follows by a straightforward induction that $\pi[\![\varphi(\boldsymbol{a})]\!] \in K_0$ for all $\varphi \in L_{\infty\omega}$ and $\boldsymbol{a} \subseteq A$, so the evaluation always takes place in the finite lattice K_0 and infinitary operations are just finite suprema and infima.

The properties of this new semantics of course depend on the definition of infinitary operations. We are mostly interested in absorptive, fully-continuous semirings, which we always extend to infinitary semirings by Theorem 3.4.2. As a sanity check, we note that the semantics is truth-preserving in these semirings if they are positive and chain-positive.

Proposition 5.3.4 (Truth-preserving)**.** *Let K be an absorptive, fully-continuous semiring. If K is positive and chain-positive, then for every model-defining K-interpretation π and $L_{\infty\omega}$-sentence ψ,*

$$\pi[\![\psi]\!] \neq 0 \quad \Longleftrightarrow \quad \mathfrak{A}_\pi \models \varphi.$$

Proof. The truth-projection $\dagger_K$ is fully continuous in chain-positive semirings, and thus commutes with infinitary sums and products (by Proposition 3.4.3). ☐

5.3.2 Embedding of LFP

We show that LFP can be embedded into infinitary logic $L^{\omega}_{\omega_1\omega}$, under semiring semantics in absorptive, fully-continuous semirings and over finite universes.

Theorem 5.3.5. *For every LFP-formula $\varphi(\boldsymbol{x})$, there is a $L^{\omega}_{\omega_1\omega}$-formula $\psi(\boldsymbol{x})$ such that $\varphi \equiv_K \psi$ for all absorptive, fully-continuous semirings K (over finite universes).*

Let K be an absorptive, fully-continuous semiring and π a K-interpretation. Consider a formula $[\mathbf{lfp}\, R\boldsymbol{x}.\, \varphi(R,\boldsymbol{x})](\boldsymbol{a})$ or $[\mathbf{gfp}\, R\boldsymbol{x}.\, \varphi(R,\boldsymbol{x})](\boldsymbol{a})$. We first follow the standard approach and inductively define formulae $\varphi_n^{\bot}(\boldsymbol{x})$ and $\varphi_n^{\top}(\boldsymbol{x})$ that capture the (least and greatest) fixed-point iterations:

- $\varphi_0^{\bot}(\boldsymbol{x}) \coloneqq (x \neq x)$ and $\varphi_0^{\top}(\boldsymbol{x}) \coloneqq (x = x)$, for some $x \in \boldsymbol{x}$,
- $\varphi_{n+1}^{\circ}(\boldsymbol{x}) \coloneqq \varphi(\varphi_n^{\circ}(\boldsymbol{x}), \boldsymbol{x})$, for $\circ \in \{\bot, \top\}$.

Here, $\varphi(\varphi_n^{\circ}(\boldsymbol{x}), \boldsymbol{x})$ is the formula that results from $\varphi(R,\boldsymbol{x})$ by replacing every literal of the form $R\boldsymbol{y}$ with $\varphi_n^{\circ}(\boldsymbol{y})$. The resulting formulae only use the variables $\boldsymbol{x}$ appearing in $\varphi(R,\boldsymbol{x})$.

It is clear by induction that

$$\pi[\![\varphi_n^\perp(\boldsymbol{a})]\!] = (F_\pi^\varphi)^n(\mathbf{0})\,(\boldsymbol{a}) \quad \text{and} \quad \pi[\![\varphi_n^\top(\boldsymbol{a})]\!] = (F_\pi^\varphi)^n(\mathbf{1})\,(\boldsymbol{a}). \tag{5.1}$$

Since LFP-formulae have closure ordinal ω in K (by Proposition 5.2.11) and infinite summation is the same as supremum, we can easily express the **lfp**-formula by the conjunction $\bigvee_{n<\omega} \varphi_n^\perp(\boldsymbol{a})$. The **gfp**-formula requires more work, since infinite conjunction is interpreted as multiplication, and not as infimum. Instead, we recall the closed-form solution $\mathbf{gfp}(F) = F^\ell(F^\ell(\mathbf{1})^{\boldsymbol{\infty}})$ of Theorem 4.3.1. The main idea is that we can express the part $F^\ell(\mathbf{1})^{\boldsymbol{\infty}}$ as an infinite conjunction. The difference between infimum and multiplication does not matter here, as the infinitary power means that we are multiplying the result infinitely often anyway. The outer part $F^\ell(\dots)$ can then again by expressed by a disjunction, which does not cause any problems. This leads to the following translation.

Lemma 5.3.6. *Let K be an absorptive, fully-continuous semiring. We have the following equivalences (over finite universes):*

(1) $[\mathbf{lfp}\, R\,\boldsymbol{x}.\, \varphi(R, \boldsymbol{x})](\boldsymbol{a}) \equiv_K \bigvee_{n<\omega} \varphi_n^\perp(\boldsymbol{a}),$

(2) $[\mathbf{gfp}\, R\,\boldsymbol{x}.\, \varphi(R, \boldsymbol{x})](\boldsymbol{a}) \equiv_K \bigvee_{n<\omega} \vartheta_n(\boldsymbol{a}),$

where the formulae $\vartheta_n(\boldsymbol{x})$ are defined inductively by

$$\vartheta_0(\boldsymbol{x}) \coloneqq \bigwedge_{n<\omega} \varphi_n^\top(\boldsymbol{x}), \qquad \vartheta_{n+1}(\boldsymbol{x}) \coloneqq \varphi(\vartheta_n(\boldsymbol{x}), \boldsymbol{x}).$$

Theorem 5.3.5 follows from this lemma by a straightforward induction on φ. The proof of the lemma requires some preparation.

Lemma 5.3.7. *Let $(a_i)_{i<\omega}$ be a decreasing ω-chain in an absorptive, fully-continuous semiring. Then,*

$$\Big(\widehat{\prod_{i<\omega}} a_i\Big)^\infty = \Big(\bigsqcap_{i<\omega} a_i\Big)^\infty.$$

Proof. Recall that the infinitary power commutes with infinite products (Lemma 3.4.4) and infima of chains (Lemma 3.3.3). By applying the definition of infinite products, we thus get:

$$\Big(\widehat{\prod_{i\in I}} a_i\Big)^\infty = \widehat{\prod_{i\in I}} a_i^\infty = \bigsqcap_{F\subset_{\text{fin}}\omega} \prod_{i\in F} a_i^\infty \overset{(*)}{=} \bigsqcap_{F\subset_{\text{fin}}\omega} a_{\max(F)}^\infty = \bigsqcap_{i<\omega} a_i^\infty = \Big(\bigsqcap_{i<\omega} a_i\Big)^\infty.$$

We show both directions of $(*)$. First, $\prod_{i\in F} a_i^\infty \le a_{\max(F)}^\infty$ by absorption. On the other hand, $(a_i)_{i<\omega}$ is a decreasing chain, and hence $\prod_{i\in F} a_i^\infty \ge \prod_{i\in F} a_{\max(F)}^\infty = a_{\max(F)}^\infty$. □

We next show that we can express $(F_\pi^\varphi)^\ell(\mathbf{1})^\infty$ as infinite product. This corresponds to the infinite conjunction $\bigwedge_{n<\omega} \varphi_n^\top(\boldsymbol{x})$ in $\vartheta_0(\boldsymbol{x})$.

Lemma 5.3.8. *Let π be a K-interpretation over finite universe A in an absorptive, fully-continuous semiring. Let $\varphi(R, \boldsymbol{x})$ be a* LFP$(\tau \cup \{R\})$*-formula and $\ell = |A^r|$, where r is the arity of R. Let $F = F_\pi^\varphi : K^\ell \to K^\ell$ be the operator induced by φ. Then,*

$$\widehat{\prod_{n<\omega}} F^n(\mathbf{1}) = F^\ell(\mathbf{1})^{\boldsymbol{\infty}}.$$

Proof. The direction "$\leq$" follows from monotonicity of F and absorption:

$$\widehat{\prod_{n<\omega}} F^n(\mathbf{1}) \leq \widehat{\prod_{\ell \leq n<\omega}} F^n(\mathbf{1}) \leq \widehat{\prod_{\ell \leq n<\omega}} F^\ell(\mathbf{1}) = F^\ell(\mathbf{1})^\infty.$$

For the other direction, recall that F corresponds to a polynomial equation system over a set $\boldsymbol{X}$ of ℓ many indeterminates (cf. Section 5.2.3), even in the case of nested fixed points. We borrow the notation for derivation trees of Chapter 4 for the proof and use indeterminates $X \in \boldsymbol{X}$ to index tuples in K^ℓ. We first prove the statement for the semiring $K = \mathbb{S}^\infty(\boldsymbol{A})$, where $\boldsymbol{A}$ is an arbitrary finite set of indeterminates. For each $X \in \boldsymbol{X}$, we have:

$$\begin{aligned}
\widehat{\prod_{n<\omega}} F^n(\mathbf{1})_X &\overset{(4.2.9)}{=} \widehat{\prod_{n<\omega}} \widehat{\sum_{T \in \mathcal{T}(X)}} \mathsf{yd}(T\|^{\mathbf{1}}_n) \overset{(\dagger)}{\geq} \widehat{\sum_{T \in \mathcal{T}(X)}} \widehat{\prod_{n<\omega}} \mathsf{yd}(T\|^{\mathbf{1}}_n) \\
&\geq \widehat{\sum_{T \in \mathcal{T}(X)}} \Big(\widehat{\prod_{n<\omega}} \mathsf{yd}(T\|^{\mathbf{1}}_n)\Big)^\infty \overset{(5.3.7)}{=} \widehat{\sum_{T \in \mathcal{T}(X)}} \Big(\prod_{n<\omega} \mathsf{yd}(T\|^{\mathbf{1}}_n)\Big)^\infty \\
&= \widehat{\sum_{T \in \mathcal{T}(X)}} \mathsf{yd}(T)^\infty \overset{(4.3.4)}{=} \widehat{\sum_{T \in \mathcal{T}(X)}} \mathsf{yd}(T\|^{\mathbf{1}}_\ell)^\infty \\
&\overset{(3.5.6)}{=} \Big(\widehat{\sum_{T \in \mathcal{T}(X)}} \mathsf{yd}(T\|^{\mathbf{1}}_\ell)\Big)^\infty \overset{(4.2.9)}{=} F^\ell(\mathbf{1})^{\boldsymbol{\infty}}_X.
\end{aligned}$$

The proof is based on Lemma 4.2.9 (Tree iteration) and Corollary 4.3.4, and the fact that infinitary power commutes with suprema in $\mathbb{S}^\infty(X)$ (Lemma 3.5.6). Step ($\dagger$) follows from monotonicity of infinite products, since $\widehat{\sum}_{T \in \mathcal{T}(X)} \mathsf{yd}(T\|^{\mathbf{1}}_n) \geq \mathsf{yd}(T\|^{\mathbf{1}}_n)$ for each tree T. Altogether, this proves equality in $\mathbb{S}^\infty(\boldsymbol{A})$.

We generalise this to arbitrary K-interpretations π as follows. Let $\boldsymbol{A} = \{A_L \mid L \in \mathrm{Lit}_A(\tau)\}$. Let $\pi^\star$ be the most general $\mathbb{S}^\infty(\boldsymbol{A})$-interpretation, with $\pi^\star(L) = A_L$, and let h be the homomorphism induced by the assignment $h(A_L) = \pi(L)$. Then $\pi = h \circ \pi^\star$ and $F_\pi = h \circ F_{\pi^\star}$ (where h is applied in each component). Since h commutes with infinite products and the infinitary power, we obtain:

$$\begin{aligned}
\widehat{\prod_{n<\omega}} F^n_\pi(\mathbf{1}) = \widehat{\prod_{n<\omega}} h(F^n_{\pi^\star}(\mathbf{1})) &= h\Big(\widehat{\prod_{n<\omega}} F^n_{\pi^\star}(\mathbf{1})\Big) \\
&= h\Big(F^\ell_{\pi^\star}(\mathbf{1})^{\boldsymbol{\infty}}\Big) = h(F^n_{\pi^\star}(\mathbf{1}))^{\boldsymbol{\infty}} = F_\pi(\mathbf{1})^{\boldsymbol{\infty}}.
\end{aligned}$$

□

We note a simple corollary to the above proof that we need to prove Theorem 5.3.5.

Corollary 5.3.9. *In the notation of the previous lemma,* $F(F^\ell(\mathbf{1})^\infty) \geq F^\ell(\mathbf{1})^\infty$.

Proof. Let $\boldsymbol{b} = F^\ell(\mathbf{1})^\infty$. We have shown $\boldsymbol{b}_X = \widehat{\sum}_{T \in \mathcal{T}(X)} \mathsf{yd}(T)^\infty$ as part of the above proof. We further know from Chapter 4 that $\mathbf{gfp}(F)_X = F^\ell(\boldsymbol{b})_X = \widehat{\sum}_{T \in \mathcal{T}(X)} \mathsf{yd}(T)$. Hence $F^\ell(\boldsymbol{b}) \geq \boldsymbol{b}$ and the claim follows by monotonicity of F. □

With the main step in Lemma 5.3.8 in place, we are now ready to prove the equivalence in Lemma 5.3.6, which then implies Theorem 5.3.5.

Proof of Lemma 5.3.6. We have already shown the equivalence (1) for least fixed points, so it remains to prove claim (2):

$$[\mathbf{gfp}\, R\boldsymbol{x}.\, \varphi(R, \boldsymbol{x})](\boldsymbol{a}) \equiv_K \bigvee_{n<\omega} \vartheta_n(\boldsymbol{a}).$$

Let π be a K-interpretation over finite universe A and let $\ell = |A^{\mathrm{arity}(R)}|$. By Lemma 5.3.8 and Eq. (5.1), we know that

$$\pi[\![\vartheta_0(\boldsymbol{a})]\!] = (F_\pi^\varphi)^\ell(\mathbf{1})^\infty\,(\boldsymbol{a}).$$

It then follows by a simple induction that

$$\pi[\![\vartheta_n(\boldsymbol{a})]\!] = (F_\pi^\varphi)^n\big((F_\pi^\varphi)^\ell(\mathbf{1})^\infty\big)\,(\boldsymbol{a}).$$

These values form an increasing chain by Corollary 5.3.9. We have $\pi[\![\vartheta_\ell(\boldsymbol{a})]\!] = \mathbf{gfp}(F_\pi^\varphi)(\boldsymbol{a})$ by Theorem 4.3.1 and hence $\pi[\![\vartheta_n(\boldsymbol{a})]\!] = \pi[\![\vartheta_\ell(\boldsymbol{a})]\!]$ for all $n \geq \ell$. This implies the desired equality:

$$\pi[\![[\mathbf{gfp}\, R\boldsymbol{x}.\, \varphi(R, \boldsymbol{x})](\boldsymbol{a})]\!] = \pi[\![\bigvee_{n<\omega} \vartheta_n(\boldsymbol{a})]\!].$$ □

5.4 Related and Future Work

We conclude this chapter with a brief overview on the development of semiring semantics and possible directions for future research.

Database provenance. Semiring semantics has originally been introduced for positive relational algebra and datalog [GKT07] to provide a unified framework for several forms of provenance analysis for database queries, in particular the *why-provenance* introduced in [BKT01]. In the following years, this framework has received a lot of attention, with extensions to XML, aggregate queries, graph databases, and probabilistic databases, to name a few, and also practical implementations. We refer to [CCT09] for an early account of different forms of database provenance, and to [Gla21, GT17b, Sen17] for more recent surveys on semiring provenance. Despite its success, we note that semiring semantics is not the only reasonable approach to provenance analysis. For instance, a recent systematic study [BBPT22] considers several alternative definitions for provenance of datalog.

We point out some contributions to this field that are relevant in our context. One issue with semiring semantics in provenance semirings, such as $\mathbb{N}[X]$ or $\mathbb{S}^\infty(X)$, is that the resulting polynomials can become exponentially large (in the number of tracked atoms of the input database). This has been addressed for datalog in [DMRT14] by introducing an intermediate representation of the provenance polynomials in the form of a circuit. These circuits are constructed by transforming each datalog rule into a circuit gadget that represents the evaluation of this rule in a semiring, with the semiring operations as gates (see below for an example). The computation of the least fixed point is then encoded by connecting these gadgets in series, simulating the fixed-point iteration. This works for absorptive semirings, since the number of iterations is then known to be bounded by the number of possible ground atoms (cf. [EKL11] and Theorem 4.3.1). (In fact, absorptive polynomials $\mathbb{S}(X)$ were introduced in [DMRT14] for this reason.) These circuits achieve a polynomial-size representation of the provenance information in absorptive semirings. A more general study of semirings for which the number of iterations is bounded has recently been put forward in [KNP+24]. These are the *p-stable* semirings that generalise the absorptive ones (which are 0-stable). Interestingly, the authors include semirings with orders that differ from the natural order. They further consider the naive evaluation, an optimised iterative strategy to evaluate datalog programs, and identify semirings for which it is applicable.

A separate issue is negation. We have presented the syntactic approach of dealing with negation through negation normal form. An alternative are the *monus semirings* that have been used in [GP10] to extend semiring provenance to positive relational algebra with set difference. These are naturally-ordered semirings in which the monus operation $a \ominus b$ is defined as the minimal c (which is assumed to exist) with $a \leq b+c$, simulating the subtraction $a-b$. The provenance semiring $\mathbb{N}[X]$ admits such a monus operation. However, it is no longer universal; a more general construction is needed to obtain the universal monus semiring.

A different approach to handle negation, and provenance in general, are the *provenance games* in [KLZ13]. Here, a notion of game provenance is defined, which is then used to define provenance for first-order logic and positive relational algebra by first translating them into games. Unlike semiring semantics, these games can naturally express negation. Interestingly, the provenance defined in this way for positive relational algebra is equivalent to semiring provenance in the most general semiring $\mathbb{N}[X]$. Based on these ideas, *provenance graphs* for first-order queries are defined in [LLG19] and utilised for a practical implementation. This provenance model is equivalent to provenance games and also captures $\mathbb{N}[X,\overline{X}]$-provenance of first-order queries.

Lastly, we point to a new development related to provenance. The Shapley value, a concept in cooperative game theory that measures the contribution of an individual player, has recently been applied to database queries [LBKS20, BKLM23], in order to measure the contribution of individual facts to a query answer. This can be seen as a quantified version of why-provenance, and it may be interesting to study connections between semiring semantics and shapley values.

Towards stratified datalog. A topic that deserves further study is semiring provenance of *stratified datalog*, which extends datalog by a limited form of negation that can express both least and greatest fixed points (but not their alternation). One way to define a semiring semantics for stratified datalog is to use semirings that can directly interpret negation, such as monus semirings or $\mathrm{PosBool}(X,\overline{X})$ (see also [DAA13] for the more general setting of logic programs under well-founded semantics, using a monus semiring similar to PosBool$(X, \overline{X})$). Here we briefly sketch an alternative definition, based on the translation into LFP. Consider the following simple example that computes the (negation of the) transitive closure:

$$\begin{aligned} Txy &:- Exy \\ Txy &:- Exz, Tzy \end{aligned} \qquad Nxy :- \neg Txy$$

We have not formally introduced (stratified) datalog in this thesis (see, e.g., [EF95] for background). Instead, we can understand its (classical) semantics by translating these rules into the following LFP-formulae:

$$\psi_T(x, y) = [\textbf{lfp}\, T\, xy.\ Exy \lor \exists z(Exz \land Tzy)](x, y),$$
$$\psi_N(x, y) = [\textbf{gfp}\, N\, xy.\ \neg Exy \land \forall z(\neg Exz \lor Nzy)](x, y),$$

where $\psi_N(x, y)$ is the negation normal form of $\neg\psi_T(x, y)$. In other words, to compute the semiring provenance of a negation $\neg Txy$, we dualise the computation of Txy. In particular, we must now compute a greatest instead of a least fixed point.

This translation induces a semiring semantics for stratified datalog that works in any absorptive, fully-continuous semiring K. What is interesting is that the circuit representation of [DMRT14] can be extended to stratified datalog due to the closed-form solution for greatest fixed points in Theorem 4.3.1. We briefly sketch the construction for our example and a K-interpretation π (i.e., an annotated graph) with node set $V = \{v, w\}$. First, the datalog rules (or the equivalent LFP-formulae) are translated into circuit gadgets for each tuple, representing their evaluation in π. For instance, here is the gadget for Tvw (on the left). The corresponding gadget for Nvw (on the right) can be obtained by dualising the circuit (i.e., by swapping the semiring operations and negating the literals).

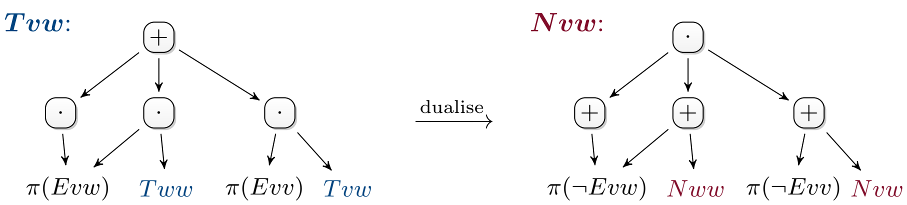

The gadget for Tvw corresponds to a single evaluation of the inner formula $Exy \lor \exists z(Exz \land Tzy)$, and thus to a single step of the fixed-point iteration. Consequently, gadgets are stacked in series to obtain the entire iteration. Since we are computing a binary relation on $n = 2$ nodes, we need $n^2 = 4$ iterations for the least fixed point (cf. [DMRT14]). The new observation is that the greatest fixed point in ψ_N can also be expressed as a circuit by Theorem 4.3.1, using two rounds of n^2 iterations and an application of the infinitary power in between. The following schematic drawing illustrates the construction (for the sake of clarity, most edges between gadgets have been omitted).

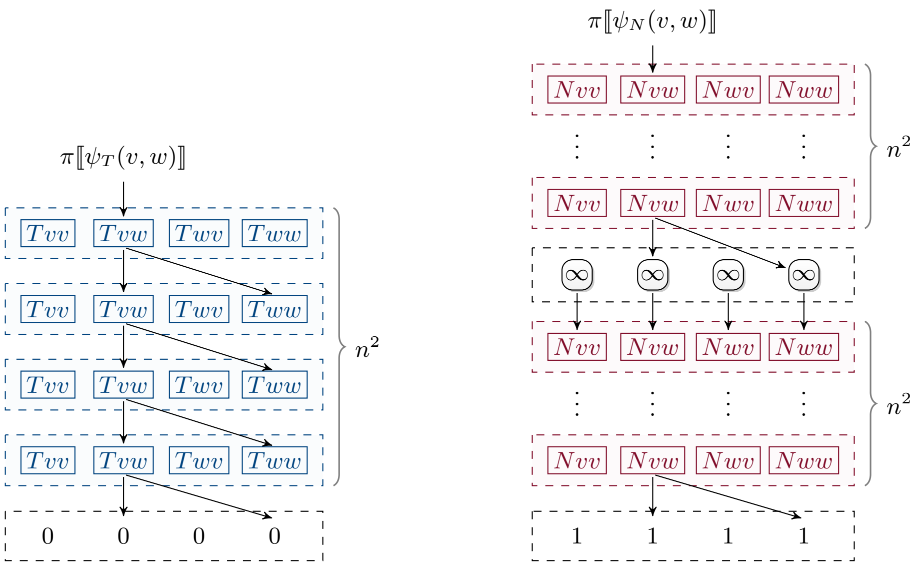

These circuits still have polynomial size (as in [DMRT14]), and may thus be of interest for provenance applications. We note, however, that we need additional unary gates for the infinitary power, and these have to be supported by applications that modify these circuits or read off provenance information.

Beyond databases. Semiring semantics has also been considered for various logics beyond database queries. Here we have presented the semantics for FO and LFP. We have also defined a semiring semantics for infinitary logic, with the goal of extending the 0-1 laws from first-order logic to LFP (cf. Chapter 7). A detailed study of infinitary logic and, perhaps, also second-order logic under semiring semantics remains to be done.

Modal logic and the guarded fragments of modal and first-order logic have been considered in [DG21]. We note that the semiring semantics for modal logic in [DG21] is different to the one obtained from the translation into FO (cf. Section 5.2.5). A detailed comparison of these semantics and their relationship to bisimulation may deserve further study. Another interesting case are description logics, which have been considered, for instance, in [DG19, BOPP20]. Unlike datalog or FO, where semiring provenance usually tracks only facts or instantiated literals, we may here want to track both *assertions* (i.e., facts) and *concept inclusions* (i.e., rules) that are used to derive a given consequence.

Semiring semantics for FO has further been used in team semantics [BHK$^+$23]. By using annotated teams (called K-*teams*) for different semirings K, several forms of team semantics can be treated in a uniform way, including classical team semantics (in the Boolean semiring), multiteams (in $\mathbb{N}$), and probabilistic teams (in $\mathbb{R}_{\geq 0}$). This semantics takes a different perspective compared to provenance. So far, we have evaluated formulae to semiring values, and then viewed these values as provenance information. Here, the logic is enriched by operators such as "$\leq$", "$\neq$", and "$\perp$?" (meaning "$= 0$") that can be used to compare the values of subformulae; for an entire formula ψ one is then only interested in the question $[\![\psi]\!] \stackrel{?}{=} 0$. It may be worthwhile to analyse the effect of these additions to FO also outside of the team semantics setting.

Chapter 6

Infinite Games

Two-player games with infinite plays are of fundamental importance in many areas of logic and computer science, especially in the formal analysis of reactive systems, where they model the nonterminating interaction between a system and its environment. In such a game, the *objective* or *winning condition* of the player who represents the system specifies the desired set of behaviours of the system. The most basic classes of such objectives are *reachability* and *safety* objectives defined by a set of states (positions) that the player should reach or avoid, which intuitively means that only finite prefixes of the infinite plays are relevant for the winning condition. Winning objectives that impose nontrivial constraints also on infinite plays, such as *liveness* or *parity* conditions, are much harder to analyse. This provides an interesting perspective for semiring provenance, which can give additional information about winning strategies in these games.

It has long been known that games and logic are intimately connected and this holds true also for semiring semantics, in two ways. First, we can characterise semiring semantics through winning strategies in the associated model-checking games. In other words, we use games to better understand semiring semantics. This has been done for first-order logic (using the notion of proof trees that is equivalent to winning strategies) in [GT17a], and for LFP in form of the Sum-of-strategies theorem [Naa19, DGNT21] that we summarise in Section 6.2 below. A related approach are the provenance games in [KLZ13], which provide an alternative to semiring provenance based on games.

On the other hand, we can also use semiring provenance to better understand games. For reachability and safety objectives, provenance information can be defined by an explicit recursion on the game arena, providing detailed information about winning strategies [GT20]. For more complicated games where such an explicit definition is difficult, we can instead use semiring semantics. Indeed, when we describe the winning region by a logical formula, we can evaluate this formula in semiring semantics to provide additional information *why* a given game is won. We explore this in a detailed case study for Büchi games (Section 6.3), which are the simplest case of infinite games with a nontrivial winning condition for infinite plays, including a genuine alteration of least and greatest fixed points on the logical side. It turns out that semiring semantics for LFP leads to meaningful provenance information for Büchi games, describing precisely all winning strategies within a broad class of strategies, which we call *absorption-dominant* strategies. This information enables applications such as computing minimal repairs to a given game graph. The idea to use semiring semantics for games is not limited to Büchi games. We end this chapter with a brief overview on more expressive games.

6.1 Games and Strategies

We begin by introducing notation for games and strategies. We always consider two-player games on finite graphs, which we represent as structures $\mathcal{G} = (V, V_0, V_1, E)$ called *game arenas*, where V is a finite set of positions with a disjoint decomposition $V = V_0 \dot\cup V_1$ into positions of Player 0 and positions of Player 1. We often use the variable $\sigma \in \{0, 1\}$ to refer to either of the two players. The edge relation $E \subseteq V \times V$ specifies the possibles moves. For a position $v \in V$, we write $vE := \{w \mid vw \in E\}$ for the immediate successors of v. A play from an initial position v_0 is a maximal (finite or infinite) path $v_0v_1v_2\dots$ through $\mathcal{G}$ where the successor $v_{i+1} \in v_iE$ is chosen by Player 0 if $v_i \in V_0$ and by Player 1 if $v_1 \in V_1$.

The basic winning condition is *move or lose*, which means that a position $v \in V_\sigma$ which is terminal (i.e., $vE = \emptyset$) is considered losing for the active Player σ. However, for infinite games we always assume $vE \neq \emptyset$ for all $v \in V$, so that all plays are infinite. We give an overview on the winning conditions we consider. We always specify them from the perspective of Player 0, so Player 1 wins exactly those plays not won by Player 0.

- In a *reachability* game $\mathcal{G} = (V, V_0, V_1, E, F)$, a play is considered winning for Player 0 once a position in $F \subseteq V$ is reached. The dual condition is a *safety* game, where Player 0 wins plays that do not visit F.
- In a *Büchi* game $\mathcal{G} = (V, V_0, V_1, E, F)$ with target set $F \subseteq V$, a play $v_0v_1v_2\dots$ is winning for Player 0 if F is visited infinitely often (i.e., $v_i \in F$ for infinitely many $i < \omega$). This objective is also called *repeated reachability* and is a special case of a *liveness* objective. The dual condition is called *co-Büchi* objective, where Player 0 wins plays that visit F only finitely often.
- A *generalised Büchi* game is of the form $\mathcal{G} = (V, V_0, V_1, E, F_1, \dots, F_k)$ with multiple target sets $F_1, \dots, F_k \subseteq V$. Player 0 wins a play if each set $F_1, \dots, F_k$ is visited infinitely often.
- A *parity* game is of the form $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ with priority labelling $\Omega\colon V \to \{0, \dots, k-1\}$ for some $k < \omega$. The winner of an infinite play $v_0v_1v_2\dots$ is determined by the sequence of priorities $\Omega(v_0)\Omega(v_1)\Omega(v_2)\dots$, with Player 0 winning if the minimal priority occurring infinitely often is even, and Player 1 if it is odd.
- A *request-response* game is of the form $\mathcal{G} = (V, V_0, V_1, E, R_1, \dots, R_k, Q_1, \dots, Q_k)$ with sets $R_i, Q_i \subseteq V$. An infinite play $v_0v_1v_2\dots$ is won by Player 0 if for each each $1 \leq i \leq k$ and each $n < \omega$, if $v_n \in R_i$ then there is $n' > n$ with $v_{n'} \in Q_i$. In other words, each visit to R_i is eventually answered by a visit to Q_i.
- A *Muller* game is of the form $\mathcal{G} = (V, V_0, V_1, E, \Omega, \mathcal{F})$, where $\Omega\colon V \to \{0, \dots, k-1\}$ is a priority labelling (as for parity games) and $\mathcal{F} \subseteq \mathcal{P}(\{0, \dots, k-1\})$. Player 0 wins a play $v_0v_1v_2\dots$ if the set of priorities occurring infinitely often in $\Omega(v_0)\Omega(v_1)\Omega(v_2)\dots$ is contained in $\mathcal{F}$.

The winning region of Player σ is the set of positions $v \in V$ such that Player σ has a winning strategy from v, i.e., a strategy that guarantees to result in a winning play, no matter what the opponent does. Such a strategy can be represented in different ways, for instance as a function $f\colon V^*V_\sigma \to V$ that assigns a next position to each partial play ending in a position of Player σ, or simply $f\colon V_\sigma \to V$ if the strategy is positional (i.e., does not need memory). Here we instead follow the approach of Grädel and Tannen [GT20] and define strategies as trees comprised of all plays consistent with the strategy. Formally, these are subtrees of the tree unraveling of a game.

Definition 6.1.1. Given a game arena $\mathcal{G} = (V, V_0, V_1, E)$, the *tree unraveling* from v_0 is the tree $\mathcal{T}(\mathcal{G}, v_0)$ whose nodes are all finite paths ρ from v_0 in $\mathcal{G}$ and whose edges are $\rho \to \rho v$ for $v \in V$. We often denote a node of $\mathcal{T}(\mathcal{G}, v_0)$ as ρv to indicate a finite path ending in $v \in V$. We write $|\rho|$ for the length of ρ and $\rho \sqsubseteq \rho'$ if ρ is a (not necessarily strict) prefix of ρ'.

Strategies can then be defined as subtrees of the tree unraveling, which allows for a more visual way to reason about strategies. An important detail is that the strategy tree only contains positions that are reachable when following the strategy. Moreover, the tree unraveling and all strategies are finitely branching, as we only consider finite game graphs.

Definition 6.1.2. A *strategy* $\mathcal{S}$ (of Player 0) from v_0 in $\mathcal{G}$ is a subtree of $\mathcal{T}(\mathcal{G}, v_0)$ induced by a node set W satisfying the following conditions:

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

We say that a play is *consistent with* $\mathcal{S}$ if it corresponds to a maximal path through $\mathcal{S}$ from the root v_0. A strategy $\mathcal{S}$ is *winning* if all consistent plays are winning (according to some winning condition).

We frequently use the following abbreviations when working with strategies. We write $\rho \in \mathcal{S}$ instead of $\rho \in W$ and we often refer to paths of the form $\rho v \in \mathcal{S}$ as *occurrences of* v in $\mathcal{S}$. When we depict strategies graphically, we represent finite paths ρv just by their last position v to ease readability (notice that in the tree unravelling, ρ can be reconstructed from v by following the path to the root). See Fig. 6.1 for an example. For $v \in V_0$, we further write $\mathcal{S}(\rho v) = w$ if $\rho v w$ is the (unique) successor of ρv in $\mathcal{S}$. If $\mathcal{S}$ is positional, we may also write $\mathcal{S}(v)$ to denote the unique successor of v chosen by $\mathcal{S}$. We write $\mathrm{Strat}_{\mathcal{G}}(v)$ and $\mathrm{WinStrat}_{\mathcal{G}}(v)$ to denote the set of all (winning) strategies of Player 0 from position $v \in \mathcal{G}$, and we drop $\mathcal{G}$ if the game is clear from the context.

6.2 The Sum-of-Strategies Theorem for LFP

It has been shown by Grädel and Tannen [GT20] that the provenance analysis for FO and posLFP is intimately connected with the provenance analysis of reachability games. Evaluation strategies to establish the truth of first-order formulae (such as the proof trees considered in [GT17a]), are really winning strategies for reachability games on acyclic game graphs (which only admit finite plays). For posLFP the situation is similar, but the associated model-checking games may have cycles and thus admit infinite plays. The winning plays for the verifying player, however, must reach a winning position (a true literal) in a finite number of steps, so winning strategies are still finite. By annotating terminal positions (i.e., literals)

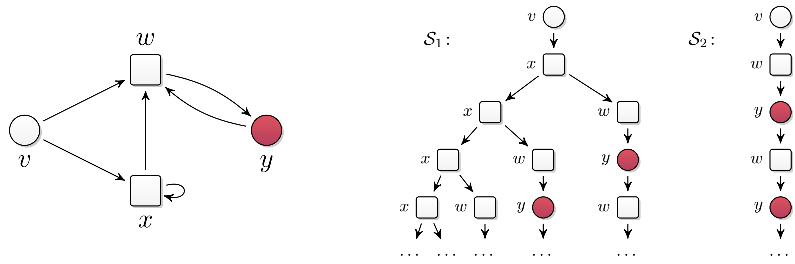

Figure 6.1: A simple Büchi game and two strategies of Player 0. Round nodes belong to Player 0, rectangular nodes to Player 1, red nodes are in F. Only $\mathcal{S}_2$ is winning.

with semiring values and propagating these values along the edges to the remaining positions, one obtains provenance values that coincide with the syntactically defined semantics $\pi[\![\psi]\!]$.

This insight has been generalised to the model-checking games for full LFP, resulting in the *Sum-of-strategies theorem* for LFP in [Naa19, DGNT21]. These model-checking games are parity games, which are considerably more complex and do not allow for a simple propagation of values from terminal positions. However, semiring semantics $\pi[\![\psi]\!]$ in absorptive, fully-continuous semirings can still be characterised in terms of the – now infinite – winning strategies: if we associate with each strategy a semiring value that results from multiplying the values of all literals occurring in the strategy, then $\pi[\![\psi]\!]$ is equal to the sum over these values for all *winning* strategies.

The proof presented here is included in this thesis for the sake of completeness, and is taken from [DGNT19] (the full version of [DGNT21]) with only minor modifications, including examples and figures. The only relevant changes are a simplified argument for the induction on the fixed-point iteration that makes use of strong distributivity, and a correction for a small inaccuracy when applying Proposition 3.5.8 to the infimum of a **gfp**-iteration.

6.2.1 Model-Checking Games

We first describe how model-checking games can be adapted to semiring semantics. In classical semantics, model-checking games are defined for a given formula ψ and a fixed structure $\mathfrak{A}$ (see, e.g., [Grä11] for background). However, the *game arena* of such a game depends only on ψ and the *universe* of the given structure. It is only the labelling of the terminal positions as winning, for either the Verifier (Player 0) or the Falsifier (Player 1), that depends on which of the literals in $\mathrm{Lit}_A(\tau)$ are true in $\mathfrak{A}$. Hence the definition readily generalizes to semiring semantics, now labelling terminal positions by semiring values.

Game arena. We briefly recall the construction of the game arena. Let ψ be an LFP-sentence in negation normal form with a relational vocabulary τ, and let A be a (finite) universe. The model checking game $\mathcal{G}(A, \psi)$ has positions $\varphi(\boldsymbol{a})$, obtained from a subformula $\varphi(\boldsymbol{x})$ of ψ by instantiating the free variables $\boldsymbol{x}$ with a tuple $\boldsymbol{a} \subseteq A$. At a disjunction $(\varphi \vee \vartheta)$, Player 0 (Verifier) moves to either φ or ϑ, and at a conjunction, Player 1 (Falsifier) makes an analogous move. At a position $\exists x\, \vartheta(\boldsymbol{a}, x)$, Verifier selects an element $b \in A$ and moves to $\vartheta(\boldsymbol{a}, b)$, whereas at positions $\forall x \vartheta(\boldsymbol{a}, x)$ the move to to the next position $\vartheta(\boldsymbol{a}, b)$ is done by Falsifier. For every subformula of φ of form $\varphi = [\mathbf{lfp}\, R\, \boldsymbol{x}.\ \vartheta(R, \boldsymbol{x})](\boldsymbol{x})$ or $\varphi = [\mathbf{gfp}\, R\, \boldsymbol{x}.\ \vartheta(R, \boldsymbol{x})](\boldsymbol{x})$ we add moves from positions $\varphi(\boldsymbol{a})$ to $\vartheta(\boldsymbol{a})$, and from positions $R\boldsymbol{a}$ back to $\vartheta(\boldsymbol{a})$, for every tuple $\boldsymbol{a}$ (thereby introducing cycles). Since these moves are unique, it makes no difference to which of the two players we assign the positions $\varphi(\boldsymbol{a})$ and $R\boldsymbol{a}$. The terminal positions of $\mathcal{G}(A, \psi)$ are always literals in $\mathrm{Lit}_A(\tau)$.

These games may have cycles and thus admit infinite plays. The winning condition for infinite plays is the *parity condition*. We assign to each fixed-point variable R a priority, which is even for greatest fixed points and odd for least fixed points, in such a way that the following property holds: if a variable T occurs within the scope of the fixed-point formula defining the variable R (so that R is the outer variable), then the priority of R is smaller or equal to the priority of T. An infinite play is won by Verifier if the least priority occurring infinitely often in the play is even, otherwise it is won by Falsifier.

Semiring values. Let K be an absorptive, fully-continuous semiring. Given a parity game $\mathcal{G}(A, \psi)$, a K-interpretation π provides a valuation of the terminal positions, which induces provenance values for plays and strategies.

Definition 6.2.1. A *finite* play $\rho = \varphi_0(\boldsymbol{a}_0), \dots, \varphi_t(\boldsymbol{a}_t)$ ends in a terminal position $\varphi_t(\boldsymbol{a}_t) \in \mathrm{Lit}_A(\tau)$ which we call the *outcome* of ρ. We identify the provenance value of ρ with the value of its outcome, i.e. we put $\pi[\![\rho]\!] := \pi[\![\varphi_t(\boldsymbol{a}_t)]\!]$. For an *infinite* play ρ we put $\pi[\![\rho]\!] := 1$ if ρ is a wining play for Verifier, and $\pi[\![\rho]\!] := 0$ otherwise.

We use the tree representation of strategies (Definition 6.1.2). Recall that $\mathrm{Strat}(\varphi(\boldsymbol{a}))$ denotes the set of evaluation strategies for the subformula $\varphi(\boldsymbol{a})$ of ψ, i.e. the set of all (not necessarily positional) strategies that Verifier has from position $\varphi(\boldsymbol{a})$ in the parity game $\mathcal{G}(A, \psi)$. We also use the notation $\mathrm{WinStrat}(\varphi(\boldsymbol{a}))$ to denote the set of winning strategies, but these are defined differently compared to classical semantics: since the game $\mathcal{G}(A, \psi)$ is independent of π, we do not know the valuations of literals and thus call a strategy $\mathcal{S}$ *winning* if all infinite plays are winning according to the parity condition (finite plays are not taken into account, i.e., are always considered winning).

For a strategy $\mathcal{S} \in \mathrm{Strat}(\varphi(\boldsymbol{a}))$, we write $\mathrm{Plays}(\mathcal{S})$ for the set of maximal paths through $\mathcal{S}$ starting at the root. These are precisely the plays from $\varphi(\boldsymbol{a})$ in $\mathcal{G}(A, \psi)$ consistent with the choices of $\mathcal{S}$, and the provenance value of a strategy is the product over the values of all these plays.

Definition 6.2.2. The provenance value of a strategy $\mathcal{S}$, evaluated in π, is defined as:

$$\pi[\![\mathcal{S}]\!] := \widehat{\prod_{\rho \in \mathrm{Plays}(\mathcal{S})}} \pi[\![\rho]\!].$$

For a literal $L \in \mathrm{Lit}_A(\tau)$, we further write $\#_L(\mathcal{S}) \in \mathbb{N} \cup \{\infty\}$ for the number of plays $\rho \in \mathrm{Plays}(\mathcal{S})$ with outcome L. Assuming that all infinite plays are winning for Verifier, we can then write the product as follows, using the infinitary power for $\#_L(\mathcal{S})$:

$$\pi[\![\mathcal{S}]\!] = \prod_{L \in \mathrm{Lit}_A(\tau)} \pi(L)^{\#_L(\mathcal{S})}.$$

The two definitions are clearly equivalent by associativity of infinite products.

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

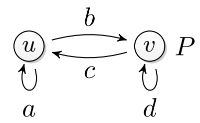

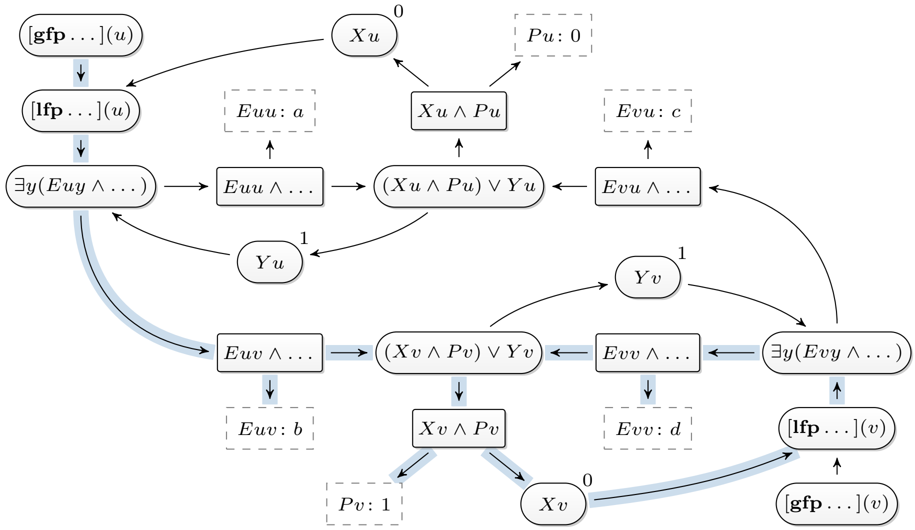

Figure 6.2: Illustration of a model-checking game. Round nodes belong to Verifier, rectangular nodes to Falsifier. Nodes are labelled by their priorities (if defined). Terminal positions (dashed border) include valuations in $\mathbb{S}^\infty(X)$.

6.2.2 Characterisation through Strategies

The following Sum-of-strategies theorem achieves a precise characterisation of semiring semantics for LFP through the provenance values of winning strategies in $\mathcal{G}(A, \psi)$.

Theorem 6.2.3 (Sum of strategies [DGNT21])**.** *Let $\psi \in$ LFP, and and let $\pi\colon \mathrm{Lit}_A(\tau) \to K$ be a K-interpretation into an absorptive, fully-continuous semiring K. Then*

$$\pi[\![\psi]\!] = \bigsqcup \{\pi[\![\mathcal{S}]\!] \mid \mathcal{S} \in \mathrm{Strat}_{\mathcal{G}(A,\psi)}(\psi)\}.$$

For a model-compatible interpretation in $\mathbb{S}^\infty(X, \overline{X})$, the above theorem states that the value of a sentence ψ is a sum of monomials $x_1^{e_1} \cdots x_k^{e_k}$. Each such monomial corresponds to a strategy $\mathcal{S}$ for Verifier that uses precisely the literals labelled by $x_1, \dots, x_k$, and each literal x_i is used precisely e_i many times (i.e., there are e_i plays with outcome x_i). By using dual indeterminates, we make sure that these literals are consistent and hence represent actual evaluation strategies of ψ. Notice that we only obtain information about strategies whose corresponding monomials are not absorbed; we study these strategies in more detail in Section 6.3.

Example 6.2.4. An example of a model-checking game for a given formula $\varphi(x)$ and a model-compatible $\mathbb{S}^\infty(X)$-interpretation π (with indeterminates $X = \{a, b, c, d\}$) is shown in Fig. 6.2. The interpretation is a small annotated graph with only two nodes to keep the size of the model-checking game reasonable. We track the edge literals with indeterminates; the literals for the unary predicate P, which holds only at v, are not tracked and simply mapped to 0 or 1. The formula is nontrivial and contains a genuine alteration of a greatest and least fixed point, which is expressed in the game through the priorities 0 and 1. It asserts that there is a path from x that infinitely often visits P.

On the given graph, witnesses for $\varphi(u)$ are infinite paths that infinitely often visit v. There are infinitely many such paths, but the simplest ones (in terms of the different edges they use) are the paths $uvvvv\dots$ and $uvuvuv\dots$ which correspond to the monomials bd^∞ and $b^\infty c^\infty$. And indeed, $\pi[\![\varphi(u)]\!] = bd^\infty + c^\infty d^\infty$. Notice that the edge a does not appear in the result, so we can conclude that its existence does not affect the truth of $\varphi(u)$.

Let us now consider the evaluation strategies for $\varphi(u)$ from the game-theoretic perspective by looking at the model-checking game in Fig. 6.2. There are four positions at which Verifier can make a decision: the two nodes labelled $\exists y(\dots)$ and the two disjunctions in the center of the figure. Hence there are 16 positional strategies in total. One of these strategies is highlighted in blue and has the provenance value bd^∞, as there is one play ending in Euv (value b) and there are arbitrarily long plays ending either in Evv (value d) or in Pv (value 1), depending on the choices of Falsifier.

Most of the other 15 positional strategies allow infinite paths with least priority 1, which means that they are not winning for Verifier and thus have provenance value 0, for instance by choosing the cycle $\exists y(\dots) \;\to\; Euu \wedge \dots \;\to\; (Xu \wedge Pu) \vee Yu \;\to\; Yu$. The only remaining winning strategy has the provenance value $b^\infty c^\infty$. One can further observe that non-positional strategies only lead to monomials with additional variables which are then absorbed, so by taking the sum (or supremum) over all winning strategies we obtain, again, $\pi[\![\varphi(u)]\!] = bd^\infty + b^\infty c^\infty$.

6.2.3 Proof of the Sum-of-Strategies Theorem

The proof of Theorem 6.2.3 is essentially an induction on ψ, relating semiring semantics $\pi[\![\psi]\!]$ with the construction of the game $\mathcal{G}(A,\psi)$. However, the cases for fixed-point formulae are more involved and require an induction on the fixed-point iteration that relates the stages of the iteration with truncations of the strategies in $\mathcal{G}(A,\psi)$. As in Chapter 4, these strategies can be infinite for greatest fixed points, which makes the proof challenging.

Notation

For the remaining section, we fix an absorptive, fully-continuous semiring K and a finite, relational signature τ. We further fix a finite universe A and simply write $\mathcal{G}(\psi)$ instead of $\mathcal{G}(A,\psi)$. We write **fp** to refer to either **lfp** or **gfp**. The relation symbols R occurring in fixed-point subformulae $[\mathbf{fp}\, R\, \boldsymbol{x}.\, \varphi(R,\boldsymbol{x})](\boldsymbol{y})$ are called *fixed-point variables* and we assume, to avoid confusion, that each variable is only used in one fixed-point subformula.

Given a fixed-point variable R (of arity r) and a finite path $\rho = v_0 v_1 \dots v_k$ in $\mathcal{G}(\psi)$, we write $|\rho|_R = |\{i \mid v_i = R\boldsymbol{a} \text{ for some } \boldsymbol{a} \in A^r\}|$ for the number of occurrences of R-atoms along the path. For a strategy $\mathcal{S}$ (whose nodes are finite paths), we further say that a node $\rho = v_0 v_1 \dots v_k \in \mathcal{S}$ is an *R-node* if $v_k = R\boldsymbol{a}$ for some $\boldsymbol{a} \in A^r$, i.e., if ρ ends in an R-atom.

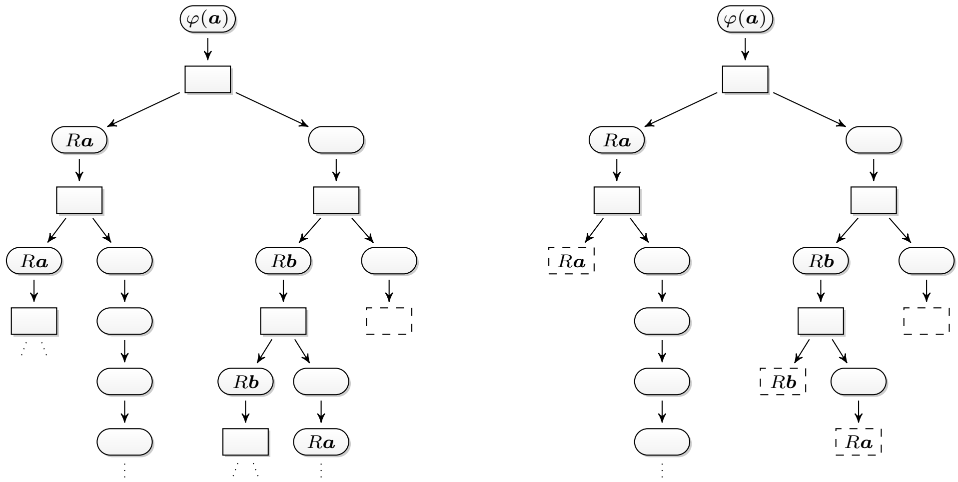

Figure 6.3: A visualisation of a strategy $\mathcal{S}$ and its $(R,2)$-truncation.

Strategy Truncations

We begin by introducing the notion of strategy truncations. For a given fixed-point-relation symbol R, the n-truncation of a strategy is obtained by cutting off the subtrees at the n-th occurrences of literals $R\boldsymbol{a}$. This is similar to the n-cut of derivation trees in Chapter 4, except that we do not cut off at a fixed height, as the n-th occurrence of R can be at a different height for each path. In fact, the n-truncation can be of infinite height if there is a path on which R occurs fewer than n times.

Definition 6.2.5. Let $\mathcal{S} = (W, F)$ be a strategy and let R be a fixed-point variable of arity r. For $n \geq 1$, the (R, n)-truncation of $\mathcal{S}$ is the subtree induced by $W' \subseteq W$ with

$$W' = \{\rho \in W \mid \text{either } |\rho|_R < n, \text{ or } \rho \text{ is an } R\text{-node with } |\rho|_R = n\}.$$

If R is clear from the context, we write $\mathcal{S}\|_n$ for the (R, n)-truncation of $\mathcal{S}$.

See Fig. 6.3 for an illustration. The provenance value of a truncation $\mathcal{S}\|_n$ is defined as for strategies (cf. Definition 6.2.2), by taking the product over the values of all maximal paths through $\mathcal{S}\|_n$. Notice that the truncation creates new leaf nodes, hence $\pi[\![\mathcal{S}\|_n]\!]$ is defined only for interpretations π over the extended signature $\tau \cup \{R\}$ Recall that for a given valuation $g\colon A^r \to K$, we write $\pi_{[R\mapsto g]}$ for the extension of π by $\pi(R\boldsymbol{a}) = g(\boldsymbol{a})$ (cf. Section 5.2).

Example 6.2.6. Consider the setting of Fig. 6.2 and let $\mathcal{S}$ be the strategy highlighted in blue. Given a valuation $g\colon \{u, v\} \to \mathbb{S}^\infty(A)$, the (X, n)-truncations of $\mathcal{S}$ have the values:

- $\pi_{[X\mapsto g]}[\![\mathcal{S}\|_1]\!] = b \cdot g(v)$,
- $\pi_{[X\mapsto g]}[\![\mathcal{S}\|_2]\!] = bd \cdot g(v)$,
- $\pi_{[X\mapsto g]}[\![\mathcal{S}\|_3]\!] = bd^{n-1} \cdot g(v)$, for $n \geq 1$.

We see that for $g = \mathbf{1}$, these values converge to the value $\pi[\![\mathcal{S}]\!] = bd^\infty$ of the original strategy.

The convergence seen in the previous example can be formalised as follows.

Lemma 6.2.7. *Let* $\varphi = [\mathbf{fp}\, R\boldsymbol{x}.\, \vartheta](\boldsymbol{y}) \in \mathrm{LFP}(\tau)$ *and* $\mathcal{S} \in \mathrm{WinStrat}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))$*. Let further* π *be a* K*-interpretation over* A *and* τ*.*

(1) *If* $\varphi = [\mathbf{lfp}\, R\boldsymbol{x}.\, \vartheta](\boldsymbol{y})$*, then* $\displaystyle\bigsqcup_{1 \le n < \omega} \pi_{[R\mapsto \mathbf{0}]}[\![\mathcal{S}\|_n]\!] = \pi[\![\mathcal{S}]\!]$.

(2) *If* $\varphi = [\mathbf{gfp}\, R\boldsymbol{x}.\, \vartheta](\boldsymbol{y})$*, then* $\displaystyle\bigsqcap_{1 \le n < \omega} \pi_{[R\mapsto \mathbf{1}]}[\![\mathcal{S}\|_n]\!] = \pi[\![\mathcal{S}]\!]$.

Proof. For (1), we show that there is a k with $\mathcal{S}\|_k = \mathcal{S}$. Suppose this would not be the case. Then for each n, some truncation must happen and hence there is a path in $\mathcal{S}$ with at least n R-nodes. But the game arena is finite and $\mathcal{S}$ is thus finitely branching, so by Kőnig's lemma there must be an infinite path through $\mathcal{S}$ with infinitely many R-nodes. But since R is the fixed-point variable with smallest priority in $\mathcal{G}(\varphi)$, this path is losing by the parity condition. Since this path also exists in the original strategy $\mathcal{S}$, this contradicts the assumption $\mathcal{S} \in \mathrm{WinStrat}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))$.

So let k be minimal with $\mathcal{S}\|_k = \mathcal{S}$. Clearly $\pi_{[R\mapsto\mathbf{0}]}[\![\mathcal{S}\|_n]\!] = \pi[\![\mathcal{S}]\!]$ for all $n \ge k$. For $n < k$, a truncation happens in $\mathcal{S}\|_n$ and hence $\pi_{[R\mapsto\mathbf{0}]}[\![\mathcal{S}\|_n]\!] = 0$. It follows that the supremum equals $\pi[\![\mathcal{S}]\!]$.

For (2), we first note that the truncations $\pi_{[R\mapsto\mathbf{1}]}[\![\mathcal{S}\|_n]\!]$ indeed form a chain, since truncated leafs receive the greatest value 1. We group literals together in the product $\pi_{[R\mapsto\mathbf{1}]}[\![\mathcal{S}\|_n]\!]$ and apply the Chain splitting lemma (3.2.8). Notice that we can ignore the R-leaves created by the truncation, as the value 1 does not affect the product. We get:

$$\bigsqcap_{n<\omega} \pi_{[R\mapsto\mathbf{1}]}[\![\mathcal{S}\|_n]\!] = \prod_{L\in\mathrm{Lit}_A} \bigsqcap_{n<\omega} \pi(L)^{\#_L(\mathcal{S}\|_n)} = \prod_{L\in\mathrm{Lit}_A} \pi(L)^{e_L}, \text{ for } e_L = \bigsqcup_{n<\omega} \#_L(\mathcal{S}\|_n).$$

It remains to prove that this is equal to $\pi[\![\mathcal{S}]\!]$, i.e., that $e_L = \#_L(\mathcal{S})$ for all literals. This is true because each node of $\mathcal{S}$ is eventually contained in $\mathcal{S}\|_n$ (for sufficiently large n). More precisely, consider a literal L. If $\#_L(\mathcal{S})$ is finite, then for sufficiently large n, we have $\#_L(\mathcal{S}\|_n) = \#_L(\mathcal{S})$ and thus $e_L = \#_L(\mathcal{S})$. If $\#_L(\mathcal{S}) = \infty$, then for each k there is a sufficiently large n such that $\#_L(\mathcal{S}\|_n) \ge k$ and thus also $e_L = \infty$. □

The Puzzle Lemma

Strategies for least fixed points are essentially finite, as the fixed-point variable cannot be visited infinitely often in a play, but strategies for greatest fixed points are infinite objects which pose a major obstacle for the proof. The following lemma is the key observation needed for **gfp**-formulae, and is the most challenging step of the proof.

Lemma 6.2.8 (Puzzle lemma)**.** *Let* $\varphi = [\mathbf{gfp}\, R\boldsymbol{x}.\, \vartheta](\boldsymbol{y}) \in \mathrm{LFP}(\tau)$ *with* R *of arity* r*, and let* $\boldsymbol{a} \in A^r$*. Let* π *be a* K*-interpretation (over* A *and* τ*). Let further* $(\mathcal{S}_i)_{1\le i<\omega}$ *be a family of strategies in* $\mathrm{WinStrat}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))$ *such that* $(\pi_{[R\mapsto\mathbf{1}]}[\![\mathcal{S}_i\|_i]\!])_{1\le i<\omega}$ *is a descending chain. Then there is a winning strategy* $\hat{\mathcal{S}} \in \mathrm{WinStrat}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))$ *with* $\pi[\![\hat{\mathcal{S}}]\!] \ge \bigsqcap_i \pi_{[R\mapsto\mathbf{1}]}[\![\mathcal{S}_i\|_i]\!]$.

This lemma intuitively states that taking the infimum over the values of strategy truncations results in a value that is meaningful, in the sense that it is realised by an actual (infinite) strategy. The following example provides some intuition why this may not be obvious. The key problem is that the strategies $\mathcal{S}_i$ can all be different. In particular, it can happen that for every i, the provenance value of the truncation $\mathcal{S}_i\big\|_i$ is larger than the value of the full strategy $\mathcal{S}_i$. The insight of the lemma is that we can always use one of the truncations $\mathcal{S}_i\big\|_i$ (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 6.2.9. Consider the following setting in $\mathbb{S}^\infty(A)$:

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

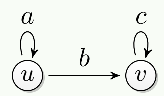

Let $\mathcal{S}_i$ be the strategy corresponding to the infinite path that cycles $i-1$ times via a, then uses edge b and finally cycles forever via c. The i-truncation then cuts off $\mathcal{S}_i$ after taking the edge b and we obtain the provenance values

$$\pi[\![\mathcal{S}_i]\!] = a^{i-1}bc^\infty \qquad \text{and} \qquad \bigsqcap_{1\le i<\omega} \pi_{[R\mapsto \mathbf{1}]}[\![\mathcal{S}_i\big\|_i]\!] = \bigsqcap_{1\le i<\omega} a^{i-1}b = a^\infty b.$$

We see that the infimum only contains the variables a and b, although there is no winning strategy with this provenance value. Instead, we obtain $\mathcal{S}$ by repeating the cycling part of any truncation $\mathcal{S}_i\big\|_i$ (without the problematic literal b). This results in the strategy $\mathcal{S}$ with value $\pi[\![\mathcal{S}]\!] = a^\infty$ that corresponds to the path always cycling via a. Notice that 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 $\mathcal{S}_i\big\|_i$) are always winning.

Proof of the Puzzle lemma. As a first step to prove the Puzzle lemma, we apply the Chain splitting lemma (Lemma 3.2.8) to the infimum and obtain:

$$\bigsqcap_{1\le i<\omega} \pi_{[R\mapsto \mathbf{1}]}[\![\mathcal{S}_i\big\|_i]\!] = \prod_{L\in \mathrm{Lit}_A} \pi(L)^{n_L} \quad \text{with} \quad n_L = \bigsqcup_{1\le i<\omega} \#_L(\mathcal{S}_i\big\|_i).$$

Literals with $n_L = \infty$ (such as the edge a in the example) can appear arbitrarily often in $\hat{\mathcal{S}}$, so they do not impose any restrictions. If $n_L < \infty$, then we must have $\#_L(\hat{\mathcal{S}}) \le n_L$ to guarantee that the provenance value of $\hat{\mathcal{S}}$ is larger than the infimum. We therefore call literals L with $n_L < \infty$ (such as the edge b in the example above) *problematic*. The outline of the proof is as follows:

- We decompose the trees $\mathcal{S}_i\big\|_i$ into *layers* based on the appearance of R-nodes.
- We choose a sufficiently large i such that there is one such layer in $\mathcal{S}_i\big\|_i$ which does not contain any problematic literals.
- We construct $\hat{\mathcal{S}}$ by first following $\mathcal{S}_i\big\|_i$ 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 shape of the puzzle pieces ensures that $\hat{\mathcal{S}}$ 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.

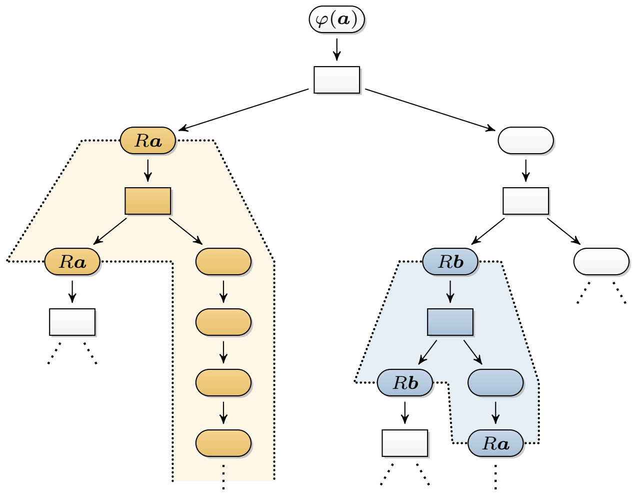

Figure 6.4: A visualisation of a strategy. The coloured nodes form the first layer (for $k = 1$). The two trees in this layer are puzzle pieces, the left one has an infinite winning path.

Decomposition into layers. Fix an i and let $\mathcal{S}_i\|_i = (W, F)$. Recall that V is the set of positions of $\mathcal{G}(\psi)$. For each $n \geq 0$, we define the sets

$$W_{\leq n} = \{\rho \in W \mid |\rho|_R \leq n\}, \quad W^+_{\leq n} = W_{\leq n} \cup \{\rho v \in W \mid \rho \in W_{\leq n},\, v \in V\}.$$

We sort the nodes $\rho \in W$ into layers based on the number of R-nodes on the path to ρ. 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

$$W_j = W^+_{\leq j \cdot k} \setminus W_{\leq (j-1)\cdot k} \quad \text{where} \quad k = |A|^r + 2.$$

Avoiding problematic literals. Let $n = \sum\{n_L \mid L \in \mathrm{Lit}_A\,,\, n_L < \infty\}$ 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 $|\rho|_R \leq (j-1)\cdot k < i$ for all $\rho \in \mathcal{S}_i\|_i$. 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 $\hat{\mathcal{S}} = \mathcal{S}_i$ 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 $\hat{\mathcal{S}}$ 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 $\pi[\![\hat{\mathcal{S}}]\!] \geq \pi_{[R\mapsto\mathbf{1}]}[\![\mathcal{S}_i\|_i]\!]$ as required.

Let T be one of the components in W_j, so T is a tree. We call a path in T *untruncated* if it is infinite or ends in a terminal position (not an R-node), so that it corresponds to a (suffix of a) play consistent with $\mathcal{S}_i$.

Definition 6.2.10. A *puzzle piece* $\mathcal{P} = (W', F')$ is a subtree of W_j such that

(a) the root of $\mathcal{P}$ is an R-node,

(b) for each inner node $\rho \in \mathcal{P}$, we have $\rho F' = \rho F$ ($\mathcal{P}$ contains all successors), and

(c) each maximal path through $\mathcal{P}$ is either untruncated or ends in an R-node.

A puzzle piece P with root ρ *matches* a node $\hat{\rho} \in W_j$ if $\rho = \rho' v$ and $\hat{\rho} = \hat{\rho}' v$ for the same position $v = R\boldsymbol{a}$. A *complete puzzle* is a set of puzzle pieces such that for each piece in the set and all R-leaves ρ of this piece, the set contains a puzzle piece that matches ρ.

First observe that T itself is a puzzle piece by construction of the layers. But it is not clear if T can be repeated, so instead we collect smaller pieces from T by the following process:

(1) Initialize $L = \{\hat{\rho}\}$ where $\hat{\rho}$ is the root of T (which is an R-node).

(2) Pick a node $\rho \in L$ and remove it from L (if L is empty, terminate).

(3) If we have already found a puzzle piece matching ρ, go back to step (2).

(4) Let $\mathcal{P}$ be the subgraph of T from ρ to the next occurrences of R-nodes. More formally, $\mathcal{P}$ is induced by the following set W' of nodes:

$$W' = \{\rho' \in T \mid \rho \sqsubseteq \rho' \text{ and there is no } R\text{-node } \rho'' \text{ with } \rho \sqsubset \rho'' \sqsubset \rho')\}.$$

Then $\mathcal{P}$ is a puzzle piece matching ρ and we add it to our set of pieces.

(5) For each $\boldsymbol{a} \in A^r$: If $\mathcal{P}$ has an R-leaf $\rho' v$ with $v = R\boldsymbol{a}$, add one such leaf $\rho' v$ to L.

(6) Go back to step (2).

We next prove that the definition of $\mathcal{P}$ in step (4) indeed yields a well-defined puzzle piece. It then follows that the process 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 $|\hat{\rho}|_R = (j-1)k+1$ and W_j contains in particular the nodes ρ with $(j-1)k < |\rho|_R \leq jk$. Assume that in (2), we picked a node ρ with $|\rho|_R = n$ (for some n). By definition of W', the piece $\mathcal{P}$ only contains nodes ρ' with $|\rho'|_R \leq n+1$. In particular, the leaves that we add to L in step (5) all satisfy $|\rho'|_R \leq n+1$. We start with $|\hat{\rho}|_R = (j-1)k+1$ and perform at most $k-2 = |A|^r$ iterations, hence we always have $|\rho|_R < jk$ for all $\rho \in L$.

This guarantees that $\mathcal{P}$ is always a puzzle piece in step (4). Inner nodes of $\mathcal{P}$ cannot be R-nodes and hence $\mathcal{P}$ always contains all successors of inner nodes, so constraint (b) of Definition 6.2.10 is satisfied. Constraint (c) is clear from the construction, as both the layers and puzzle pieces result from cutting only through R-nodes.

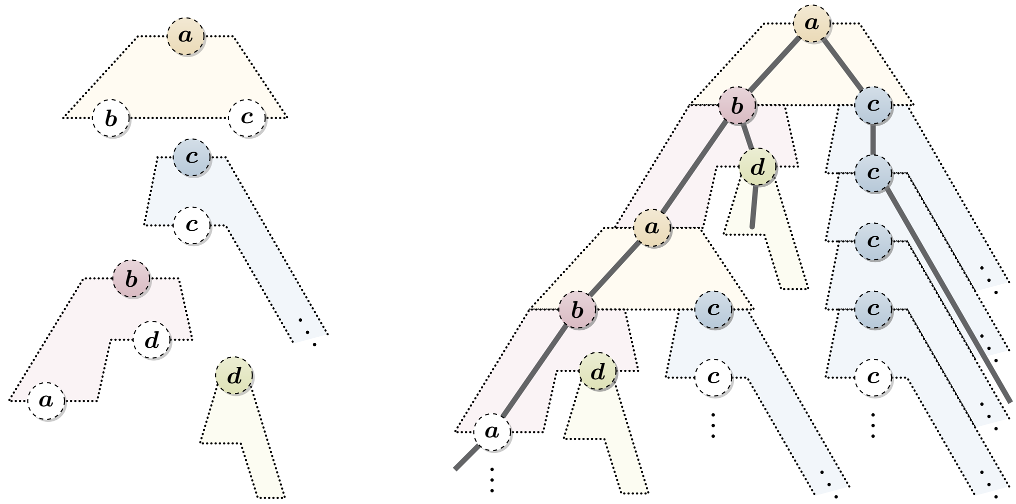

Figure 6.5: Schematic illustration of the pieces in a complete puzzle and their infinite repetition. We abbreviate R-nodes $R\boldsymbol{a}$ (at which we join pieces) by just $\boldsymbol{a}$. The grey lines indicate three paths: One through infinitely many pieces, a finite one, and an infinite one visiting only finitely many pieces (from left to right). All three are winning by construction.

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 Fig. 6.5; 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 $\hat{\mathcal{S}}$. Since puzzle pieces can contain infinite paths or even infinitely many R-leaves, we construct $\hat{\mathcal{S}}$ 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}.
- Given S_n, we construct S_{n+1} as follows. Recall that for $\rho = v_0 \dots v_l \in S_n$, we write $|\rho| = l$ for its length (which equals the depth of ρ in S_n). Consider the set

$$X = \{\rho \in S_n \mid \rho \text{ is an } R\text{-leaf of } S_n \text{ with } |\rho| = n\}.$$

  Because S_n is finitely branching (as we construct it from subtrees of $\mathcal{S}_i\|_i$), this set is finite. Moreover, each $\rho \in X$ 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 ρ. The tree S_{n+1} results from S_n by replacing all leaves $\rho \in X$ with some puzzle piece matching ρ. 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).
- We define $\hat{\mathcal{S}} = \bigcup_{n<\omega} S_n$, so $\hat{\mathcal{S}}$ contains no more R-leaves.

Then $\hat{\mathcal{S}}$ is a well-defined strategy, as the successors of each node are chosen according to the well-defined strategy $\mathcal{S}_i$. Moreover, we have $\pi[\![\hat{\mathcal{S}}]\!] \geq \pi_{[R\mapsto \mathbf{1}]}[\![\mathcal{S}_i\|_i]\!] \geq \prod_i \pi_{[R\mapsto \mathbf{1}]}[\![\mathcal{S}_i\|_i]\!]$ as desired, because the repetition of puzzle pieces does not contain any problematic literals. It remains to argue that $\hat{\mathcal{S}}$ is a winning strategy. Consider a play consistent with $\hat{\mathcal{S}}$ and the corresponding maximal path through $\hat{\mathcal{S}}$. If the path is finite, it ends in a leaf of $\hat{\mathcal{S}}$ which corresponds to a leaf of $\mathcal{S}_i$ and is therefore winning. If the path visits infinitely many puzzle pieces, then it visits infinitely many R-nodes and is thus winning by the parity condition (recall that R has the smallest priority in $\mathcal{G}(\varphi)$). 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. It must then from some point on stay in one puzzle piece, thus corresponding to (the suffix of) an infinite path in $\mathcal{S}_i\|_i$ which must be winning. This completes the proof of the Puzzle lemma. ☐

Main Proof

We are now ready to prove the Sum-of-strategies theorem Theorem 6.2.3. Recall that we want to show:

$$\pi[\![\psi]\!] = \bigsqcup\{\pi[\![\mathcal{S}]\!] \mid \mathcal{S} \in \mathrm{Strat}(\psi)\}.$$

We first prove this for interpretations in the semiring $\mathbb{S}^\infty(X)$, and then generalise it to all absorptive, fully-continuous semirings. This allows us to exploit the properties of $\mathbb{S}^\infty(X)$ for the proof, namely strong distributivity (Corollary 3.5.10) and the characterisation of infima through monomial chains (Proposition 3.5.8). More precisely, we prove the following statement.

Lemma 6.2.11. *Let A be a finite universe and X a finite set of indeterminates. Let $\varphi(\boldsymbol{a})$ be an instantiated formula of signature τ, and let π be a $\mathbb{S}^\infty(X)$-interpretation over A and τ. Then,*

$$\pi[\![\varphi(\boldsymbol{a})]\!] = \bigsqcup\{\pi[\![\mathcal{S}]\!] \mid \mathcal{S} \in \mathrm{WinStrat}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))\}.$$

Proof. The proof is by induction on φ. Notice that the signature and the game $\mathcal{G}(\varphi)$ depend on the formula φ to make the induction work. To ease the presentation, we abbreviate WinStrat as WS in the following.

- $\varphi(\boldsymbol{a}) = L \in \mathrm{Lit}_A(\tau)$. Then $\mathcal{G}(\varphi(\boldsymbol{a}))$ consists only of a terminal position and the unique (trivial) strategy, which is winning, satisfies $\pi[\![\mathcal{S}]\!] = \pi(L)$.
- $\varphi(\boldsymbol{a}) = \varphi_1(\boldsymbol{a}) \vee \varphi_2(\boldsymbol{a})$. The game $\mathcal{G}(\varphi(\boldsymbol{a}))$ is shown on the right. Each strategy $\mathcal{S}$ for $\mathcal{G}(\varphi(\boldsymbol{a}))$ makes a unique choice at $\varphi(\boldsymbol{a})$ and thus either consists of a strategy $\mathcal{S}_1$ for $\mathcal{G}(\varphi_1(\boldsymbol{a}))$ or a strategy $\mathcal{S}_2$ for $\mathcal{G}(\varphi_2(\boldsymbol{a}))$. Using the induction hypothesis (IH), we thus have:

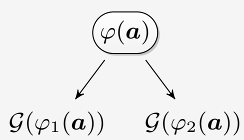

$$\begin{aligned}\pi[\![\varphi(\boldsymbol{a})]\!] &= \pi[\![\varphi_1(\boldsymbol{a})]\!] \sqcup \pi[\![\varphi_2(\boldsymbol{a})]\!] \\ &\overset{\text{(IH)}}{=} \bigsqcup\{\pi[\![\mathcal{S}_1]\!] \mid \mathcal{S}_1 \in \mathrm{WS}_{\mathcal{G}(\varphi_1)}(\varphi_1(\boldsymbol{a}))\} \sqcup \bigsqcup\{\pi[\![\mathcal{S}_2]\!] \mid \mathcal{S}_2 \in \mathrm{WS}_{\mathcal{G}(\varphi_2)}(\varphi_2(\boldsymbol{a}))\} \\ &= \bigsqcup\{\pi[\![\mathcal{S}]\!] \mid \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))\}.\end{aligned}$$

- $\varphi(\boldsymbol{a}) = \varphi_1(\boldsymbol{a}) \wedge \varphi_2(\boldsymbol{a})$. Now each strategy $\mathcal{S}$ for $\mathcal{G}(\varphi(\boldsymbol{a}))$ consists of both a strategy $\mathcal{S}_1$ for $\mathcal{G}(\varphi_1(\boldsymbol{a}))$ and a strategy $\mathcal{S}_2$ for $\mathcal{G}(\varphi_2(\boldsymbol{a}))$, so that $\pi[\![\mathcal{S}]\!] = \pi[\![\mathcal{S}_1]\!] \cdot \pi[\![\mathcal{S}_2]\!]$. Since multiplication is continuous and commutes with (arbitrary) suprema,

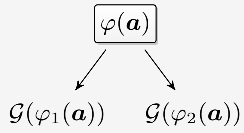

$$\begin{aligned}
\pi[\![\varphi(\boldsymbol{a})]\!] &= \pi[\![\varphi_1(\boldsymbol{a})]\!] \cdot \pi[\![\varphi_1(\boldsymbol{a})]\!] \\
&\overset{\text{(IH)}}{=} \bigsqcup\{\pi[\![\mathcal{S}_1]\!] \mid \mathcal{S}_1 \in \mathrm{WS}_{\mathcal{G}(\varphi_1)}(\varphi_1(\boldsymbol{a}))\} \cdot \bigsqcup\{\pi[\![\mathcal{S}_2]\!] \mid \mathcal{S}_2 \in \mathrm{WS}_{\mathcal{G}(\varphi_2)}(\varphi_2(\boldsymbol{a}))\} \\
&= \bigsqcup\{\pi[\![\mathcal{S}_1]\!] \cdot \pi[\![\mathcal{S}_2]\!] \mid \mathcal{S}_i \in \mathrm{WS}_{\mathcal{G}(\varphi_i)}(\varphi_i(\boldsymbol{a})) \text{ for } i \in \{1,2\}\} \\
&= \bigsqcup\{\pi[\![\mathcal{S}]\!] \mid \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))\}.
\end{aligned}$$

- $\varphi(\boldsymbol{a}) = \exists y\, \vartheta(\boldsymbol{a}, y)$ and $\varphi(\boldsymbol{a}) = \forall y\, \vartheta(\boldsymbol{a}, y)$. These are interpreted as finite sums and products (recall that A is finite) and the game consists of $|A|$ successors of $\varphi(\boldsymbol{a})$. The argument is completely analogous to the cases for $\vee$ and $\wedge$, respectively.

- $\varphi = [\mathbf{fp}\, R\, \boldsymbol{x}.\, \vartheta(R, \boldsymbol{x})](\boldsymbol{y})$. Recall that ϑ induces the operator $F^{\vartheta}_{\pi} \colon g \mapsto g'$ on functions $g, g' \colon A^r \to \mathbb{S}^\infty(X)$ (where r is the arity of R), such that $g'(\boldsymbol{a}) = \pi_{[R \mapsto g]}[\![\vartheta(\boldsymbol{a})]\!]$. Let $(g_n)_{n<\omega}$ be the fixed-point iteration induced by φ. If φ is a **lfp**-formula, then $g_0 = \mathbf{0}$ and $\pi[\![\varphi(\boldsymbol{a})]\!] = \bigsqcup_n g_n(\boldsymbol{a})$, and if φ is a **gfp**-formula then $g_0 = \mathbf{1}$ and $\pi[\![\varphi(\boldsymbol{a})]\!] = \bigsqcap_n g_n(\boldsymbol{a})$.

  Claim 1: We first show, by induction on $n \geq 1$, that for all $\boldsymbol{a} \in A^r$,

$$g_n(\boldsymbol{a}) = \bigsqcup\{\pi_{[R \mapsto g_0]}[\![\mathcal{S}\|_n]\!] \mid \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))\}.$$

  To this end, consider the strategies from $\varphi(\boldsymbol{a})$ in $\mathcal{G}(\varphi)$. Notice that R is the outermost fixed-point variable which is assigned the smallest priority in $\mathcal{G}(\varphi)$ (if other fixed-point variables are used outside of $\varphi(\boldsymbol{a})$, then these are part of the signature of φ and are treated as terminal positions in $\mathcal{G}(\varphi)$). A strategy $\mathcal{S}$ from $\varphi(\boldsymbol{a})$ first moves to $\vartheta(\boldsymbol{a})$. From there, the game proceeds in the same way as the game $\mathcal{G}(\vartheta)$, until an R-node is reached. In $\mathcal{G}(\vartheta)$, any R-node is a terminal position, but in $\mathcal{G}(\varphi)$ each play reaching a position $R\boldsymbol{b}$ must continue by $R\boldsymbol{b} \to \vartheta(\boldsymbol{b})$ and must then follow some winning strategy $\mathcal{S}_{\boldsymbol{b}}$ from $\vartheta(\boldsymbol{b})$ in $\mathcal{G}(\varphi)$, as illustrated below.

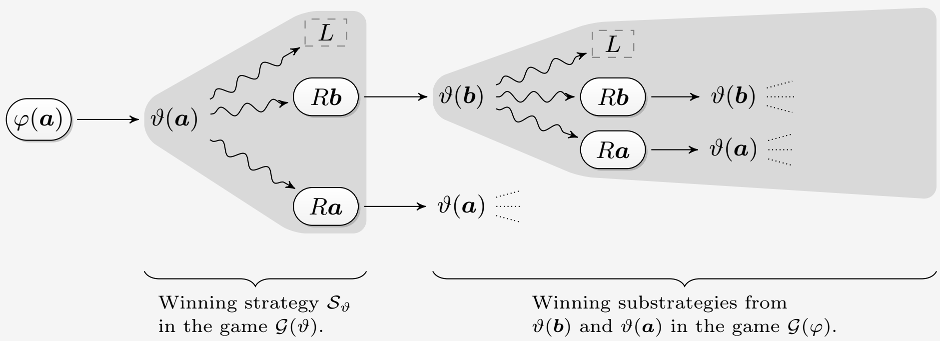

  Notice that $\mathcal{S}$ visits exactly one R-node before reaching the substrategies $\mathcal{S}_{\boldsymbol{b}}$, so that the $n+1$-truncation of $\mathcal{S}$ consists of the n-truncation of these substrategies. We write $R(\mathcal{S}_\vartheta)$ for these first R-nodes, i.e., $R(\mathcal{S}_\vartheta)$ is the set of nodes $\rho \in \mathcal{S}_\vartheta$ that are R-leaves.

We can now prove the claim. For $n = 1$, the truncation $\mathcal{S}\|_1$ consists only of $\mathcal{S}_\vartheta$ and the claim follows by induction on ϑ (for the interpretation $\pi_{[R\mapsto\mathbf{0}]}$ or $\pi_{[R\mapsto\mathbf{1}]}$). For the induction step, we also apply the induction hypothesis on ϑ:

$$g_{n+1}(\boldsymbol{a}) = \pi_{[R\mapsto g_n]}[\![\vartheta(\boldsymbol{a})]\!] = \bigsqcup\{\pi_{[R\mapsto g_n]}[\![\mathcal{S}_\vartheta]\!] \mid \mathcal{S}_\vartheta \in \mathrm{WS}_{\mathcal{G}(\vartheta)}(\vartheta(\boldsymbol{a}))\}.$$

Now fix a strategy $\mathcal{S}_\vartheta$. Using the induction hypothesis for g_n and strong distributivity of $\mathbb{S}^\infty(X)$ (recall that summation is the same as supremum), we get:

$$\begin{aligned}
\pi_{[R\mapsto g_n]}[\![\mathcal{S}_\vartheta]\!] &= \widehat{\prod_{\substack{\rho\in\mathrm{Plays}(\mathcal{S}_\vartheta),\\ \rho\notin R(\mathcal{S}_\vartheta)}}} \pi[\![\rho]\!] \cdot \widehat{\prod_{\substack{\rho\in\mathrm{Plays}(\mathcal{S}_\vartheta)\\ \text{ending in } R\boldsymbol{b}}}} \bigsqcup\{\pi_{[R\mapsto g_0]}[\![\mathcal{S}\|_n]\!] \mid \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\vartheta(\boldsymbol{b}))\} \\
&= \bigsqcup_{f\in\mathcal{F}} \Bigg(\widehat{\prod_{\substack{\rho\in\mathrm{Plays}(\mathcal{S}_\vartheta),\\ \rho\notin R(\mathcal{S}_\vartheta)}}} \pi[\![\rho]\!] \cdot \widehat{\prod_{\substack{\rho\in\mathrm{Plays}(\mathcal{S}_\vartheta)\\ \text{ending in } R\boldsymbol{b}}}} \pi_{[R\mapsto g_0]}[\![f(\rho)\|_n]\!] \Bigg) \\
&= \bigsqcup\{\pi_{[R\mapsto g_0]}[\![\mathcal{S}\|_{n+1}]\!] \mid \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a})) \text{ has prefix } \mathcal{S}_\vartheta\}.
\end{aligned}$$

Here, $\mathcal{F}$ is the set of choice functions f such that for each $\rho \in R(\mathcal{S}_\vartheta)$ ending in $R\boldsymbol{b}$ for some $\boldsymbol{b} \in A^r$, we have $f(\rho) \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\vartheta(\boldsymbol{b}))$. Clearly, each $f \in \mathcal{F}$ induces a unique winning strategy from $\varphi(\boldsymbol{a})$ that first uses $\mathcal{S}_\vartheta$ and then continues from the R-leaves according to f, and conversely each winning strategy with prefix $\mathcal{S}_\vartheta$ induces an $f \in \mathcal{F}$. The claim then follows by taking the supremum over all prefix strategies $\mathcal{S}_\vartheta$.

Claim 2: We are now ready to prove that, for all $\boldsymbol{a} \in A^r$,

$$\pi[\![\varphi(\boldsymbol{a})]\!] = \bigsqcup\{\pi[\![\mathcal{S}]\!] \mid \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))\}.$$

For $\varphi = [\mathbf{lfp}\, R\boldsymbol{x}.\, \vartheta(R,\boldsymbol{x})](\boldsymbol{y})$, this follows by Lemma 6.2.7:

$$\begin{aligned}
\pi[\![\varphi(\boldsymbol{a})]\!] = \bigsqcup_{n<\omega} g_n(\boldsymbol{a}) &= \bigsqcup_{n<\omega} \bigsqcup\{\pi_{[R\mapsto\mathbf{0}]}[\![\mathcal{S}\|_n]\!] \mid \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))\} \\
&= \bigsqcup\Big\{ \bigsqcup_{n<\omega} \pi_{[R\mapsto\mathbf{0}]}[\![\mathcal{S}\|_n]\!] \;\Big|\; \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))\Big\} \\
&\overset{(6.2.7)}{=} \bigsqcup\{\pi[\![\mathcal{S}]\!] \mid \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))\}.
\end{aligned}$$

For $\varphi = [\mathbf{gfp}\, R\boldsymbol{x}.\, \vartheta(R,\boldsymbol{x})](\boldsymbol{y})$, we only obtain one direction from Lemma 6.2.7:

$$\begin{aligned}
\pi[\![\varphi(\boldsymbol{a})]\!] = \bigsqcap_{n<\omega} g_n(\boldsymbol{a}) &= \bigsqcap_{n<\omega} \bigsqcup\{\pi_{[R\mapsto\mathbf{1}]}[\![\mathcal{S}\|_n]\!] \mid \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))\} \\
&\geq \bigsqcup\Big\{ \bigsqcap_{n<\omega} \pi_{[R\mapsto\mathbf{1}]}[\![\mathcal{S}\|_n]\!] \;\Big|\; \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))\Big\} \\
&\overset{(6.2.7)}{=} \bigsqcup\{\pi[\![\mathcal{S}]\!] \mid \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))\}.
\end{aligned}$$

The other direction requires the Puzzle lemma (6.2.8) and a bit of preparation. Let $X^\star = \{X_L \mid L \in \mathrm{Lit}_A(\tau)\}$ and consider the $\mathbb{S}^\infty(X^\star)$-interpretation $\pi^\star$ with $\pi^\star(L) = X_L$.

Our reasoning so far also applies to $\pi^\star$, and we have the additional property that each value $\pi^\star[\![\mathcal{S}\|_n]\!]$ is a monomial (this fixes a small inaccuracy in [DGNT19], where this property was claimed for π instead of $\pi^\star$). Let h be the fully-continuous homomorphism induced by the assignment $X_L \mapsto \pi(L)$, so that $\pi = h \circ \pi^\star$. This allows us to continue the proof for $\pi^\star$ and later transfer the result back to π.

Let $P_n := \bigsqcup\{\pi^\star_{[R\mapsto\mathbf{1}]}[\![\mathcal{S}\|_n]\!] \mid \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))\}$ and observe that this defines a decreasing ω-chain. Recall that by Proposition 3.5.8, the infimum of such a chain can be described by infima of monomial chains. Let $\mathcal{M}$ be the set of descending ω-chains $(m_n)_{n<\omega}$ of monomials with the property that $m_n \in P_n$ for all n. Then by Proposition 3.5.8,

$$\pi^\star[\![\varphi(\boldsymbol{a})]\!] = \prod_{n<\omega} P_n = \bigsqcup\Big\{ \prod_{n<\omega} m_n \,\Big|\, (m_n)_{n<\omega} \in \mathcal{M}\Big\}.$$

By construction, each monomial is of the form $m_n = \pi^\star_{[R\mapsto\mathbf{1}]}[\![\mathcal{S}\|_n]\!]$ for some strategy $\mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))$. The Puzzle lemma then yields, for each chain $(m_n)_{n<\omega}$, a strategy $\hat{\mathcal{S}}$ such that $\pi^\star[\![\hat{\mathcal{S}}]\!] \geq \prod_n m_n$. It follows that

$$\pi^\star[\![\varphi(\boldsymbol{a})]\!] \leq \bigsqcup\{\pi^\star[\![\mathcal{S}]\!] \mid \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))\}.$$

To close the proof, we transfer the result back to the original interpretation π. Since h commutes with semiring semantics (by Proposition 5.1.4), with suprema (by Corollary 3.2.6), as well as as with the infinite product in $\pi[\![\mathcal{S}]\!]$ (by Proposition 3.4.3), we get:

$$\begin{aligned}\pi[\![\varphi(\boldsymbol{a})]\!] = h(\pi^\star[\![\varphi(\boldsymbol{a})]\!]) &\leq \bigsqcup\{h(\pi^\star[\![\mathcal{S}]\!]) \mid \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))\} \\ &= \bigsqcup\{\pi[\![\mathcal{S}]\!]) \mid \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a}))\}.\end{aligned}$$ □

Generalisation

It remains to prove Theorem 6.2.3 (Sum of strategies) for absorptive, fully-continuous semirings other than $\mathbb{S}^\infty(X)$. This follows by repeating the argument we just did in the case $\varphi = [\mathbf{gfp}\, R\boldsymbol{x}.\, \vartheta(R,\boldsymbol{x})](\boldsymbol{y})$. Let K be an absorptive, fully-continuous semiring, and let π be a K-interpretation (over A and τ). Define $\pi^\star$ and h as above. Using the properties of h and Lemma 6.2.11, we get:

$$\begin{aligned}\pi[\![\varphi(\boldsymbol{a})]\!] = h(\pi^\star[\![\varphi(\boldsymbol{a})]\!]) &= \bigsqcup\{h(\pi^\star[\![\mathcal{S}]\!]) \mid \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a})) \\ &= \bigsqcup\{\pi[\![\mathcal{S}]\!] \mid \mathcal{S} \in \mathrm{WS}_{\mathcal{G}(\varphi)}(\varphi(\boldsymbol{a})).\end{aligned}$$ □

6.3 A Case Study: Semiring Provenance for Büchi Games

The Sum-of-strategies theorem for LFP shows that we can use model-checking games to better understand semiring semantics. We now take the opposite direction: we demonstrate the usefulness of the semiring provenance framework for the analysis of infinite games, by using semiring semantics of LFP to perform provenance analysis of Büchi games. The case of Büchi games is of particular interest, as these are the simplest infinite games with a winning condition that requires a fixed-point alternation. Indeed, given that the objective of Player 0 is to ensure that the play hits the target set F infinitely often, we may informally describe their winning region as the *largest* set Y of positions from which they can enforce a (further) visit to $Y \cap F$ after at least one move. On the other hand, the set of positions from which Player 0 can enforce a visit to some target set is the *smallest* set of positions that either are already in the target set, or from which Player 0 can enforce the play to come closer to it. Thus, the winning region of Player 0 can be described as a greatest fixed point, inside of which there is a least fixed point, and it is well-known that this fixed-point alternation cannot be avoided (see, e.g., [BW18]).

In order to define a semiring value for (a position x in) a Büchi game, we take the fixed-point formula $\mathsf{win}_0(x)$ that defines the winning region for Player 0 in Büchi games and evaluate it in a suitable semiring, most importantly in the semiring of generalised absorptive polynomials $\mathbb{S}^\infty(X)$. The resulting semiring value provides detailed information about a certain class of winning strategies – those that win with minimal effort. This class of strategies, which we call *absorption-dominant*, may be of independent interest, and we discuss how these strategies relate to positional and the more general *persistent* strategies.

We make the connection to winning strategies precise by proving a sum-of-strategies theorem for Büchi games. Besides being of theoretical interest, this result allows to study a number of interesting questions concerning the available winning strategies in a Büchi game.

- *Strategy tracking*. We can track edges of a Büchi game $\mathcal{G}$ by using the interpretation π_{strat} that maps each edge to a unique indeterminate. The resulting semiring value $\pi_{\mathsf{strat}}[\![\mathsf{win}_0(v)]\!]$ for a position v is then a polynomial whose monomials correspond precisely to the absorption-dominant winning strategies from v.

  From these monomials we can derive whether Player 0 wins from v, which edges are used by each absorption-dominant strategy, and how often they appear in the strategy tree. This information can also be used to identify and count positional strategies, and to answer questions such as: can Player 0 still win if a certain edge is removed?

- *Repairing a game*. Instead of analysing strategies in a fixed game, we may also reason about modifications or synthesis of (parts of) the game. For example, assuming Player 0 loses from v, what are minimal modifications (called repairs) to the game arena that would let Player 0 win from v?

  Such questions can be answered by evaluating $\mathsf{win}_0(v)$ in the semiring $\mathrm{PosBool}(X, \overline{X})$, where the dual-indeterminates are used to take into account also negative information (i.e., absent edges in the graph). We illustrate this in detail for the task of computing minimal repairs.

- *Cost computation*. A typical application of semiring provenance in databases is cost analysis. In our setting, this means that each move comes with a cost (a nonnegative

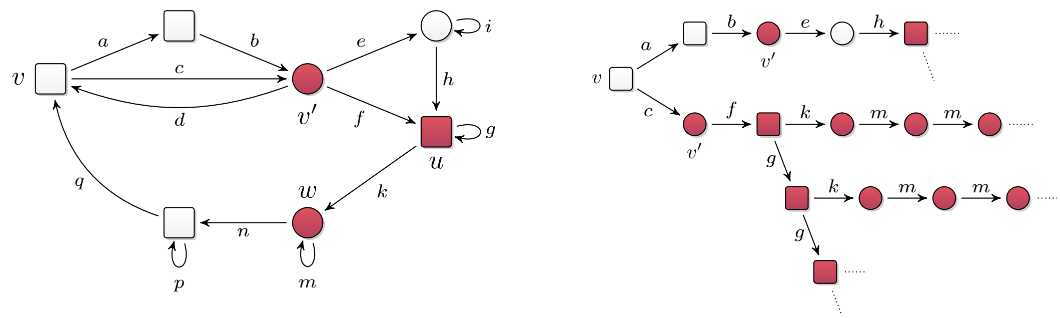

(a) Rectangular nodes belong to Pl. 1, round nodes to Pl. 0, red nodes are in F.

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

Figure 6.6: Running example of a Büchi game and a winning strategy.

real number) and we ask for the minimal cost of a winning strategy. For reachability and safety games that admit only finite plays, this cost can be obtained by using an interpretation in the tropical semiring (cf. [GT20]).

The same interpretation also induces a cost measure for strategies in Büchi games. However, we argue that this is not an intuitive notion of cost for infinite games, and we propose two alternatives. This reveals a limitation of the semiring provenance approach: we show that there is no semiring such that the evaluation of $\mathsf{win}_0(x)$ describes, without further computations, either of the two alternative cost measures.

The case study presented in this section is joint work with Erich Grädel and Niels Lücking, and has previously been published in [GLN21, GLN24] (with the exception of Theorem 6.3.14).

6.3.1 Running Example

Recall that a Büchi game has the form $\mathcal{G} = (V, V_0, V_1, E, F)$, where $F \subseteq V$ is the set of target positions. We assume that $vE \neq \emptyset$ for all positions v, so that all plays are infinite and Player 0 wins if, and only if, F is visited infinitely often during a play. We again represent strategies as trees, as defined in Definition 6.1.2. In the following, $\mathcal{G}$ always refers to a Büchi game, and we may write $v \in \mathcal{G}$ instead of $v \in V$.

Example 6.3.1. Figure 6.6a shows our running example of a Büchi game. Player 0 has essentially three different positional winning strategies from v, by either choosing edge d, or edges e, h, m or f, m. Notice that for the first strategy, we did not specify moves for all positions in V_0 as some positions cannot be reached when edge d is played; this is the main reason why we represent strategies as trees. Figure 6.6b depicts such a tree representation of a strategy. This strategy is a typical example of a winning strategy that is not positional, but still minimal if we take edge multiplicities (see Definition 6.3.2) into account.

6.3.2 Strategies with Minimal Effort

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

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

Definition 6.3.2. Given an edge $e = vw \in E$ in a Büchi game $\mathcal{G}$ and a strategy $\mathcal{S}$ in $\mathcal{G}$, we denote by $\#_e(\mathcal{S}) = |\{\rho v \in \mathcal{S} \mid \rho v \to \rho v w \text{ is an edge in } \mathcal{S}\}| \in \mathbb{N} \cup \{\infty\}$ the number of times (possibly infinite) the edge e occurs in $\mathcal{S}$. With each strategy $\mathcal{S}$ we associate its *edge profile*, the vector $\#_E(\mathcal{S}) = (\#_e(\mathcal{S}))_{e \in E}$.

Example 6.3.3. Consider the following Büchi game:

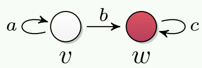

Player 0 wins by first looping n times at position v (for any fixed $n \in \mathbb{N}$) and then moving to w, corresponding to the edge profile $(n, 1, \infty)$. Clearly, looping at v is a redundant move, so we regard the strategy with $n = 0$ as the simplest one (that wins with the least effort).

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

Definition 6.3.4. Let $\mathcal{S}_1, \mathcal{S}_2$ be two strategies in a Büchi game $\mathcal{G} = (V, V_0, V_1, E, F)$. We say that $\mathcal{S}_1$ *absorbs* $\mathcal{S}_2$, denoted $\mathcal{S}_1 \succeq \mathcal{S}_2$, if $\#_e(\mathcal{S}_1) \leq \#_e(\mathcal{S}_2)$ for all edges $e \in E$. If additionally $\#_e(\mathcal{S}_1) < \#_e(\mathcal{S}_2)$ for some $e \in E$, we say that $\mathcal{S}_1$ *strictly absorbs* $\mathcal{S}_2$, denoted $\mathcal{S}_1 \succ \mathcal{S}_2$. They are *absorption-equivalent*, denoted $\mathcal{S}_1 \equiv \mathcal{S}_2$, if both $\mathcal{S}_1 \succeq \mathcal{S}_2$ and $\mathcal{S}_2 \succeq \mathcal{S}_1$.

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

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

Example 6.3.5. Consider the strategy $\mathcal{S}$ as depicted in Fig. 6.6b. It is not positional, as the choice for position v' is not unique (both e and f occur in $\mathcal{S}$). It is, however, absorption-dominant. As there are two paths to v', every strategy must either use e or f twice, or use both edges. If e (or f) is used twice, then the strategy cannot absorb $\mathcal{S}$, and one can verify

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

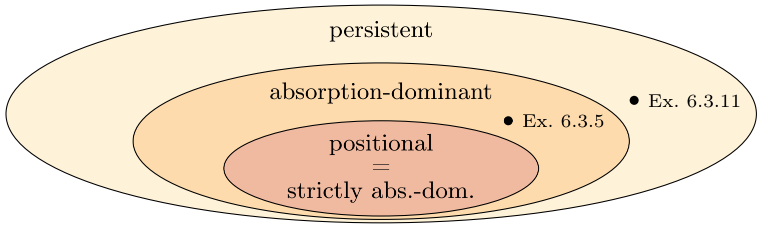

Figure 6.7: Venn diagram depicting classes of winning strategies.

that $\mathcal{S}$ absorbs all strategies using both e and f.

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

Strategies such as the one in Fig. 6.6b are not positional, but satisfy the weaker property that within each *play*, the strategy makes a unique decision for each position $v \in V_0$. This notion of strategies has been introduced as *persistent* strategies in [MT02] in the context of LTL on game graphs, and has been further studied in [Dup03]. Persistent strategies have also been called *weakly positional* in [GT20].

We say that a strategy *plays positionally* from a position $v \in V_0$ if the strategy makes a unique choice at position v (not depending on the history of the play). A strategy that plays positionally from all positions in V_0 is positional. With this notation, we now clarify the relation between the different notions of strategies; a summary is shown in Fig. 6.7. We first observe that if a strategy $\mathcal{S}$ does not play positionally from v, we can always obtain a strategy $\mathcal{S}'$ with $\mathcal{S}' \succeq \mathcal{S}$ by swapping the choices at v, which leads to:

Proposition 6.3.6. *Strictly absorption-dominant strategies coincide with positional strategies.*

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

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

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

Lemma 6.3.7. *Let $v \in \mathcal{G}$ be a position. There are only finitely many absorption-dominant strategies from v up to absorption-equivalence.*

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

Proposition 6.3.8. *Let $\mathcal{S} \in \mathrm{Strat}_{\mathcal{G}}(v)$ be absorption-dominant from v, and let $w \in V_0$ be a position. If w occurs infinitely often in $\mathcal{S}$, then $\mathcal{S}$ plays positionally from w.*

Proof. Consider the infinitely many substrategies at occurrences of w in $\mathcal{S}$. By Lemma 6.3.7, there is one such substrategy $\mathcal{S}_w$ such that infinitely many of the substrategies are absorption-equivalent to $\mathcal{S}_w$. This means that every edge occurring in $\mathcal{S}_w$ also occurs in infinitely many substrategies and hence infinitely often in the full strategy $\mathcal{S}$. Consider the subgame of $\mathcal{G}$ containing only edges occurring in $\mathcal{S}_w$. This subgame is nonterminating (i.e., each node has a successor) and hence there is a positional strategy $\mathcal{S}_{\mathrm{pos}} \in \mathrm{Strat}_{\mathcal{G}}(w)$ using only edges that occur in $\mathcal{S}_w$ and hence infinitely often in $\mathcal{S}$.

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

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

Corollary 6.3.9. *Every absorption-dominant strategy is persistent.*

Proof. Let $\mathcal{S} \in \mathrm{Strat}_{\mathcal{G}}(v)$ be absorption-dominant from v. Assume towards a contradiction that $\mathcal{S}$ is not persistent, so there is a position $w \in V_0$ and a play of the form $\rho_1 w \rho_2 w \rho_3$ such that $\mathcal{S}$ makes different decisions at w, say $\mathcal{S}(\rho_1 w) = w_1$ and $\mathcal{S}(\rho_1 w \rho_2 w) = w_2$ with $w_1 \neq w_2$. By Proposition 6.3.8, w can only occur finitely often in $\mathcal{S}$. Hence the edge ww_1 also occurs finitely often, say n times. Let $\mathcal{S}'_w$ be the substrategy of $\mathcal{S}$ from $\rho_1 w \rho_2 w$. Now consider the strategy $\mathcal{S}' \in \mathrm{Strat}_{\mathcal{G}}(v)$ that behaves like $\mathcal{S}$, but switches to $\mathcal{S}'_w$ at $\rho_1 w$. By construction, $\mathcal{S}'$ uses each edge at most as often as $\mathcal{S}$. Moreover, one occurrence of the edge ww_1 is removed, so this edge occurs at most $n-1$ times in $\mathcal{S}'$. Hence $\mathcal{S}' \succ \mathcal{S}$, contradicting the absorption-dominance of $\mathcal{S}$. □

Remark 6.3.10. Notice that the previous statements (6.3.6 – 6.3.9) talk about arbitrary strategies, which are not necessarily winning. They are thus independent of the winning condition and apply to all games with finite game arena.

Example 6.3.11. For strictness, consider the game on the left (a modified part of Fig. 6.6a) and the winning strategies induced by the edge profiles shown on the right:

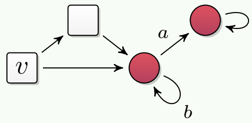

$$\#_{a,b}(\mathcal{S}_1) = (2, 0),$$
$$\#_{a,b}(\mathcal{S}_2) = (0, \infty),$$
$$\#_{a,b}(\mathcal{S}_3) = (1, \infty).$$

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

As a consequence of Corollary 6.3.9, all moves after the first repeated position are determined by persistence. We can thus represent absorption-dominant strategies in a compact way and strengthen Lemma 6.3.7 as follows.

Corollary 6.3.12. *Let $\mathcal{G}$ be a game with $n = |V|$ positions. Every strategy $\mathcal{S} \in \mathrm{Strat}_{\mathcal{G}}(v)$ that is absorption-dominant from v can be uniquely represented by a subtree of the tree unraveling of height at most n. In particular, the number of absorption-dominant strategies is finite.*

As a last step in our analysis of absorption-dominant strategies, we establish the connection to the Büchi winning condition. That is, we show that whenever Player 0 has a winning strategy $\mathcal{S}$, this strategy can be simplified to an absorption-dominant strategy $\mathcal{S}'$. We prove this statement by first providing an alternative characterisation of absorption-dominant strategies (which may be of independent interest).

Proposition 6.3.13. *Let $\mathcal{S} \in \mathrm{Strat}_{\mathcal{G}}(v)$. Then $\mathcal{S}$ is absorption-dominant from v if, and only if, the following conditions holds for each position $w \in V_0$:*

- *if w occurs infinitely often in $\mathcal{S}$, then $\mathcal{S}$ plays positionally from w,*
- *if w occurs only finitely often in $\mathcal{S}$, then w occurs at most once in each play consistent with $\mathcal{S}$.*

Proof. First let $\mathcal{S}$ be absorption-dominant from v. The first condition holds by Proposition 6.3.8. The second condition follows from Corollary 6.3.9: since $\mathcal{S}$ is persistent, a second occurrence of w in a play would lead to a repeating cycle and w would occur infinitely often.

For the other direction, suppose that $\mathcal{S}$ satisfies the two conditions but there is a strategy $\mathcal{S}' \in \mathrm{Strat}_{\mathcal{G}}(v)$ with $\mathcal{S}' \succ \mathcal{S}$. We recursively compare $\mathcal{S}$ and $\mathcal{S}'$ as follows. First assume v occurs infinitely often in $\mathcal{S}$. Then $\mathcal{S}$ is positional by the first condition and Proposition 6.3.6 contradicts the assumption $\mathcal{S}' \succ \mathcal{S}$. Otherwise v occurs only finitely often in $\mathcal{S}$. If $v \in V_1$, then there must be some successor $w \in vE$ so that the substrategies $\mathcal{S}'_{vw}$ and $\mathcal{S}_{vw}$ from this successor satisfy $\mathcal{S}'_{vw} \succ \mathcal{S}_{vw}$ and we continue recursively with these strategies. If $v \in V_0$, then this is the only occurrence of v in $\mathcal{S}$. Let $v \to w$ be the move chosen by $\mathcal{S}$. If $\mathcal{S}'$ would choose a different successor $w' \neq w$, then the edge vw' would occur in $\mathcal{S}'$ but not in $\mathcal{S}$ (since v does not occur anywhere else in $\mathcal{S}$). This would contradict $\mathcal{S}' \succ \mathcal{S}$, so $\mathcal{S}'$ must choose the same successor w. We again continue with the respective substrategies $\mathcal{S}'_{vw} \succ \mathcal{S}_{vw}$ from this successor.

Since the number of positions is finite, after finitely many steps a position v occurring infinitely often must be reached, so we always obtain a contradiction. Hence no such strategy $\mathcal{S}'$ can exist and $\mathcal{S}$ is absorption-dominant. □

Theorem 6.3.14. *For each winning strategy $\mathcal{S} \in \mathrm{WinStrat}_{\mathcal{G}}(v)$, there is a winning strategy $\mathcal{S}' \in \mathrm{WinStrat}_{\mathcal{G}}(v)$ that is absorption-dominant from v and satisfies $\mathcal{S}' \succeq \mathcal{S}$.*

Proof. Let $V_\infty \subseteq V$ be the set of positions occurring infinitely often in $\mathcal{S}$, and V_{fin} the positions occurring only finitely often. Let further $\mathcal{S}_{\text{fin}}$ be the finite prefix of $\mathcal{S}$ containing all occurrences of V_{fin}.

Consider the subgame $\mathcal{G}'$ induced by the positions V_∞ (notice that $\mathcal{G}'$ is nonterminating and no moves of Player 1 are removed). Each occurrence of $w \in V_\infty$ that happens in $\mathcal{S}$ after the prefix $\mathcal{S}_{\text{fin}}$ induces a winning strategy from w in $\mathcal{G}'$, so Player 0 wins from all $w \in \mathcal{G}'$. By positional determinacy, there are corresponding positional winning strategies $(\mathcal{P}_w)_{w \in V_\infty}$ so that $\mathcal{P}_w \in \mathrm{WinStrat}_{\mathcal{G}'}(w)$, and hence also $\mathcal{P}_w \in \mathrm{WinStrat}_{\mathcal{G}}(w)$ in the original game. We can further assume that these strategies are consistent with each other (i.e., for each position $w \in V_\infty \cap V_0$ they all choose the same successor w').

We now recursively construct a strategy $\mathcal{S}'$ from $\mathcal{S}$ that satisfies the conditions in Proposition 6.3.13 and is thus absorption-dominant. Recall that $\mathcal{S}$ starts at position v.

(1) If $v \in V_\infty$, we set $\mathcal{S}' \coloneqq \mathcal{P}_v$ for the positional strategy $\mathcal{P}_v \in \mathrm{WinStrat}_{\mathcal{G}'}(v)$.

(2) If $v \in V_{\text{fin}} \cap V_1$ with substrategies $\mathcal{S}_1, \ldots, \mathcal{S}_k$ for the immediate successors of v, then $\mathcal{S}'$ consists of the root v followed by the recursively obtained strategies $\mathcal{S}'_1, \ldots, \mathcal{S}'_k$.

(3) If $v \in V_{\text{fin}} \cap V_0$ and v occurs only once in $\mathcal{S}$, let w be the chosen successor and $\mathcal{S}_w$ the corresponding substrategy (in which v does not occur). Then $\mathcal{S}'$ consists of the root v followed by the recursively obtained substrategy $\mathcal{S}'_w$.

(4) If $v \in V_{\text{fin}} \cap V_0$ and v occurs more than once (but finitely often) in $\mathcal{S}$, we pick a last occurrence $\rho \in \mathcal{S}$ of v. Let $\mathcal{S}_\rho$ be the substrategy from ρ, obtain $\mathcal{S}'_\rho$ by case (3) and set $\mathcal{S}' \coloneqq \mathcal{S}'_\rho$.

First observe that this construction always terminates, as the prefix $\mathcal{S}_{\text{fin}}$ is finite and hence case (1) applies at finite recursion depth. The resulting strategy is composed of winning strategies, and is thus also winning (for case (4), we use that the winning condition depends only on the suffix of a play). We further have $\mathcal{S}' \succeq \mathcal{S}$, as the positional strategies in (1) only contain edges that occur infinitely often in $\mathcal{S}$, and steps (2)–(4) do not add any new moves. It remains to show that the conditions for Proposition 6.3.13 hold.

If a position $w \in V_0$ occurs in the positional substrategies $\mathcal{P}_v$, then $w \in V_\infty$ and hence w can only occur in the positional part of $\mathcal{S}'$. Since positional strategies are absorption-dominant, the conditions hold for w. Otherwise we must have $w \in V_{\text{fin}}$ and w can only occur finitely often in $\mathcal{S}'$. Then cases (3) and (4) ensure that an occurrence of w cannot be followed by a second one. Hence $\mathcal{S}'$ is indeed absorption-dominant. □

Interpretation & application	π_{strat} *strategy tracking*	π_{repair} *reverse analysis of moves*	π_{target} *target synthesis*
Semiring	$\mathbb{S}^\infty(X)$	$\mathrm{PosBool}(X, \overline{X})$	$\mathrm{PosBool}(X, \overline{X})$
$\pi(Evw)$	$X_{vw}/0$	X_{vw} (or $1/0$)	$1/0$
$\pi(\neg Evw)$	$1/0$	$\overline{X_{vw}}$ (or $1/0$)	$1/0$
$\pi(Fv)$	$1/0$	$1/0$	X_v
$\pi(\neg Fv)$	$1/0$	$1/0$	$\overline{X_v}$
other literals	$1/0$	$1/0$	$1/0$

Figure 6.8: Semiring interpretations used for the analysis of Büchi games (the notation $a/0$ indicates the value a if the literal is true, and 0 if it is false in the given game).

6.3.3 Computing Strategies with Semiring Semantics

We now describe how semiring semantics for LFP (see Section 5.2) can be used to define semiring values for positions in a Büchi game. Formally, we view a Büchi game $\mathcal{G} = (V, V_0, V_1, E, F)$ as a τ-structure with universe V and relational signature $\tau = \{E, F, V_0, V_1\}$. The set of instantiated literals $\mathrm{Lit}_V(\tau)$ then contains, for instance, Ev_1v_2 and $\neg Fv_1$, where $v_1, v_2 \in V$. A K-interpretation (into an appropriate semiring K) can thus provide annotations for the moves and the target set (and also for V_0 and V_1).

We define the values of positions in $\mathcal{G}$ by evaluating the standard LFP-formula for the winning region of Player 0 in semiring semantics, and we show that in absorptive, fully-continuous semirings these values can be described by a sum-of-strategies theorem.

The Semiring Interpretation

In order to analyse the moves in winning strategies, we label edges with indeterminates X (cf. Fig. 6.6a) and use an $\mathbb{S}^\infty(X)$-interpretation π_{strat} that maps each edge literal Evw to the corresponding label. We assume the game arena to be fixed (for now) and do not wish to track information about the target set F or the active player at a certain node, hence we simply map all other literals, such as Fv, V_0v and $\neg Evw$, to 0 or 1, depending on whether they are true or false in the fixed game. The resulting interpretation is model-defining and hence behaves very similar to the original game, except that we remember which edges are used in the evaluation of a formula.

Definition 6.3.15. Let $\mathcal{G} = (V, V_0, V_1, E, F)$ be a Büchi game and let $X = \{X_{vw} \mid vw \in E\}$ be a set of indeterminates for the edges. We define the $\mathbb{S}^\infty(X)$-interpretation $\pi_{\mathsf{strat}} \colon \mathrm{Lit}_V(\tau) \to \mathbb{S}^\infty(X)$ as follows (depending on $\mathcal{G}$):

$$\pi_{\mathsf{strat}}(Evw) = X_{vw} \text{ for all edges } vw \in E,$$

$$\pi_{\mathsf{strat}}(L) = \begin{cases} 1, & \text{if } \mathcal{G} \models L, \\ 0, & \text{if } \mathcal{G} \not\models L, \end{cases} \quad \text{for all other literals } L \in \mathrm{Lit}_V(\tau).$$

For the applications in Section 6.3.5, we also consider interpretations that track negative edge literals or the target set F. See Fig. 6.8 for an overview.

The Formula

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

$$\mathsf{win}_0(x) \coloneqq \big[\mathbf{gfp}\, Yy.\ [\mathbf{lfp}\, Zz.\ \varphi(Y,Z,z)](y)\big](x),$$

where

$$\begin{aligned}\varphi(Y,Z,z) \coloneqq &\Big(Fz\ \wedge\ ((V_0z \wedge \exists u(Ezu \wedge Yu)) \vee (V_1z \wedge \forall u(Ezu \to Yu)))\Big)\\ &\vee \Big(\neg Fz\ \wedge\ ((V_0z \wedge \exists u(Ezu \wedge Zu)) \vee (V_1z \wedge \forall u(Ezu \to Zu)))\Big).\end{aligned}$$

Given a K-interpretation π for a Büchi game $\mathcal{G} = (V, V_0, V_1, E, F)$, semiring semantics of the above formula induces[2] the following fixed-point computation. To simplify the presentation, we introduce two families of variables, $\mathbf{Y} = (Y_v)_{v\in V}$ and $\mathbf{Z} = (Z_v)_{v\in V}$ that take values in K. We can then express the resulting semiring valuation as $\pi[\![\mathsf{win}_0(v)]\!] = Y_v^*$ where $\mathbf{Y}^* = (Y_v^*)_{v\in V}$ is the *greatest* solution to the equation system

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

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

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

The interpretation π_{strat} tracks only *moves* of winning strategies, and hence maps most of the literals to 0 or 1. We can thus simplify the equations depending no the position v:

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

A good way to think about the least and greatest solutions is the fixed-point iteration. In this context, this means that for a given tuple $\mathbf{Y}$, we start by setting $Z_v = 0$ and then repeatedly apply the above equations (i.e., the induced operator $F^{\varphi}_{\pi_{\mathsf{strat}}}$) until a fixed point is reached, which then becomes the new value for $\mathbf{Y}$. Notice that, although the game arena is finite, this iteration can be infinite in case of an infinite semiring K.

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

Example 6.3.16. Recall the simple game from Example 6.3.3 (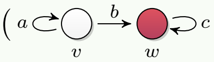).

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

$$\begin{array}{lccccccccccccc}
\mathbf{Y}: & \begin{pmatrix}1\\1\end{pmatrix} & & & & & \begin{pmatrix}bc\\c\end{pmatrix} & & & & & \begin{pmatrix}bc^2\\c^2\end{pmatrix} & & \begin{pmatrix}bc^n\\c^n\end{pmatrix} \\
& \searrow & & & & \nearrow & & \searrow & & & \nearrow & \searrow & \nearrow & \searrow \\
\mathbf{Z}: & & \begin{pmatrix}0\\0\end{pmatrix} \overset{F^\varphi}{\longmapsto} & \begin{pmatrix}0\\c\end{pmatrix} \overset{F^\varphi}{\longmapsto} & \begin{pmatrix}bc\\c\end{pmatrix} & & & & \begin{pmatrix}0\\0\end{pmatrix} \overset{F^\varphi}{\longmapsto} \begin{pmatrix}0\\c^2\end{pmatrix} \overset{F^\varphi}{\longmapsto} & \begin{pmatrix}bc^2\\c^2\end{pmatrix} & & \cdots & & \cdots
\end{array}$$

We obtain the overall result $\pi_{\mathsf{strat}}[\![\mathsf{win}_0(v)]\!] = Y_v^* = \prod_n bc^n = bc^\infty$ corresponding to the unique absorption-dominant strategy using edge b once and c infinitely often (cf. Example 6.3.3).

Connection to Strategies

By mapping edges to semiring values, we can track edges through the fixed-point computation. In Example 6.3.16, the resulting semiring value revealed how often each edge is used in the unique absorption-dominant winning strategy. We now generalize this observation. For simplicity, we only consider K-interpretations π that are *edge tracking* for a given game $\mathcal{G}$. That is, they may assign arbitrary values to positive edge literals Evw, but all other literals are mapped to 0 or 1 in accordance with $\mathcal{G}$. To make the connection to strategies explicit, we first define semiring values for strategies based on the appearance of edges.

Definition 6.3.17. Let $\mathcal{S}$ be a strategy in a Büchi game $\mathcal{G} = (V, V_0, V_1, E, F)$. Let K be an absorptive, fully-continuous semiring and π an edge-tracking K-interpretation on $\mathcal{G}$. The *K-value* of $\mathcal{S}$ is the product of the values for all edges appearing in $\mathcal{S}$. Formally,

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

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

The semiring value of $\mathsf{win}_0(x)$ can then be expressed as the sum over the values of all winning strategies. A direct proof is not completely straightforward, as fixed-point iterations and strategy trees can both be infinite (even if $\mathcal{G}$ is finite). Instead, we make use of the Sum-of-strategies theorem for LFP (Theorem 6.2.3) by comparing $\mathcal{G}$ with the model-checking game of $\mathsf{win}_0(x)$.

Theorem 6.3.18 (Sum of Strategies)**.** *Let $\mathcal{G}$ be a Büchi game and v a position in $\mathcal{G}$. Let K be an absorptive, fully-continuous semiring and π an edge-tracking K-interpretation. Then,*

$$\pi[\![\mathsf{win}_0(v)]\!] = \bigsqcup \big\{ \pi[\![\mathcal{S}]\!] \;\big|\; \mathcal{S} \in \mathrm{WinStrat}_{\mathcal{G}}(v) \textit{ is absorption-dominant from } v \big\}.$$

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

Example 6.3.19. For the edge-tracking interpretation π_{strat} induced by the edge labels in Fig. 6.6a, we obtain

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

There are four monomials, corresponding to four equivalence classes of absorption-dominant strategies. Each monomial reveals the edges that appear in the corresponding strategies, so we see that the first three monomials belong to positional (and hence uniquely defined) strategies. The last monomial belongs to the non-positional strategy shown in Fig. 6.6b (and its switched version, see Example 6.3.5). The values of all other strategies are strictly absorbed by one of these monomials.

6.3.4 Proof of the Sum-of-Strategies Theorem

To prove that Theorem 6.3.18 follows from Theorem 6.2.3 (Sum-of-strategies for LFP), we show that the model-checking game for $\mathsf{win}_0(v)$ on a Büchi game $\mathcal{G}$ has the same structure as the game arena $\mathcal{G}$ itself.

The Model-Checking Game

We use the construction of model-checking games for LFP as defined in Section 6.2.1. For a Büchi game $\mathcal{G}$ and a position v, let $\mathsf{MC}(\mathcal{G}, v)$ denote the model-checking game for $\mathsf{win}_0(x)$ on $\mathcal{G}$, with explicit starting position $\mathsf{win}_0(v)$. Recall the definition of $\mathsf{win}_0(x)$ from Section 6.3.3 (split into subformulae to ease referencing and with implication already rewritten):

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

where

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

$$\vartheta_1(z) := \big((V_0z \wedge \exists u(Ezu \wedge Yu)) \vee (V_1z \wedge \forall u(\neg Ezu \vee (Ezu \wedge Yu)))\big),$$
$$\vartheta_2(z) := \big((V_0z \wedge \exists u(Ezu \wedge Zu)) \vee (V_1z \wedge \forall u(\neg Ezu \vee (Ezu \wedge Zu)))\big).$$

This formula only has a single fixed-point alternation, and hence the parity game $\mathsf{MC}(\mathcal{G}, v)$ only has two priorities. We can thus equivalently view $\mathsf{MC}(\mathcal{G}, v)$ as a Büchi game (with Verifier as Player 0 and positions with even priority as target set F).

A depiction of the complete model-checking game $\mathsf{MC}(\mathcal{G}, v)$ is shown in Fig. 6.9 (with some unavoidable omissions due to space reasons). To simplify the presentation, the positions for fixed-point formulae (e.g., $[\mathbf{lfp}\,Z\,z.\ \varphi](v)$) and the corresponding fixed-point variables (e.g., Zv) are identified. This does not affect the gameplay, as any play has to take the unique move $[\mathbf{lfp}\,Z\,z.\ \varphi](v) \to Zv$ anyway.

We denote strategies in $\mathsf{MC}(\mathcal{G}, v)$ by the letter $\mathcal{M}$, to avoid confusion with the strategies in $\mathcal{G}$ (which we denote by $\mathcal{S}$). Recall that the semiring value $\pi[\![\mathcal{M}]\!]$ of a strategy in $\mathsf{MC}(\mathcal{G}, v)$ is defined as the product over the values π assigns to the literals occurring (as leafs) in $\mathcal{M}$ (cf. Definition 6.2.2). Theorem 6.2.3 then states that

$$\pi[\![\mathsf{win}_0(v)]\!] = \bigsqcup\{\pi[\![\mathcal{M}]\!] \mid \mathcal{M} \text{ is a winning strategy in } \mathsf{MC}(\mathcal{G}, v)\}. \qquad (*)$$

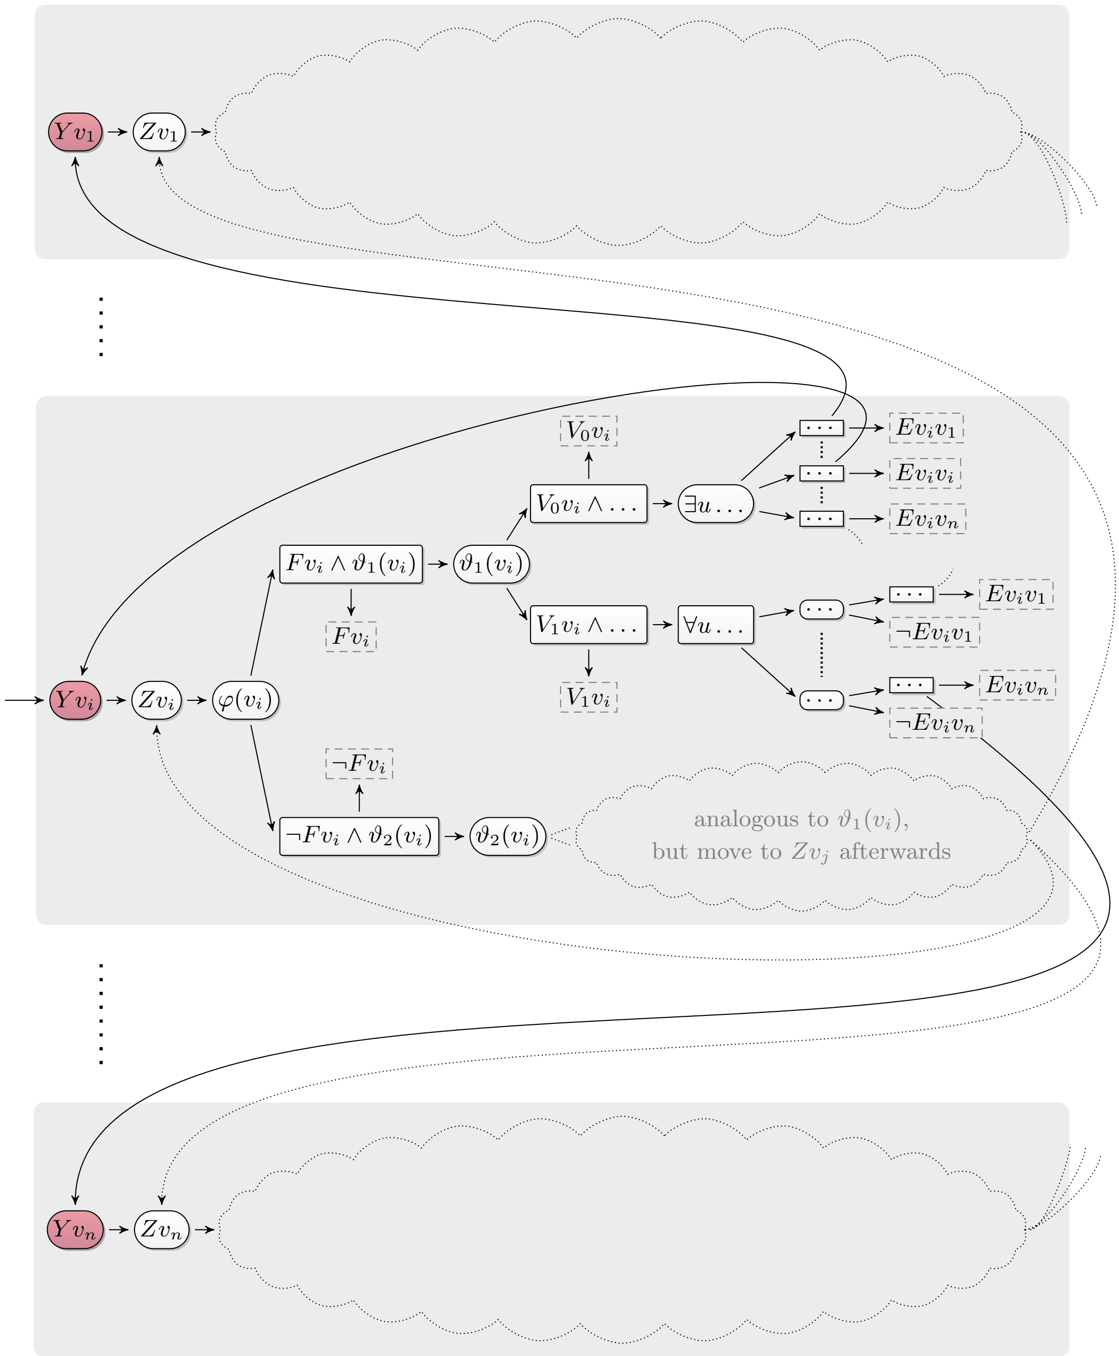

Figure 6.9: Illustration of the model-checking game $\mathsf{MC}(\mathcal{G}, v_i)$ for a Büchi game with positions $V = \{v_1, \dots, v_n\}$. Rounded nodes belong to Verifier, rectangular nodes to Falsifier. Nodes with dashed border are terminal positions representing literals, red nodes are target positions. Dotted edges and clouds indicate omitted parts.

In order to apply this to edge-tracking interpretations in Büchi games, we show that the winning strategies $\mathcal{M}$ in $\mathsf{MC}(\mathcal{G}, v)$ correspond to winning strategies $\mathcal{S}$ in $\mathcal{G}$ with the same semiring value.

From the Model-Checking Game to the Büchi Game

We describe a series of simplifications to $\mathsf{MC}(\mathcal{G}, v)$ without affecting the sum of strategy values in $(*)$. Let π be an edge-tracking K-interpretation for $\mathcal{G}$, so most of the literals are mapped to 0 or 1. We can then remove the corresponding terminal positions in $\mathsf{MC}(\mathcal{G}, v)$ and in some cases also their predecessors. For instance, consider a position $\varphi(\boldsymbol{a})$ in $\mathsf{MC}(\mathcal{G}, v)$ from which Falsifier can move to a literal L with $\pi(L) = 0$. Then every strategy $\mathcal{M}$ that visits $\varphi(\boldsymbol{a})$ must also visit L, thus having value $\pi[\![\mathcal{M}]\!] = 0$, so we can ignore this strategy for the sum in $(*)$. Hence, replacing the position $\varphi(\boldsymbol{a})$ by its successor L does not change the sum. On the other hand, if $\pi(L) = 1$, then visiting L does not affect the value of $\mathcal{M}$ and hence we can remove L. Similar reasoning applies to positions of Verifier. Moreover, we can always skip over non-target positions with a unique successor, as they neither affect gameplay nor the values of strategies.

With these insights, we can simplify the model-checking game in Fig. 6.9 quite a bit. If, say, $v_i \in F$ and $v_i \in V_0$, then the central part of the picture simplifies to:

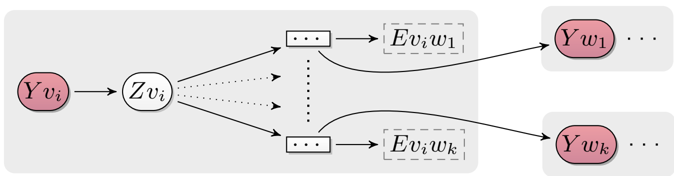

Here, $v_iE = \{w_1, \dots, w_k\}$ are the successors of v_i in $\mathcal{G}$, so Verifier's moves from Zv_i are only those corresponding to actual edges of $\mathcal{G}$. The other situations are similar, here is the case $v_i \notin F$, $v_i \in V_1$:

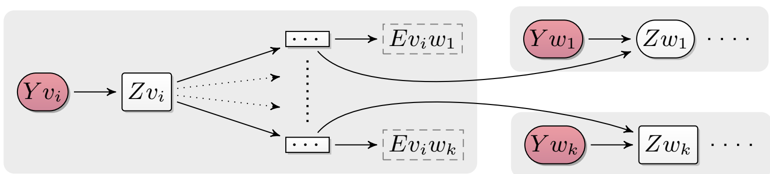

where again $\{w_1, \dots, w_k\}$ are the successors of v_i in $\mathcal{G}$. Notice that Zv_i belongs to Falsifier (as a result of skipping several positions). In general, Zv_i belongs to Verifier precisely if v_i belongs to Player 0 in $\mathcal{G}$.

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

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

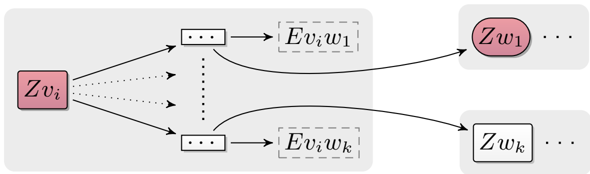

Proof of Theorem 6.3.18

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

Lemma 6.3.20. *Let $\mathcal{M}$ be a winning strategy in $\widetilde{\mathsf{MC}}(\mathcal{G}, v)$. Let $\mathcal{S}$ be a winning strategy in $\mathcal{G}$ so that $\mathcal{M}$ and $\mathcal{S}$ are corresponding strategies. Then $\pi[\![\mathcal{M}]\!] = \pi[\![\mathcal{S}]\!]$.*

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

This closes the proof of the Sum-of-strategies theorem for Büchi games: since our modifications did not affect the supremum over all winning strategies, we get

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

and by Theorem 6.3.14, restricting the supremum to winning strategies $\mathcal{S}$ that are absorption-dominant from v does not change the overall value. $\square$

6.3.5 Applications of Semiring Semantics

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

Strategy Analysis

We begin with the question what information we can derive from the Sum-of-strategies theorem. To this end, we fix a Büchi game $\mathcal{G}$ and focus on the $\mathbb{S}^\infty(X)$-interpretation π_{strat} with $X = \{X_{uv} \mid u, v \in \mathcal{G}\}$. The values $\pi_{\mathsf{strat}}[\![\mathcal{S}]\!]$ are monomials and we can read off the number of occurrences of each edge in $\mathcal{S}$ from the exponents, i.e., the monomial is a representation of the edge profile $\#_E(\mathcal{S})$. In particular, $\pi_{\mathsf{strat}}[\![\mathcal{S}_1]\!] \succeq \pi_{\mathsf{strat}}[\![\mathcal{S}_2]\!]$ if, and only if, $\mathcal{S}_1 \succeq \mathcal{S}_2$. The fact that absorptive polynomials are always finite (cf. Lemma 3.5.2) is thus another way to see that the number of absorption-dominant strategies is finite.

What can we learn from the polynomial $\pi_{\mathsf{strat}}[\![\mathsf{win}_0(v)]\!]$? First, $\pi_{\mathsf{strat}}[\![\mathsf{win}_0(v)]\!] \neq 0$ holds if, and only if, Player 0 has a winning strategy from v. By Theorem 6.3.18, we can further derive information about all absorption-dominant strategies. More precisely, we learn which edges each absorption-dominant strategy uses and how often they appear in the strategy tree. Knowing the edge profile immediately reveals whether the strategy is positional and what the positional choices are. By counting monomials, we can thus count the positional strategies, as well as the absorption-dominant strategies up to absorption-equivalence.

We can further answer questions such as: can Player 0 still win if we remove edge e? This is the case if, and only if, the polynomial $\pi_{\mathsf{strat}}[\![\mathsf{win}_0(v)]\!]$ contains a monomial without the variable X_e (if there is a winning strategy without e, then there is also an absorption-dominant strategy and hence a monomial without X_e). Going further, a more interesting question is: can Player 0 still win if edge e may only be used finitely often in each play? The answer is not immediately obvious. Consider for example the strategy $\mathcal{S}$ in Fig. 6.6b. The edge k occurs infinitely often in the strategy tree and we get $\pi_{\mathsf{strat}}[\![\mathcal{S}]\!] = abcefhg^\infty k^\infty m^\infty$. However, k is clearly played only once in each play consistent with $\mathcal{S}$, whereas edge m is played infinitely often. We cannot distinguish edges k and m just from $\pi_{\mathsf{strat}}[\![\mathcal{S}]\!]$, but we can do so if we compute $\pi_{\mathsf{strat}}[\![\mathsf{win}_0(w)]\!]$ for all positions $w \in V$, by the following criterion (notice that all of these values are computed anyway for the fixed-point iteration).

Proposition 6.3.21. *Let $\mathcal{S} \in \mathrm{WinStrat}_{\mathcal{G}}(v)$ be absorption-dominant from v, and let $e = uw \in E$ be an edge with $\#_e(\mathcal{S}) = \infty$. Then there is a unique (positional) strategy $\mathcal{S}_w \in \mathrm{WinStrat}_{\mathcal{G}}(w)$ such that $\pi_{\mathsf{strat}}[\![\mathcal{S}_w]\!] \succeq \pi_{\mathsf{strat}}[\![\mathcal{S}]\!]$. Moreover, $\mathcal{S}$ admits a play in which e occurs infinitely often if, and only if, e occurs in $\mathcal{S}_w$.*

Proof. Consider the strategy tree $\mathcal{S}$ and let ρw be an occurrence of w in $\mathcal{S}$. By assumption, w occurs infinitely often in $\mathcal{S}$. But then, for all successors $\rho w w'$ of ρw, also w' occurs infinitely often in $\mathcal{S}$. Indeed, either $w \in V_1$ and every occurrence of w must be followed by an occurrence of w'; or $w \in V_0$ and $\mathcal{S}$ plays positionally from w by Proposition 6.3.8, so again every occurrence of w is followed by an occurrence of w' in $\mathcal{S}$. By induction, it follows that $\mathcal{S}$ plays positionally from w and from all positions occurring below ρw. In particular, the substrategy $\mathcal{S}_w$ that $\mathcal{S}$ plays from ρw (and any other occurrence of w) is positional. As a substrategy, we trivially have $\pi_{\mathsf{strat}}[\![\mathcal{S}_w]\!] \succeq \pi_{\mathsf{strat}}[\![\mathcal{S}]\!]$. To see that $\mathcal{S}_w$ is unique with this property, notice that a strategy $\mathcal{S}' \in \mathrm{Strat}_{\mathcal{G}}(w)$ deviating from $\mathcal{S}_w$ must play an edge which does not occur in $\mathcal{S}_w$ and hence also not in $\mathcal{S}$, so $\pi_{\mathsf{strat}}[\![\mathcal{S}']\!] \not\succeq \pi_{\mathsf{strat}}[\![\mathcal{S}]\!]$.

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

Example 6.3.22. Consider the strategy $\mathcal{S}$ in Fig. 6.6b with $\pi_{\mathsf{strat}}[\![\mathcal{S}]\!] = abcefhg^\infty k^\infty m^\infty$ and the edge k from u to w. Since edge n does not occur in $\pi_{\mathsf{strat}}[\![\mathcal{S}]\!]$, the only winning strategy $\mathcal{S}_w$ from w we need to consider is the strategy that always stays at w, with $\pi_{\mathsf{strat}}[\![\mathcal{S}_w]\!] = m^\infty \succeq \pi_{\mathsf{strat}}[\![\mathcal{S}]\!]$. As k does not occur in $\mathcal{S}_w$, we conclude that it occurs only finitely often (and hence at most once) in each play consistent with $\mathcal{S}$.

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

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

Corollary 6.3.23. *From the polynomial $\pi_{\mathsf{strat}}[\![\mathsf{win}_0(v)]\!]$, we can efficiently (in the size of the polynomial) derive the following information:*

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

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

Reverse Analysis

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

More precisely, let $\mathcal{G} = (V, V_0, V_1, E, F)$ be a Büchi game and $v \in \mathcal{G}$ a position from which Player 1 wins. Let $E^- \subseteq E$ and $E^+ \subseteq V^2 \setminus E$ be sets of edges we are allowed to delete or add, respectively. We call a set of edges in $E^\pm := E^- \cup E^+$ a *repair* if Player 0 wins when these edges are deleted or added. Our goal is to determine all (preferably minimal) repairs.

Semiring interpretation. We achieve this by evaluating $\mathsf{win}_0(v)$ in a different interpretation that also takes negative information into account. Following the approach in [GT17a, XZAT18] and Section 5.1.3 we use dual-indeterminate polynomials for this purpose, and since we also need absorption to deal with fixed points the relevant semirings are $\mathbb{S}^\infty(X, \overline{X})$ and $\mathrm{PosBool}(X, \overline{X})$. Finding repairs does not require information on edge multiplicities, so we can work with the simpler semiring $\mathrm{PosBool}(X, \overline{X})$ that results from $\mathbb{S}^\infty(X, \overline{X})$ by dropping exponents.

We replace π_{strat} by a $\mathrm{PosBool}(X, \overline{X})$-interpretation $\pi^{\pm}_{\mathsf{repair}}$ with $X = \{X_e \mid e \in E^{\pm}\}$ defined as follows. If $vw \in E^{\pm}$, we set $\pi^{\pm}_{\mathsf{repair}}(Evw) = X_{vw}$ and $\pi^{\pm}_{\mathsf{repair}}(\neg Evw) = \overline{X_{vw}}$, all other literals are mapped to 0 or 1 according to $\mathcal{G}$ (cf. Fig. 6.8). The resulting interpretation is no longer model-defining or edge-tracking, but is instead model-compatible (cf. Definition 5.1.7).

Back and forth between monomials and models. Let $X^{\pm} = \{X_e \mid e \in E^+\} \cup \{\overline{X_e} \mid e \in E^-\}$. Given $Y \subseteq X^{\pm}$, we further write $E(Y) = \{e \mid X_e \in Y \text{ or } \overline{X_e} \in Y\}$ for the set of edges mentioned in Y. Monomials in $\mathrm{PosBool}(X, \overline{X})$ do not have exponents, so we represent them simply as sets $m \subseteq X \cup \overline{X}$. By examining what combinations of indeterminates from $X^{\pm}$ occur in the monomials of $\pi^{\pm}_{\mathsf{repair}}[\![\mathsf{win}_0(v)]\!]$, we can read off all minimal repairs as follows.

Proposition 6.3.24. *In the above setting, the following holds:*

(1) *Let $m \in \pi^{\pm}_{\mathsf{repair}}[\![\mathsf{win}_0(v)]\!]$ be a monomial. Then the set $E(m \cap X^{\pm})$ is a repair.*

(2) *Let $R \subseteq E^{\pm}$ be a repair. Then there is a monomial $m \in \pi^{\pm}_{\mathsf{repair}}[\![\mathsf{win}_0(v)]\!]$ such that $E(m \cap X^{\pm}) \subseteq R$. If R is minimal, then $E(m \cap X^{\pm}) = R$.*

Before proving Proposition 6.3.24, we illustrate the computation of minimal repairs in a small example.

Example 6.3.25. In the following game, Player 1 wins from v. We are interested in the minimal repairs with $E^+ = \{c\}$ and $E^- = \{a, b\}$.

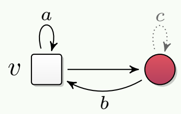

$$\pi^{\pm}_{\mathsf{repair}}[\![\mathsf{win}_0(v)]\!] = \overline{X_a} X_c + \overline{X_a} X_b$$

Evaluating $\mathsf{win}_0(v)$ in $\pi^{\pm}_{\mathsf{repair}}$ results in two monomials. The first yields the repair $\{a, c\}$, the second yields the minimal repair $\{a\}$ (notice that $X_b \notin X^{\pm}$, as edge b is already present).

The reason why we get two monomials is that we track also positive usage of edge b by X_b, but are only interested in the negative indeterminate $\overline{X_b}$ for the repairs. This is not avoidable, as we need X_b so that $\pi^{\pm}_{\mathsf{repair}}$ is model-compatible and homomorphisms into $\mathbb{B}$ work as expected (see the following proof).

Proof of Proposition 6.3.24. We prove both statements by considering homomorphisms into the Boolean semiring $\mathbb{B}$. For the first statement, let $m \in \pi^{\pm}_{\mathsf{repair}}[\![\mathsf{win}_0(v)]\!]$ be a monomial and let $h\colon X \cup \overline{X} \to \mathbb{B}$ be the unique function that respects dual-indeterminates and satisfies

- $h(x) = 1$, for all $x \in m$

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

Then, h lifts to a fully-continuous semiring homomorphism $h\colon \mathrm{PosBool}(X, \overline{X}) \to \mathbb{B}$ with $h(m) = 1$. Moreover, $h \circ \pi^{\pm}_{\mathsf{repair}}$ is a model-defining Boolean interpretation which induces a (classical) model $\mathcal{G}'$. Since semiring semantics are compatible with fully-continuous homomorphisms, we have $h \circ \pi^{\pm}_{\mathsf{repair}}[\![\mathsf{win}_0(v)]\!] = h(\pi^{\pm}_{\mathsf{repair}}[\![\mathsf{win}_0(v)]\!]) \geq h(m) = 1$ and hence $\mathcal{G}' \models \mathsf{win}_0(v)$. By the choice of h, the model $\mathcal{G}'$ is equal to $\mathcal{G}$ except that we add all edges $e \in X^+$ with $X_e \in m$, and remove all $e \in X^-$ with $\overline{X_e} \in m$. Hence $\mathcal{G}'$ results from $\mathcal{G}$ by modifying the edges $E(m \cap X^{\pm})$, and since $\mathcal{G}' \models \mathsf{win}_0(v)$ this is indeed a repair.

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

Target synthesis

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

Similar to the computation minimal repairs, we can solve this task with an interpretation π_{target} over the dual-indeterminate semiring $\mathrm{PosBool}(X, \overline{X})$ which interprets most literals, including edge literals, by Boolean values, but tracks the target set F using corresponding pairs of dual-indeterminates $\pi_{\mathsf{target}}(Fv) = X_v$ and $\pi_{\mathsf{target}}(\neg Fv) = \overline{X_v}$ for each position $v \in V$ (cf. Fig. 6.8). We can then derive all possible minimal choices for F from the polynomial $\pi_{\mathsf{target}}[\![\mathsf{win}_0(u)]\!]$, as illustrated in the following example.

Example 6.3.26. Consider the following game arena. What are the minimal choices of the target set F so that Player 0 wins from position a, or b?

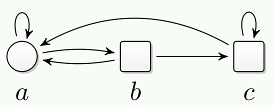

$$\pi_{\mathsf{target}}[\![\mathsf{win}_0(a)]\!] = X_a + \overline{X_a} X_b X_c$$
$$\pi_{\mathsf{target}}[\![\mathsf{win}_0(b)]\!] = X_a X_b X_c + X_a \overline{X_b} X_c + \overline{X_a} X_b X_c$$

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

In general, the positive indeterminates X_v in each monomial of $\pi_{\mathsf{target}}[\![\mathsf{win}_0(u)]\!]$ induce one possible choice of the target set F, and conversely every minimal choice of F occurs as a monomial; this follows by the same arguments as in Proposition 6.3.24.

We remark that we can do slightly better in this case, as we do not need to track negative information for the synthesis problem and hence do not need the negative indeterminates $\overline{X_v}$. In the above example, we would then obtain the polynomials $X_a + X_b X_c$ and $X_a X_c + X_b X_c$ for positions a and b, respectively, which correspond exactly to the minimal choices for F. This can be achieved by setting $\pi_{\mathsf{target}}(\neg Fv) = 1$ for all $v \in V$ and observing that this has the same effect as omitting the subformula $\neg Fv$ from win_0. Since the resulting formula is equivalent in Boolean semantics, the reasoning in the proof of Proposition 6.3.24 can be adapted to this setting.

Complexity

The previous applications show that once we have computed the polynomial $\pi_{\mathsf{strat}}[\![\mathsf{win}_0(v)]\!]$ or $\pi^{\pm}_{\mathsf{repair}}[\![\mathsf{win}_0(v)]\!]$, it is easy to derive information about strategies or minimal repairs – but how efficient is the computation of the polynomial in the first place? Recall that for a fixed formula, the LFP model-checking problem for a classical structure can be solved in polynomial time in the size of the structure (here the number of positions). The analogous problem in semiring semantics, to compute the value $\pi_{\mathsf{strat}}[\![\mathsf{win}_0(v)]\!]$, is more involved due to the computation in more expressive semirings than the Boolean one. For absorptive, fully-continuous semirings such as $\mathbb{S}^\infty(X)$ and $\mathrm{PosBool}(X, \overline{X})$, the techniques of Chapter 4 can be used to compute $\pi_{\mathsf{strat}}[\![\mathsf{win}_0(v)]\!]$ with polynomially many semiring operations (in the number of positions of the game).

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

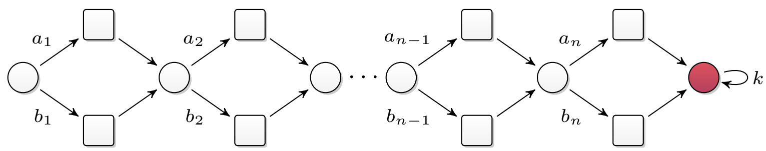

Algorithmically, the computation of the polynomials $\pi_{\mathsf{strat}}[\![\mathsf{win}_0(v)]\!]$ or $\pi^{\pm}_{\mathsf{repair}}[\![\mathsf{win}_0(v)]\!]$ is thus infeasible in many cases. It should also be noted that the questions addressed in Corollary 6.3.23 can be solved by direct methods, and just for finding *some* winning strategy or repair, these will in many cases be more efficient than computing the polynomials. Thus, the main benefit of semiring semantics in $\mathbb{S}^\infty(X)$ or $\mathrm{PosBool}(X, \overline{X})$ does not lie in a more efficient algorithmic method to compute some specific winning strategy or repair, but rather in providing a general and compact description of many important strategies at once, from which we can directly derive the answers to a number of different questions concerning the strategy analysis of a Büchi game. This is attractive in cases where the relevant polynomials are reasonably small.

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

Limitations for Cost Computation

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

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

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

(1) Player 0 only has to pay once for an edge, no matter how often it occurs (think of an "unlocking fee"), or

(2) Player 0 only has to pay for the maximal cost of any play consistent with $\mathcal{S}$, but not for all plays simultaneously.

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

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

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

Discussion

This case study has shown that semiring provenance can successfully be extended to infinite games, by utilising semiring semantics for fixed-point logics. On one side, we have used the Sum-of-strategies theorem for LFP to prove that evaluating the winning-region formula $\mathsf{win}_0(x)$ gives detailed information about all winning strategies, up to absorption. On the other side, interpretations in the dual-indeterminate semiring $\mathrm{PosBool}(X, \overline{X})$ can be used for repair and synthesis problems, based on the theoretical foundations of the Fundamental property for LFP and the universality of $\mathrm{PosBool}(X, \overline{X})$, which together enable reasoning about classes of (repaired or synthesised) models through homomorphisms into the Boolean semiring.

Not all provenance applications work well for infinite games, as we have demonstrated for cost computations in the tropical semiring. Although the sum-of-strategies theorem still guarantees that the value of $\mathsf{win}_0(x)$ is the sum over the costs of all strategies, this notion of costs may be less intuitive than for simpler types of games. It is worth noting that this issue only concerns how multiplicities of edges are handled. Applications such as access control, where the semiring has an idempotent multiplication, still work as expected.

Finally, a peculiarity of our approach needs discussion: we have defined semiring values for Büchi games through one specific LFP-formula $\mathsf{win}_0(x)$. Our results seem to justify this choice, but what is the reason that $\mathsf{win}_0(x)$ is the appropriate formula? And is this approach robust if we choose an equivalent but different formula?

Robustness depends on the semiring. In general, formulae that are logically equivalent (in classical semantics) need not be equivalent in semiring semantics. For instance, replacing the edge literals Ezu in $\mathsf{win}_0(x)$ with the classically equivalent $Ezu \wedge Ezu$ leads to different edge multiplicities in the resulting polynomial (when evaluated in π_{strat}), so the Sum-of-strategies theorem no longer holds. A more subtle example are the two formulae for generalised Büchi games in Section 6.4 below. In contrast, the reverse analysis in $\mathrm{PosBool}(X, \overline{X})$ is very robust, as the homomorphism-based reasoning is independent of the formula (to be precise, we only need that the formula expresses the desired property in classical semantics). This can be attributed to the idempotence of both semiring operations and absorption, which mimic the laws of classical semantics.

Given that the strategy analysis in $\mathbb{S}^\infty(X)$ depends on choosing the right formula $\mathsf{win}_0(x)$, it may be desirable to define semiring values for Büchi games directly without the detour to fixed-point logic, as it is done for games with simpler winning condition (i.e., reachability and safety) in [GT20]. However, defining these semiring values is more involved here due to the fixed-point alternation in the Büchi winning condition, and such a definition will inevitably end up looking very similar to $\mathsf{win}_0(x)$. In other words, LFP provides the language to conveniently express such nested fixed points, and hence we may view $\mathsf{win}_0(x)$ not just as any formula describing the winning region, but rather as the canonical description of the (semiring) semantics of Büchi games.

6.4 Towards More Expressive Games

Following the detailed case study for Büchi games, we briefly discuss how semiring provenance can be extended to infinite games with more general winning conditions. We have seen that the reverse analysis approach using model-compatible $\mathrm{PosBool}(X, \overline{X})$-interpretations works for all LFP formulae and does not make any assumptions about the game. It can thus be applied to all games (on finite game arenas) with LFP-definable winning regions, such as all infinite games with ω-regular winning conditions using a bounded number of priorities [CGLP15], which includes parity, generalised Büchi, and Muller games.

On the other hand, the strategy analysis relies on the Sum-of-strategies theorem which requires careful reasoning about the LFP-formula for the winning region. This reasoning readily generalises to parity games, and we sketch the proof of a Sum-of-strategies theorem. Because of the positional determinacy of parity games, the results on absorption-dominant strategies apply as well. This changes for generalised Büchi games, where the additional target sets require winning strategies with memory that are neither positional nor absorption-dominant. Moreover, we give two natural LFP-formulae for the winning region, of which only produces the correct edge multiplicities for a sum-of-strategies theorem. As an alternative to the LFP-based approach, we sketch how memory reductions can be used to define semiring values for more complex infinite games.

6.4.1 Parity Games

Parity games are of fundamental importance as the model-checking games for LFP and the modal μ-calculus, which includes logics such as LTL, CTL and CTL* relevant for formal verification. They also pose an interesting open problem, as the complexity of solving parity games (without bounding the number of priorities) is known to be in $\mathrm{NP} \cap \text{co-NP}$, but it is not known whether they can be solved in polynomial time. Only recently a quasipolynomial-time breakthrough was achieved [CJK$^+$17, CDF$^+$19] which also lead to quasipolynomial algorithms for evaluating alternating fixed points ([ANP21, HS21], cf. Section 4.6).

Given the importance of parity games, we seek to extend semiring provenance from Büchi to parity games, and this turns out to be straightforward. Recall that a parity game with k priorities has the form $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ with priorities $\Omega\colon V \to \{0, \dots, k-1\}$ and with the winning condition that Player 0 wins an infinite play if the smallest infinitely often occurring priority is even. Büchi games are the special case with $k = 2$ priorities. We interpret parity games as structures with universe V and signature $\tau = \{V_0, V_1, E, P_1, \dots, P_k\}$, where P_i is interpreted as the set of nodes with priority i.

Absorption-Dominant Strategies

The notion of absorption-dominant strategies and its relation to positional and persistent strategies are independent of the winning condition and also apply here. What deserves attention is whether every winning game also permits an absorption-dominant winning strategy. This is indeed the case: it is well-known that parity games are positionally determined and moreover, the parity condition of an infinite play is independent of a finite prefix of the play; these are the two requirements used in the proof of Theorem 6.3.14.

Corollary 6.4.1. *For each winning strategy $\mathcal{S} \in \mathrm{WinStrat}_{\mathcal{G}}(v)$ in a parity game $\mathcal{G}$, there is a winning strategy $\mathcal{S}' \in \mathrm{WinStrat}_{\mathcal{G}}(v)$ that is absorption-dominant from v with $\mathcal{S}' \succeq \mathcal{S}$.*

Sum-of-Strategies Theorem

The winning region of a parity game with k priorities can be expressed by the following LFP-formula (see, e.g., [CGLP15, Wal02]):

$$\mathsf{win}_0^k(x) := [\mathbf{gfp}\, X_0 x.\ [\mathbf{lfp}\, X_1 x \cdots [\mathbf{fp}\, X_{k-1} x.\ \varphi(X_0, \dots, X_{k-1}, x)](x) \cdots](x)\,](x),$$

where

$$\varphi(\overline{X}, x) := \bigvee_{0 \le i < k} \Big(P_i x \ \wedge\ \big((V_0 x \wedge \exists y(Exy \wedge X_i y)) \vee (V_1 x \wedge \forall y(Exy \to X_i y))\big)\Big).$$

We again translate the implication $Exy \to X_i y$ to the negation-normal form $\neg Exy \vee (Exy \wedge X_i y)$ and evaluate $\mathsf{win}_0^k(x)$ in the same edge-tracking interpretation π_{strat} that we used for Büchi games (cf. Fig. 6.8), with the additional literals $P_i v$ mapped to 0 or 1. To see that also here we obtain information about the edge profiles of all absorption-dominant strategies, we claim that the following sum-of-strategies theorem holds.

Theorem 6.4.2 (Sum of strategies)**.** *Let $\mathcal{G}$ be a parity game with k priorities, and let v be a position in $\mathcal{G}$. Let K be an absorptive, fully-continuous semiring and π an edge-tracking K-interpretation. Then,*

$$\pi[\![\mathsf{win}_0^k(v)]\!] = \bigsqcup \big\{ \pi[\![\mathcal{S}]\!] \;\big|\; \mathcal{S} \in \mathrm{WinStrat}_{\mathcal{G}}(v) \text{ is absorption-dominant from } v \big\}.$$

The proof is analogous to the one for Büchi games, by using the Sum-of-strategies theorem for LFP and comparing the model-checking game for $\mathsf{win}_0^k(v)$ with the original game. We sketch this comparison for a small example. Figure 6.10a shows the parity game $\mathcal{G}$ and the edge-tracking interpretation π_{strat} (which yields the value $\pi_{\mathsf{strat}}[\![\mathsf{win}_0^3(v)]\!] = a^\infty b^\infty c^\infty d^\infty$ indicating that there is only one absorption-dominant winning strategy). The model-checking game $\mathsf{MC}(\mathcal{G}, v)$ for the formula $\mathsf{win}_0^3(x)$ is sketched in Fig. 6.10b. Since an edge-tracking interpretation such as π_{strat} maps most literals to 0 or 1, large parts of this game can be omitted without affecting the values of strategies. For the game $\mathcal{G}$ and π_{strat}, this results in the simplified model-checking game shown in Fig. 6.10c. It is now easy to see that the winning strategies in $\mathcal{G}$ are in correspondence with the winning strategies in (the simplified version of) $\mathsf{MC}(\mathcal{G}, v)$, and that π assigns the same semiring values to corresponding strategies. It follows that Theorem 6.4.2 is implied by the Sum-of-strategies theorem for LFP.

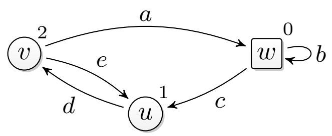

(a) Sample parity game $\mathcal{G}$ and edge-tracking interpretation π_{strat}.

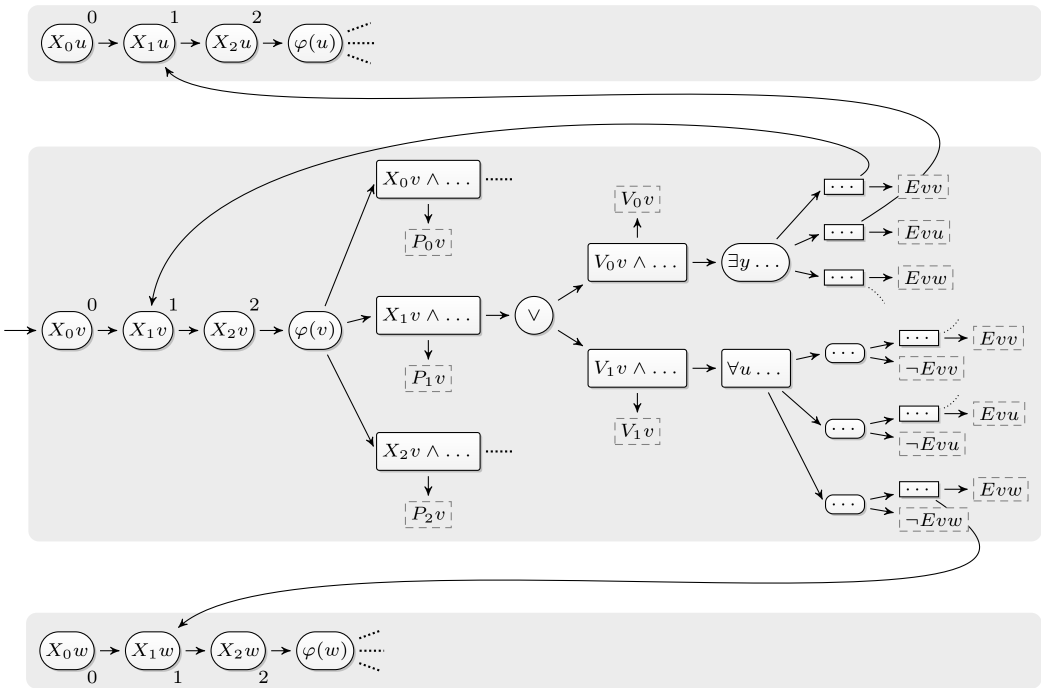

(b) Illustration of the model-checking game $\mathsf{MC}(\mathcal{G}, v)$ (with several omissions).

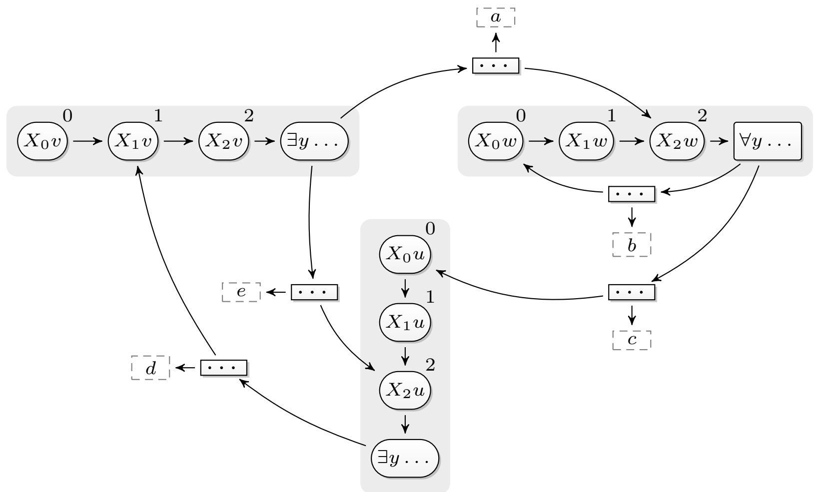

(c) Simplified version of $\mathsf{MC}(\mathcal{G}, v)$, with terminal positions L labelled by $\pi_{\mathsf{strat}}(L)$.

Figure 6.10: The proof of the Sum-of-strategies theorem for parity games.

6.4.2 Games with Finite Memory

We next consider generalised Büchi games as an examples for games without positional determinacy. These games have the form $\mathcal{G} = (V, V_0, V_1, E, F_1, \dots, F_k)$ and Player 0 has to visit all target sets $F_1, \dots, F_k$ infinitely often in a winning play. Büchi games are the special case with $k = 1$. A simple running example with two target sets $F_1 = \{w_1\}$ and $F_2 = \{w_2\}$ is shown in Fig. 6.11a. Any winning play has to alternate infinitely often between the moves b and c at u, so there is no persistent winning strategy, and hence also no absorption-dominant or positional one (so Theorem 6.3.14 does not hold anymore). Under the edge-tracking interpretation indicated in Fig. 6.11a, all winning strategies have value $ab^\infty c^\infty d^\infty$. We discuss two ways to define semiring values for generalised Büchi games.

Semiring Values via LFP

Winning regions in generalised Büchi games with a fixed number k of target sets are LFP-definable. We view such a game as τ-structure with $\tau = \{V_0, V_1, E, P_1, \dots, P_k\}$. The following helper formula is useful to define the winning regions. It says that Player 0 can force a visit to X in the next step:

$$\varphi_{\text{move}}(x, X) \coloneqq (V_0 x \wedge \exists y(Exy \wedge Xy)) \vee (V_1 x \wedge \forall y(Exy \to Xy)).$$

One possibility to describe the winning region is to assert that Player 0 can, at any time, enforce a visit to each target set (and back to a position in the greatest fixed point):

$$\textsf{win-all}_0^k(x) \coloneqq \Big[\mathbf{gfp}\, Y\, y.\, \bigwedge_{i=1}^{k} [\mathbf{lfp}\, Z\, z.\, (F_i z \wedge \varphi_{\text{move}}(z, Y)) \vee (\neg F_i z \wedge \varphi_{\text{move}}(z, Z)](y)\Big](x).$$

A second possibility is to assert that the target sets are visited in order. That is, we can first reach a position in F_1 from which we can reach a position in F_2 and so on, and after reaching F_k this process is repeated ad infinitum through the **gfp**:

$$\textsf{win-seq}_0^k(x) \coloneqq \big[\mathbf{gfp}\, Y\, y.\, \varphi^1_{\text{reach}}(y, Y)\big](x),$$

where the reachability formulae are defined as follows, for $i < k$:

$$\begin{aligned} \varphi^i_{\text{reach}}(y, Y) &\coloneqq [\mathbf{lfp}\, Z\, z.\, (F_i z \wedge \varphi^{i+1}_{\text{reach}}(z, Y)) \vee (\neg F_i z \wedge \varphi_{\text{move}}(z, Z))](y), \\ \varphi^k_{\text{reach}}(y, Y) &\coloneqq [\mathbf{lfp}\, Z\, z.\, (F_k z \wedge \varphi_{\text{move}}(z, Y)) \vee (\neg F_k z \wedge \varphi_{\text{move}}(z, Z))](y). \end{aligned}$$

The formulae $\textsf{win-all}_0^k(x)$ and $\textsf{win-seq}_0^k(x)$ both describe the winning region and are thus equivalent in classical semantics. However, there is a subtle difference in semiring semantics.

Example 6.4.3. For the $\mathbb{S}^\infty(X)$-interpretation $\pi_{\textsf{strat}}$ shown in Fig. 6.11a, we obtain:

$$\begin{aligned} \pi_{\textsf{strat}}[\![\textsf{win-all}_0^2(v)]\!] &= a^2 b^\infty c^\infty d^\infty, \\ \pi_{\textsf{strat}}[\![\textsf{win-seq}_0^2(v)]\!] &= a b^\infty c^\infty d^\infty. \end{aligned}$$

Both correctly describe the edges used by the winning strategies (which are all absorption-equivalent), but they differ in the multiplicity of edge a. The reason is that $\textsf{win-all}_0^2(v)$ asserts that both F_1 and F_2 are reachable from v, and since the paths to F_1 and F_2 both use edge

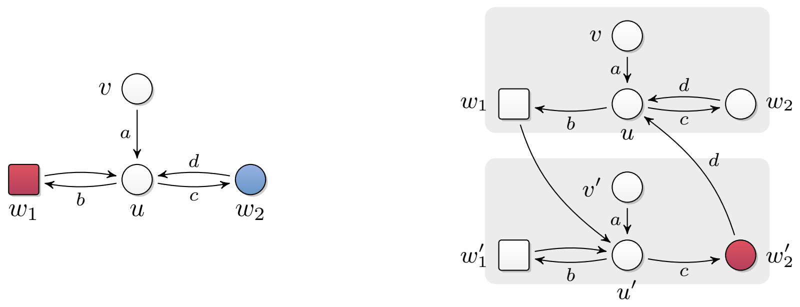

(a) Red nodes belong to F_1, blue nodes to F_2.

(b) Visiting F_i switches to the i-th copy, target positions are in the last copy.

Figure 6.11: A generalised Büchi game and its reduction to a Büchi game.

a, we get a^2 in the result. In contrast, $\mathsf{win\text{-}seq}_0^2(v)$ is similar to the reduction in Fig. 6.11b. Here, a is used only once to reach w_1 and the evaluation of all further **lfp**-formulae takes place in the cyclic part of the game.

This example shows that care has to be taken in selecting the appropriate LFP-formula to define semiring values. Proving a sum-of-strategies theorem also becomes more difficult, as the model-checking game for $\mathsf{win\text{-}seq}_0^k(x)$ is a parity game (or a Büchi game, as only two priorities are needed), but not a generalised Büchi game. In order to compare the winning strategies in these games, we thus have to first apply the memory reduction from the generalised Büchi game to a Büchi game discussed below (cf. Fig. 6.11b). We note (without proof) that this indeed works out and results in the following sum-of-strategies theorem, since $\mathsf{win\text{-}seq}_0^k(x)$ is essentially a formalisation of the reduction. Notice that we cannot restrict the supremum to the absorption-dominant winning strategies, as these may not exist (but of course values of some strategies may be absorbed, as K is absorptive).

Theorem 6.4.4 (Sum of strategies)**.** *Let $\mathcal{G}$ be a generalised Büchi game with k target sets, and let v be a position in $\mathcal{G}$. Let K be an absorptive, fully-continuous semiring and π an edge-tracking K-interpretation. Then,*

$$\pi[\![\mathsf{win\text{-}seq}_0^k(v)]\!] = \bigsqcup \{\pi[\![\mathcal{S}]\!] \mid \mathcal{S} \in \mathrm{WinStrat}_{\mathcal{G}}(v)\}.$$

Since we are already using the memory reduction for the proof, we can also use it to define semiring values more directly, without fixed-point logic.

Semiring Values via Reductions

A winning strategy for the game in Fig. 6.11a has to switch between using the edges b and c at u, so an implementation of such a strategy would needs at least two memory states (or one bit of memory) to remember the previous decision. We can thus reduce the game to a Büchi game by creating two copies as shown in Fig. 6.11b, corresponding to the two memory states. We switch between copies after visiting the respective target sets. The target set of the resulting Büchi game is then the set F_2 in the second component, since upon reaching this set both F_1 and F_2 have been visited.

This idea can be formalised with the notion of memory reductions (cf. [DJW97, Grä16]).

Definition 6.4.5. Let $\mathcal{G}$ be an infinite game with arena $G = (V, V_0, V_1, E)$. A *memory automaton* for G is a finite automaton $\mathcal{M} = (M, m_0, \delta)$ with state set M, initial state $m_0 \in M$ and transition function $\delta\colon M \times V \to M$.

The arena $G \times \mathcal{M}$ is defined as the usual product: $G \times \mathcal{M} = (V \times M, V_0 \times M, V_1 \times M, E')$ with moves $E' = \{((v, m), (v', \delta(v, m))) \mid (v, v') \in E,\ m \in M\}$.

Any play $\rho = v_0 v_1 v_2 \dots$ from v_0 in G induces a unique play from (v_0, m_0) in $G \times \mathcal{M}$ by starting at the initial state and applying the transition function of $\mathcal{M}$. That is, $\rho' = (v_0, m_0)(v_1 m_1)(v_2 m_2) \dots$ with $m_{i+1} = \delta(m_i, v_i)$ in each step. Conversely, each play in $G \times \mathcal{M}$ induces a unique play in G by projection to the first component (i.e., forgetting about the memory states). If an appropriate winning condition is imposed on $G \times \mathcal{M}$, so that ρ and ρ' are won by the same player, then we can reduce $\mathcal{G}$ to a game with arena $G \times \mathcal{M}$. Since the memory is made explicit, such a game can have a simpler winning condition that permits positional winning strategies.

Definition 6.4.6. Let $\mathcal{G}$ be an infinite game with arena G, and let $\mathcal{M} = (M, m_0, \delta)$ be a memory automaton for G. We say that $\mathcal{G}$ reduces to an infinite game $\mathcal{G}'$ via memory $\mathcal{M}$, denoted $\mathcal{G} \leq_{\mathcal{M}} \mathcal{G}'$, if the following conditions hold:

- $\mathcal{G}'$ has arena $G \times \mathcal{M}$,
- every play ρ' in $\mathcal{G}'$ from a position (v, m_0) has the same winner as the projected play ρ in $\mathcal{G}$.

The correspondence of plays in $\mathcal{G}$ and $\mathcal{G}'$ also holds for strategies (by applying the reasoning on plays to each node in the strategy tree), and hence Player 0 wins from v in $\mathcal{G}$ if, and only if, they win from (v, m_0) in $\mathcal{G}'$.

Example 6.4.7. The generalised Büchi game $\mathcal{G}$ in Fig. 6.11a reduces to the Büchi game $\mathcal{G}'$ in Fig. 6.11b with memory $\mathcal{M} = (\{1, 2\}, 1, \delta)$, where

$$\delta(m, v) = \begin{cases} m, & \text{if } v \notin F_m, \\ 3 - m, & \text{if } v \in F_m. \end{cases}$$

Positional winning strategies for Player 0 in $\mathcal{G}'$ induce non-positional (and even non-persistent) winning strategies in $\mathcal{G}$ by projection.

For an edge-tracking interpretation that ignores the memory states, winning strategies in $\mathcal{G}'$ receive the same value as their projections in $\mathcal{G}$, and hence the sum of the strategy values is the same for both games.

Lemma 6.4.8. *Let $\mathcal{G} \leq_{\mathcal{M}} \mathcal{G}'$ with $\mathcal{M} = (M, m_0, \delta)$. Let further π be an edge-tracking K-interpretation for $\mathcal{G}$ into an absorptive, fully-continuous semiring. We extend π to the* *signature of $\mathcal{G}'$ by setting $\pi(E(v,m)(v',m')) = \pi(Evv')$. Then for each $v \in \mathcal{G}$,*

$$\bigsqcup\{\pi[\![\mathcal{S}]\!] \mid \mathcal{S} \in \mathrm{WinStrat}_{\mathcal{G}}(v)\} = \bigsqcup\{\pi[\![\mathcal{S}']\!] \mid \mathcal{S}' \in \mathrm{WinStrat}_{\mathcal{G}'}((v,m_0))\}.$$

Proof sketch. Recall that each partial play $\rho = v_0 v_1 \dots v_n$ from $v \in \mathcal{G}$ induces the unique partial play $\rho' = (v_0,m_0)(v_1,m_1)\dots(v_n,m_n)$ from $(v,m_0) \in \mathcal{G}'$ according to the transition function δ. Conversely, ρ results from ρ' by projection.

By applying this to all nodes in the strategy tree, each strategy $\mathcal{S} \in \mathrm{Strat}_{\mathcal{G}}(v)$ induces a unique strategy $\mathcal{S}' \in \mathrm{Strat}_{\mathcal{G}'}((v,m_0))$ and vice versa. Since $\mathcal{G} \leq_{\mathcal{M}} \mathcal{G}'$, either none or both of $\mathcal{S}$ and $\mathcal{S}'$ are winning for Player 0. Since π ignores the memory states, we further have $\pi[\![\mathcal{S}]\!] = \pi[\![\mathcal{S}']\!]$, and hence also the suprema are equal. □

This leads to the following approach to define semiring values for positions in infinite games with finite memory. Let $\mathcal{G}$ be an infinite game (permitting finite-memory winning strategies) and π an edge-tracking interpretation in an absorptive, fully-continuous semiring.

(1) Reduce $\mathcal{G}$ to a parity game $\mathcal{G}'$ via finite memory.

(2) Extend π to $\mathcal{G}'$ by ignoring the memory information.

(3) Obtain semiring values for all positions $(v,m_0) \in \mathcal{G}'$ according to Section 6.4.1, by evaluating the appropriate LFP-formula $\mathsf{win}_0^k(v)$ in π.

(4) Define the semiring value for position $v \in \mathcal{G}$ as the value for $(v,m_0) \in \mathcal{G}'$. By the previous lemma and Theorem 6.4.2, this valuation satisfies a sum-of-strategies theorem.

This general strategy is applicable to generalised Büchi games (via the reduction used in Fig. 6.11) and also to request-response games (see [WHT03]) and Muller games (via the standard *latest appearance record*-reduction to parity games, see e.g. [GTW02]). The downside is, of course, that the potentially large arena $G \times \mathcal{M}$ has to be constructed to evaluate $\mathsf{win}_0^k(v)$.

Chapter 7

Zero-One Laws

So far, we have looked at semiring semantics for logics, in particular fixed-point logics, from two angles: we have studied its foundations, such as universal semirings and the computation of greatest fixed points, and we have taken the perspective of provenance analysis by using semiring semantics to obtain information about the evaluation of formulae and infinite games.

This chapter takes a different perspective and puts semiring semantics itself at the centre of attention, by investigating its model-theoretic properties. Since we are mostly working with interpretations over finite universes, we are interested in understanding how the classical results of finite model theory behave for semiring interpretations. This is worthwhile both to advance the theory of semiring semantics, and also for a more fine-grained analysis of the classical results. Indeed, we can think of different classes of semirings as arranged on a ladder: starting from the Boolean semiring, the next steps take us to min-max and then to lattice semirings, followed by fully idempotent and absorptive semirings (which are incomparable classes), then moving to the idempotent ones, and eventually reaching the class of all naturally ordered semirings (see Fig. 7.1). Each step can be seen as giving up on certain properties of classical logic, such as multiplicative idempotence ($\varphi \land \varphi \equiv \varphi$) or absorption ($\varphi \lor (\varphi \land \vartheta) \equiv \varphi$). Since semiring semantics in the Boolean semiring coincides with standard semantics, the classical theorems of finite model theory all hold at the first step of the ladder, but many fail for arbitrary semirings. So we may ask for each theorem: how far can we climb up on the ladder of semiring semantics? This can give new insights about the assumptions needed for these results. Moreover, proving generalisations for larger classes of semirings often requires new techniques, which sometimes also transfer back to the classical setting.

The development of finite model theory for semiring semantics is a very recent line of work, with some fundamental results and a number of interesting open questions for future research. We give an overview on what is known so far in Section 7.1. In this thesis we focus on a well-known result of finite model theory: the classical 0-1 laws about the asymptotic behaviour of logic on random structures of increasing size. These are particularly interesting for us as they apply not only to first-order logic, but also to fixed-point logic.

Zero-one laws for FO and $L^{\omega}_{\infty\omega}$. We briefly recall some basic facts about 0-1 laws for first-order logic on random structures. For a finite relational vocabulary τ, a finite universe $[n] = \{0, \dots, n-1\}$, a constant p with $0 < p < 1$, we consider the probability spaces $\mathrm{Str}_{n,p}(\tau)$ of random τ-structures with universe $[n]$, obtained by the random experiment

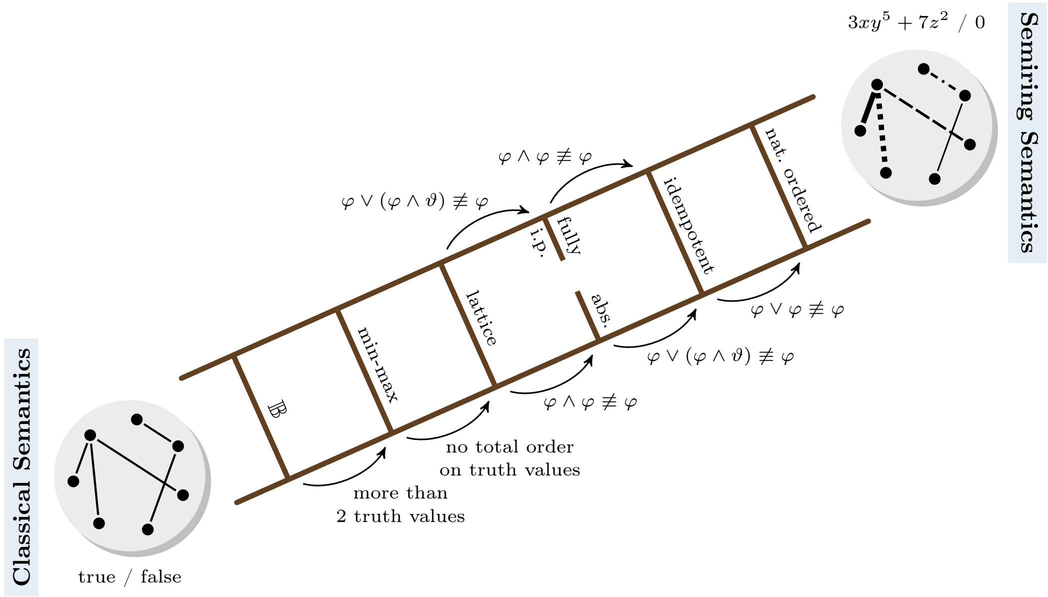

Figure 7.1: Climbing the ladder form classical semantics to semiring semantics.

which, independently for each relational atom $\alpha = Ri_1 \dots i_k$ (where $R \in \tau$ has arity k, and $i_1, \dots i_k \in [n]$), makes a random choice whether α shall be true (with probability p), or false (with probability $1 - p$). The most common such distribution is the *uniform* one, with $p = \frac{1}{2}$, which gives to each possible τ-structure over $[n]$ the same probability. Beyond the case where p is constant, there has been intensive research on probability spaces $\mathrm{Str}_{n,p}(\tau)$, where the probabilities of atomic facts depend on the size of the universe, i.e. are given by a function $p \colon \mathbb{N} \to [0, 1]$; however, here we consider only constant p, so that atomic probabilities are the same for each n.

Given a first-order sentence $\psi \in \mathrm{FO}(\tau)$ we define $\mu_{n,p}(\psi)$ to be the probability that a random structure from $\mathrm{Str}_{n,p}(\tau)$ is a model of ψ, and we are interested in the behaviour of the sequence $(\mu_{n,p}(\psi))_{n<\omega}$ as n tends to infinity (see Fig. 7.2 for an illustration). A fundamental result, proved in [GKLT69] and [Fag76] is the celebrated 0-1 law for first-order logic:

Theorem 7.0.1. *For every sentence $\psi \in \mathrm{FO}(\tau)$ the asymptotic probability $\lim_{n\to\infty} \mu_{n,p}(\psi)$ exists, and is either 0 or 1. Moreover, the sequence $(\mu_{n,p}(\psi))_{n<\omega}$ converges exponentially fast to this limit.*

Informally, we say that each sentence $\psi \in \mathrm{FO}(\tau)$ is either almost surely true or almost surely false on finite structures. There are several possibilities to prove the 0-1 law. In the original proof of Glebskii et al. [GKLT69] a surprisingly elementary (in the words of Compton [Com89]) quantifier elimination argument was used. Later, Fagin [Fag76] presented a different proof using the machinery of model theory, based on the theory of *extension axioms*. These say, informally, that every configuration of k points can be extended in any consistent way to a configuration of $k + 1$ points. For undirected graphs, for instance, this means that for any collection $v_1, \dots, v_k$, of k nodes and any $i \leq k$ there is a further node w which is adjacent to $v_1, \dots, v_i$, but not to $v_{i+1}, \dots, v_k$.

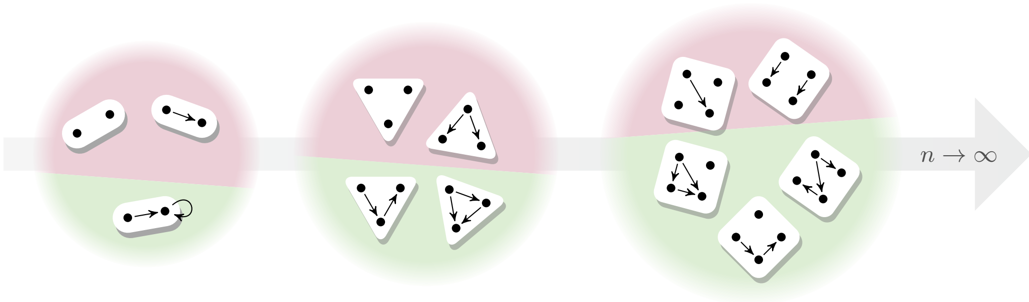

Figure 7.2: Schematic illustration of the classical 0-1 law. The fraction of structures on n elements satisfying $\psi = \exists x \exists y \exists z (Exy \wedge Eyz)$ is indicated in green (not drawn to scale). For $n \to \infty$, this fraction converges to either 0 or 1 (in this case to 1).

Fagin's proof relies on the following facts:

- Each extension axiom is almost surely true on random structures.
- The theory T of all extension axioms is ω-categorical, i.e. it has a unique countable model, up to isomorphism, which is sometimes called the *random τ-structure* or, in the case of undirected graphs, the *Rado graph*.
- It follows that T is complete, i.e. either $T \models \psi$ or $T \models \neg\psi$, for every sentence $\psi \in \mathrm{FO}(\tau)$. By compactness it then follows that either ψ or $\neg\psi$ is a consequence of finitely many extension axioms, and is therefore almost surely true on random τ-structures.
- Moreover, it follows that ψ is almost surely true on finite τ-structures if, and only if, ψ is true in the countable random τ-structure.

This argument has been extended to finite-variable infinitary logic $L^{\omega}_{\infty\omega}$ by Kolaitis and Vardi [KV92]. They observed that on the random countable τ-structure, $L^{k}_{\infty\omega}$ collapses to $L^{k}_{\omega\omega} = \mathrm{FO}^{k}$ (the k-variable fragment), and this collapse is witnessed by an almost surely true FO^{k}-formula. (In fact, they gave three different proofs of the 0-1 law for $L^{\omega}_{\infty\omega}$, the other ones via pebble games and via quantifier elimination.) This result marks the boundary of 0-1 laws for infinitary logic: if infinitely many variables are permitted, then one can express even cardinality with an infinite disjunction $\bigvee_{n \text{ even}} \varphi_n$, but this property does not have an asymptotic probability.

More general zero-one laws. Following its discovery for first-order logic, the 0-1 law has been extended in many directions. The first extensions to (partial) fixed-point logic have since been subsumed by the 0-1 law for $L^{\omega}_{\infty\omega}$ (see [KV92] for an overview). The situation becomes more interesting for second-order logic: while the 0-1 law fails in general (already for the monadic fragment, in a strong way [KS85]), it does hold for certain fragments of existential second-order logic, which can be precisely characterised in terms of quantifier prefix classes [KV00]. A different direction are more general probability distributions where the probability p is a function of n [Spe93] which captures, for instance, random sparse graphs. These and further extensions to more general kinds of convergence laws and to specific classes of mathematical structures are surveyed in [Com89].

Such results often give a simple and direct argument for proving that properties for which these convergence laws fail cannot be expressed in such logics. A simple example is the fact that no first-order sentence (and no sentence in infinitary logic, as long as the number of variables is finite) can distinguish between finite structures of even and odd cardinality. More practically, 0-1 laws have also been put to use for studying query answering in the context of uncertain data (see e.g. [Lib18]). Following the work presented here, 0-1 laws have recently been generalised also to many-valued logics over lattice algebras [BCN23], which extends our results on lattice semirings.

Zero-one laws for semiring semantics. It is a natural question whether semiring semantics admits results that are analogous to the 0-1 laws of finite model theory. Fixing a probability distribution over a given semiring K, the notion of a random structure generalizes in a rather straightforward way to the notion of a random K-interpretation, so the typical questions studied for logic on random structures make sense also in the context of semiring semantics. Notice that the 0-1 law splits the relational first-order and finite-variable infinitary sentences into two classes: those that are almost surely true and those that are almost surely false (on finite structures). This leads to the following questions:

- Is there a similar split for valuations in other semirings than the Boolean one? For instance, given a finite semiring K, can we partition $L^{\omega}_{\infty\omega}(\tau)$ into classes $(\Phi_j)_{j\in K}$ such that for each $j \in K$, every sentence in Φ_j evaluates almost surely to j, under random semiring interpretations into K?
- Are all these classes Φ_j nonempty, or do the almost sure valuations concentrate on just a few values, for instance on 0 and 1?
- How are these partitions, if they exist, related if we compare different semirings? For instance, are the almost surely false sentences always the same, no matter which semiring we consider?
- Are there similar results for infinite semirings? More generally, what kind of algebraic conditions do we have to impose on the underlying semiring to obtain results that are analogous to the traditional 0-1 law?
- Finally, there also are questions of complexity: how difficult is it to compute the almost sure valuation of a given first-order sentence (assuming that it exists)?

Our methods to answer such questions combine on the one hand techniques that are adapted from traditional studies of logic on random structures, such as extension properties of atomic types, and on the other side specific ideas of semiring semantics, such as the use of polynomials with indeterminates for tracking the literals.

The general picture that emerges from our analysis of random K-interpretations for a number of different semirings K shows that there indeed is a 0-1 law for first-order and infinitary logic, saying that with probabilities converging to 1 exponentially fast, the valuation $\pi[\![\psi]\!]$ of a sentence $\psi \in L^{\omega}_{\infty\omega}$ almost surely concentrates on one specific value $j \in K$. The induced partition $(\Phi_j)_{j\in K}$ and the details of our proofs depend on the semiring K, but the general technique using algebraic representations of formulae is the same, and may be applicable also to further semirings. More precisely, we make the following contributions to the above questions.

- For finite lattice semirings, we prove a 0-1 law for infinitary sentences by means of descriptions of formulae by polynomials, which can be seen as a quantifier elimination argument (Section 7.2.3). Interestingly, it is neither the case that all semiring elements occur as almost sure valuations, nor that only the values 0 and 1 occur. Instead, the 0-1 law induces a partition of $L^{\omega}_{\infty\omega}(\tau)$ into three classes Φ_0, Φ_1, and Φ_ε of sentences that, respectively, almost surely evaluate to 0, to 1, and to $\varepsilon = \bigsqcap\{j \in K \mid j \neq 0\}$. Notice that ε is the smallest element greater than 0, if such an element exists (as, for instance, in finite min-max semirings). In other words, 1 and ε are the largest and smallest values we consider as *true* (whereas 0 represents *false*). For all other values $j \in K$ we have that $\Phi_j = \emptyset$. Over most semirings the three classes Φ_0, Φ_1, Φ_ε are distinct, but there are a few cases where we have only two classes because Φ_ε collapses to Φ_0 (as in lattices semirings without a smallest positive element), or to Φ_1 (in the Boolean semiring).

- Since the polynomials depend only on the formula and not on the semiring, the split into Φ_0, Φ_1 and Φ_ε is the same for all finite lattice semirings (unless Φ_ε collapses), and the classes Φ_1 and Φ_ε contain precisely the sentences that are almost surely true in classical semantics.

  We further show that the problem of determining into which class a first-order sentence falls remains Pspace-complete, as in classical semantics (Section 7.2.4).

- We show that the 0-1 law can be generalised to infinite lattice semirings, under mild assumptions on the probability distributions (Section 7.3). The proof uses the same polynomials as for finite lattice semirings, and hence sentences in Φ_0 and Φ_1 have the same almost sure valuations also in infinite lattice semirings.

- The results on lattice semirings also imply a partial 0-1 law for absorptive semirings (Section 7.4.1). An important instance where the analysis is somewhat different are the natural semirings $\mathbb{N}$ and $\mathbb{N}^\infty$, where multiplication is increasing rather than decreasing (as for absorptive semirings). We show that the 0-1 law still holds for the natural semirings (Section 7.4.2), but the proof relies on more general ∞-expressions instead of polynomials and there are rather trivial constructions showing that every number $j \in \mathbb{N}$ appears as almost sure valuation.

The 0-1 laws for first-order logic have been obtained in collaboration with Erich Grädel, Hayyan Helal, and Richard Wilke, and have been published in [GHNW22a, GHNW22b]. The generalisation to infinitary logic has not been published before.

7.1 Model Theory for Semiring Semantics

We summarise the development of model theory for semiring semantics of first-order logic. Recall the definition of semiring semantics $\pi[\![\psi]\!]$ for FO-formulae ψ from Section 5.1. In reminiscence of classical structures, we always work with K-interpretations π, for some commutative semiring K, that are *model-defining*, and unless mentioned otherwise we always assume a finite relational signature τ and *finite universes*. In particular, two sentences φ and ψ are K-equivalent, denoted $\varphi \equiv_K \psi$, if $\pi[\![\varphi]\!] = \pi[\![\psi]\!]$ holds for all model-defining K-interpretations π over finite universe.

7.1.1 Survey

One of the fundamental questions in model theory is whether two given structures can be distinguished in some logic, and how one can prove that this is not the case. For semiring semantics, it is clear that fewer equivalences between formulae hold for more general classes of semirings (cf. Fig. 7.1), but it is not obvious what this means for equivalences between models. We first note that the concepts of elementary equivalence and isomorphism naturally extend to semiring semantics. Let $\pi\colon \mathrm{Lit}_A(\tau) \to K$ and $\pi'\colon \mathrm{Lit}_B(\tau) \to K$ be two K-interpretations over the same signature τ, but different finite universes A and B. We say that π and π' are *isomorphic*, denoted $\pi \cong \pi'$, if there is a bijection $f\colon A \to B$ that is compatible with the semiring values of literals, i.e., $\pi(R\boldsymbol{a}) = \pi'(Rf(\boldsymbol{a}))$ and $\pi(\neg R\boldsymbol{a}) = \pi'(\neg Rf(\boldsymbol{a}))$ for all $R \in \tau$ and $\boldsymbol{a} \subseteq A$. We call π and π' *elementary equivalent*, denoted $\pi \equiv \pi'$, if $\pi[\![\varphi]\!] = \pi'[\![\varphi]\!]$ holds for all sentences $\varphi \in \mathrm{FO}(\tau)$. If we restrict φ to sentences with quantifier rank at most m, we obtain the more fine-grained *m-equivalence* $\pi \equiv_m \pi'$.

Elementary equivalence vs. isomorphism. It is well-known that first-order logic, in classical semantics, is strong enough to capture isomorphism of finite structures, as each finite structure $\mathfrak{A}$ can be described (up to isomorphism) by a characteristic sentence $\psi_{\mathfrak{A}}$. This situation changes in interesting ways when we move to semiring semantics, as studied by Grädel and Mrkonjić [GM21]. It turns out that there are examples of non-isomorphic but elementary equivalent interpretations π, π' for any fully-idempotent semiring with at least three elements (and universe with at least four elements). Hence even the simplest min-max semiring different from $\mathbb{B}$ exhibits such a counterexample. The proof is based on a new technique of *separating homomorphisms*, where a set H of semiring homomorphisms $h\colon K \to L$ is (diagonally) separating if for all $a, b \in K$ with $a \neq b$, there is a homomorphism $h \in H$ so that also $h(a) \neq h(b)$. This means that any sentence φ that distinguishes two K-interpretations π, π' must also distinguish some of the L-interpretations $h \circ \pi$, $h \circ \pi'$ for $h \in H$, and if one can show that this is not the case then $\pi \equiv \pi'$ must hold.

Intuitively, first-order logic under semiring semantics is weaker compared to classical semantics, as it cannot talk directly about semiring values and lacks compositionality. Indeed, we cannot write a formula such as $Ra = \frac{1}{2}$ to enforce that $\pi(Ra) = \frac{1}{2}$; we can only compare the values of sentences such as $\exists x\, Rx$. In min-max semirings, we can thus only make assertions about the minimal or maximal value of a certain relation, and this is precisely the limitation that underlies the initial counterexample in [GM21]. Moreover, a conjunction $\varphi_1 \wedge \varphi_2$ does not express the same information as evaluating both φ_1 and φ_2 separately, as we only get information about the minimal value of the two formulae; the other value is lost.

This leads to further observations that differ from the classical case: for a given interpretation π in the Viterbi semiring (and the isomorphic tropical semiring), there is a set of finitely many sentences that describes π up to isomorphism, but an additional lower bound implies that this set cannot be reduced to a single characteristic sentence. Further results are obtained for various polynomial semirings, to the effect that elementary equivalence implies isomorphism in the semirings $\mathbb{V}$, $\mathbb{T}$, $\mathbb{N}$, and $\mathbb{N}[X]$, whereas counterexamples exists for all fully-idempotent semirings (except for $\mathbb{B}$) and also for $\mathrm{Why}(X)$, $\mathbb{S}(X)$, and $\mathbb{B}[X]$. This includes many of the semirings we are interested in, and the tools developed in [GM21] may apply to further semirings. However, a complete classification remains yet to be achieved.

Ehrenfeucht-Fraïssé games. An important tool to prove elementary equivalence and also m-equivalence of classical relational structures $\mathfrak{A}$ and $\mathfrak{B}$ are the Ehrenfeucht-Fraïssé games (EF-games) $G(\mathfrak{A}, \mathfrak{B})$ and $G_m(\mathfrak{A}, \mathfrak{B})$. These games are played by two players: first Spoiler chooses an element in one of the two structures, then Duplicator responds with an element in the other structure. Duplicator wins $G_m(\mathfrak{A}, \mathfrak{B})$ if after m such rounds, the chosen pairs of elements define a partial isomorphism. The game $G(\mathfrak{A}, \mathfrak{B})$ follows the same rules, but begins with Spoiler choosing the number m of rounds that will be played. The well-known Ehrenfeucht-Fraïssé theorem states that Duplicator wins $G_m(\mathfrak{A}, \mathfrak{B})$ if, and only if, $\mathfrak{A} \equiv_m \mathfrak{B}$. We refer to [EF95] for more background.

EF-games can be applied to semiring interpretations without modifications, denoted $G_m(\pi, \pi')$; the semiring values are important only for verifying the partial isomorphism in the end. This leads to the question for which semirings the Ehrenfeucht-Fraïssé theorem applies, and if the rules of the game can be adapted in cases where it fails. This has been studied in detail by Brinke, Grädel, and Mrkonjić [BGM24] who provide a convenient overview on the soundness and completeness of EF-games over various semirings (see [BGM24, Fig. 1]).

We summarise some of their observations. For semirings where elementary equivalence does not imply isomorphism, such as the fully-idempotent ones (except for $\mathbb{B}$), the games $G_m(\pi, \pi')$ are not complete, i.e., $\pi \equiv_m \pi'$ does not imply that Duplicator has a winning strategy (since for sufficiently large m, such a winning strategy would construct an isomorphism). On the other hand, it turns out that G_m is sound (i.e., a winning strategy for Duplicator implies $\pi \equiv_m \pi'$) precisely if the semiring is fully-idempotent, so that the games G_m cannot be both sound and complete for any semiring other than $\mathbb{B}$. (The game G is also considered in [BGM24], but for finite structures it is trivially sound by choosing m as the size of the universe, so we do not consider it here.)

Despite these restrictions, EF-games can still be a useful tool for inexpressibility results. Indeed, if we want to prove that FO cannot distinguish two structures (with different properties), then the soundness of G_m is sufficient. This is used, for instance, to prove Hanf locality (see the next paragraph). An alternative explored in [BGM24] is to consider more powerful games: bijection (or counting) games BG_m are sound for $\equiv_m$ on any semiring, but are often too strong to be complete, with the notable exception of $\mathbb{N}$ and $\mathbb{N}[X]$. Instead, a new kind of model comparison game is proposed, the homomorphism game inspired by separating homomorphisms, which is both sound and complete for $\equiv_m$ on all lattice semirings.

We remark that the results summarised so far paint a rather complex picture. The soundness of the games G_m is easy to characterise: we can climb the ladder up to the fully-idempotent semirings. Similarly, the new homomorphism games work for all lattice semirings. Completeness of G_m and the strength of elementary equivalence are less clear. Intuitively, more general semirings such as $\mathbb{N}[X]$ can be more informative (a polynomial represents more information than, say, a value in a min-max semiring), and this counteracts the lack of compositionality. However, it seems that the analysis really depends on the specific semiring, and cannot be carried out for entire classes (as in Fig. 7.1).

Locality. A second well-known tool to prove inexpressibility results are the locality properties of first-order logic, most notably Hanf locality and Gaifman locality. In classical semantics, Hanf's locality theorem is a way to prove that two structures $\mathfrak{A}$ and $\mathfrak{B}$ are elementary equivalent, by showing that the local neighbourhoods in each structure are similar. Gaifman's theorem states that each sentence can be brought (effectively, but with nonelementary blowup in the formula size) into a Gaifman normal form, which intuitively makes statements only about a finite number of distinct local neighbourhoods.

These locality properties have been generalised to semiring semantics by Bizière, Grädel, and myself [BGN23a]. The first step is an appropriate notion of locality. In classical semantics, distances in the Gaifman graph $G(\mathfrak{A})$ are used, where two elements $a, b \in A$ are connected by en edge if they appear together in a true atom. These atoms are now mapped to semiring values, so different notions of distance are possible. However, the approach taken here is to keep the distance notion Boolean. That is, the Gaifman graph $G(\pi)$ of a model-defining semiring interpretation π is defined as the Gaifman graph $G(\mathfrak{A}_\pi)$ of the induced Boolean structure, and distances $d(a, b)$ for $a, b \in A$ are defined as distances in $G(\pi)$. There is one caveat: semiring interpretations can assign nontrivial values also to negative atoms, say to $\neg Eab$ for a graph. But then, the value of $\neg Eab$ can be relevant even if a and b are far apart in $G(\pi)$, and hence no locality properties can hold. We thus assume that π *tracks only positive information*, which means that all negative atoms are mapped to the neutral semiring values 0 or 1 which do not affect the evaluation of a formula.

For such interpretations, the classical proof of Hanf's locality theorem can easily be generalised to semiring semantics in fully-idempotent semirings. Indeed, the proof makes use of the Ehrenfeucht-Fraïssé theory by constructing a back-and-forth system, and we know that the soundness of these systems (or, equivalently, of EF-games) holds in fully-idempotent semirings. Semirings with non-idempotent operations, such as the natural semiring $\mathbb{N}$, permit counting by quantification, to the effect that Hanf's locality theorem fails.

The generalisation of Gaifman's theorem is much more involved. We first need an appropriate syntactical notion of Gaifman normal forms. In the classic theory, these are Boolean combinations of basic local sentences of the form $\exists x_1 \dots \exists x_m (\bigwedge_{i \neq j} d(x_i, x_j) > 2r \wedge \bigwedge_i \varphi^{(r)}(x_i))$. Here, $\varphi^{(r)}(x)$ is an r-local formula around x, which means that all quantifiers are relativised to the r-neighbourhood of x, i.e., to elements y with $d(x, y) \leq r$. The property $d(x, y) \leq r$ is not expressible in semiring semantics (at least not by a formula that only takes the values 0 and 1), so instead relativised quantifiers $\exists y{\in}B_r(x)\, \vartheta(x, y)$ and $\forall y{\in}B_r(x)\, \vartheta(x, y)$ are explicitly added to the syntax, where y ranges only over the elements in the r-neighbourhood of x. Moreover, universal basic local sentences (the duals of existential ones) are permitted so that Gaifman normal forms can be brought into negation normal form.

With these preparations, one might expect that the classical proof carries over to fully-idempotent semirings, as for Hanf's theorem, but this is not the case: in contrast to the classical theorem, Gaifman normal forms need not exist for formulae with free variables over any naturally ordered semiring different from $\mathbb{B}$, as witnessed by the formula $\varphi(x) = \exists y(Uy \wedge y \neq x)$ over signature $\tau = \{U\}$ with a single unary relation symbol. However, the main result of [BGN23a] is a positive one: a Gaifman normal form theorem for *sentences* over min-max and lattice semirings. Interestingly, the proof is surprisingly difficult, already for min-max semirings. It follows the strategy of Gaifman's original proof [Gai82] and is essentially a constructive elimination of quantifier alternations (the more modern proof in [EF95] uses characteristic sentences, which are not available in semiring semantics). The reason for the difficulty of this construction is that we have to avoid negation (due to its different behaviour in semiring semantics), whereas Gaifman's original proof uses negation to encode case distinctions in the constructed formulae.

In fact, a careful analysis of the proof reveals that it can be reformulated as a two-step proof. The first step is purely about classical semantics, and shows that every *sentence* has a Gaifman normal form which *does not introduce new negations*. This step fails for formulae with free variables. Indeed, the earlier counterexample $\varphi(x)$ has the classical Gaifman normal form $\exists y \exists z(y \neq z \wedge Uy \wedge Uz) \vee (\neg Ux \wedge \exists y\, Uy)$ which introduces the negation $\neg Ux$ (and this cannot be avoided). The second step then uses separating homomorphisms to lift this result to min-max semirings, and can be slightly adapted to also capture lattice semirings. This proof is an example where new ideas for the generalisation to semiring semantics have lead to a strengthening of the original result in classical semantics.

7.1.2 Compactness

A major obstacle for generalising classical model theory to semiring semantics is the compactness theorem of first-order logic, which is used, for instance, in Fagin's proof of the first-order 0-1 law. In this section, we summarise unpublished joint work with Dawar, Grädel, and Mrkonjić [DGMN24] on possible formulations of compactness in semirings, including compactness theorems for absorptive semirings.

Infinite Interpretations

The proof of the classical compactness theorem involves the construction of an infinite model, and hence compactness, as well as many other model-theoretic results, requires a proper treatment of infinite semiring interpretations. In this thesis, we have laid some of the algebraic foundations by discussing desirable properties of infinite operations as well as appropriate universal semirings (Section 3.4). We refer to [BGMN24] and Mrkonjić's forthcoming thesis [Mrk24] for a more detailed analysis of infinitary operations and their application to provenance on infinite structures. (We should also note that some results on EF-games in [BGM24] include infinite interpretations.)

In the following, we only need infinite operations for absorptive, fully-continuous semirings, which are defined as in Section 3.4.2.

Notions of Compactness

The classical compactness theorem states that a set $\Phi \subseteq \mathrm{FO}(\tau)$ is satisfiable if, and only if, all finite subsets $\Phi_0 \subseteq \Phi$ are satisfiable, and further that the entailment $\Phi \models \psi$ holds for a sentence ψ if, and only if, there is a finite subset $\Phi_0 \subseteq \Phi$ with $\Phi_0 \models \psi$. When we move to semiring interpretations, there are several ways to include semiring values.

One option for the satisfiability statement would be to replace Φ with a set of pairs (φ, k), for a sentence φ and a semiring value $k \in K$, asserting that $\pi[\![\varphi]\!] = k$ (one can also consider $\leq$ or $\geq$ instead of $=$). The goal then is to show that if all finite subsets of pairs can be satisfied by a K-interpretation π, then there is also a K-interpretation that satisfies all pairs simultaneously. It appears to be a challenging open problem whether this "equation system"-like formulation holds for semiring semantics.

Here we take a different approach based on a semantics for entailment. Let K be an absorptive, fully-continuous semiring equipped with infinite operations. For a family $\Phi = (\varphi_i)_{i \in I}$ of FO-sentences, we first define $\pi[\![\Phi]\!] := \pi[\![\bigwedge_{i \in I} \varphi_i]\!]$ as the (infinite) product over the values $\pi[\![\varphi_i]\!]$ of its sentences (we write Φ as a family to account for non-idempotent multiplication). We then define the following notation:

- $\Phi \equiv_K 0$ holds if $\pi[\![\Phi]\!] = 0$ for all model-defining K-interpretations π,
- $\Phi \models_K \psi$ holds if $\pi[\![\Phi]\!] \leq \pi[\![\psi]\!]$ for all model-defining K-interpretations π.

This is motivated by the fact that 1 is the greatest element in absorptive semirings. For instance, in the Viterbi semiring $\mathbb{V} = ([0,1], \max, \cdot, 0, 1)$ we interpret 1 as unanimously true, and values $0 < x < 1$ as true to some limited degree. For $\Phi \models_K \psi$, we thus want that $\pi[\![\Phi]\!] = 1$ implies that also $\pi[\![\psi]\!] = 1$, since a true formula should never entail something "less true". We remark that this relies on absorption, so for non-absorptive semirings such as $\mathbb{N}$, a different semantics of entailment may be more appropriate.

Definition 7.1.1. An absorptive, fully-continuous semiring K has

- *weak compactness*, if for every family Φ with $\Phi \equiv_K 0$, there is a finite subfamily $\Phi_0 \subseteq \Phi$ with $\Phi_0 \equiv_K 0$;
- *strong compactness*, if for every family Φ and sentence ψ with $\Phi \models_K \psi$, there is a finite subfamily $\Phi_0 \subseteq \Phi$ with $\Phi_0 \models_K \psi$.

While the two classical properties are equivalent, here we only have the implication that strong compactness implies weak compactness.

Results for Absorptive Semirings

A first observation is that any absorptive, fully-continuous semiring K has weak compactness. The reason is that due to absorption, K-interpretations are compatible with $\mathbb{B}$-interpretations (i.e., classical structures) in the following sense. As usual, any model-defining K-interpretation π induces a Boolean structure (or $\mathbb{B}$-interpretation) $\mathfrak{A}_\pi$. Conversely, each $\mathbb{B}$-interpretation can be viewed also as a K-interpretation (using the values 0 and 1 in K), without affecting the evaluation of formulae. We can thus apply the classical compactness theorem to conclude weak compactness.

Strong compactness is more complicated and fails in several absorptive semirings. The idea to construct a counterexample is to simulate universal quantification in ψ with an infinite set Φ of all n-fold conjunctions. Over infinite universes, both evaluate to an infinite product, but each finite $\Phi_0 \subseteq \Phi$ evaluates only to a finite product.

More precisely, assume that the semiring K contains an element s with $s^\infty \neq s^n$ for all $n < \omega$. Then let $\vartheta = \exists y\, Py$ and consider $\Phi = \{\bigwedge_{i<n} \vartheta \mid n \geq 1\}$ and $\psi = \forall x\, \vartheta$. One can check that $\Phi \models_K \psi$, but an infinite interpretation with $\pi[\![\vartheta]\!] = s$ witnesses that $\Phi_0 \not\models_K \psi$ for any finite $\Phi_0 \subseteq \Phi$. This example may appear somewhat artificial due to the seemingly superfluous quantifiers and conjunctions, but these become important when multiplication is not idempotent. Suitable elements s exist in the semirings $\mathbb{T}$, $\mathbb{V}$, and $\mathbb{S}^\infty(X)$, and a similar argument yields a counterexample also for the Łukasiewicz semiring $\mathbb{L}$.

These counterexamples are all over infinite semirings, and one can indeed show that strong compactness holds in all finite absorptive semirings. The proof requires some work, but the key idea is to encode a K-interpretation π as a classical structure over an extended signature, as in Section 7.5, so that the classical compactness theorem can be used (see [Mrk24]).

Lifting to Infinite Semirings

For min-max semirings, strong compactness can be lifted from finite to infinite semirings. The idea is to use separating homomorphisms $h\colon K \to \mathbb{B}$ so that a K-interpretation π with $\pi[\![\Phi]\!] > \pi[\![\psi]\!]$ (contradicting $\Phi \models_K \psi$) induces a $\mathbb{B}$-interpretation $h \circ \pi$ with the same property. A hypothetical counterexample in K would then induce a counterexample in $\mathbb{B}$. However, there are two issues with this approach. First, $h \circ \pi$ may not be model-defining, even if π is, and hence may not induce a Boolean structure. Second, this argument requires h to commute with infinitary operations, so that $(h \circ \pi)[\![\varphi]\!] = h(\pi[\![\varphi]\!])$ for all φ and (infinite) interpretations π, but separating homomorphisms with this property do not always exist (consider the min-max semiring over the real interval $[0, 1]$).

We modify this idea as follows. Instead of homomorphisms into $\mathbb{B}$, we consider homomorphisms into a three-element min-max semiring $M_3 := (\{0, \varepsilon, 1\}, \max, \min, 0, 1)$, thereby avoiding the first issue. And instead of requiring h to commute with infinite operations, we use case distinctions on the linear order.

Theorem 7.1.2. *Every (finite or infinite) min-max semiring has strong compactness.*

We know that strong compactness holds for finite min-max semirings (which are absorptive). Let K be an infinite min-max semiring and assume $\Phi \models_K \psi$. Since M_3 can be embedded into K (this is trivial due to the idempotent operations), this implies $\Phi \models_{M_3} \psi$. And since M_3 is finite, strong compactness guarantees a finite $\Phi_0 \subseteq \Phi$ with $\Phi_0 \models_{M_3} \psi$. We can view Φ_0 equivalently as sentence $\bigwedge \Phi_0$, and applying the following lemma yields $\Phi_0 \models_K \psi$, which proves the theorem.

Lemma 7.1.3. *Let K be an infinite min-max semiring, and let $\psi, \vartheta \in \mathrm{FO}$ be sentences. Then $\psi \models_K \vartheta \iff \psi \models_{M_3} \vartheta$.*

Proof. The direction "$\Longrightarrow$" follows by embedding M_3 into K. We prove the other direction by contraposition. Assume that there is a model-defining K-interpretation π with $\pi[\![\psi]\!] > \pi[\![\vartheta]\!]$ and (possibly infinite) universe A. We construct a model-defining M_3-interpretation π_3 with $\pi_3[\![\psi]\!] > \pi_3[\![\vartheta]\!]$ (over the same universe), which proves the lemma. Given an element $s \in K \setminus \{0\}$, we define the following semiring homomorphism:

$$h_{\geq s}\colon K \to M_3, \quad x \mapsto \begin{cases} 1, & x \geq s, \\ \varepsilon, & 0 < x < s, \\ 0, & x = 0. \end{cases}$$

Notice that $h_{\geq s}$ is a semiring homomorphism for any $s \neq 0$, but is not necessarily continuous and may thus not commute with infinite operations. We distinguish two cases.

Case 1: first assume that there is an element $t \in K$ with $\pi[\![\psi]\!] > t > \pi[\![\vartheta]\!]$. We then define $\pi_3 := h_{\geq t} \circ \pi$ and claim that for all $\varphi(\boldsymbol{x}) \in \mathrm{FO}$ and tuples $\boldsymbol{a} \subseteq A$,

(1) if $\pi[\![\varphi(\boldsymbol{a})]\!] > t$, then $\pi_3[\![\varphi(\boldsymbol{a})]\!] = 1$;

(2) if $\pi[\![\varphi(\boldsymbol{a})]\!] < t$, then $\pi_3[\![\varphi(\boldsymbol{a})]\!] \leq \varepsilon$.

We prove this claim by induction on φ. The cases for literals, $\vee$, and $\wedge$ are straightforward, using the fact that K is totally ordered. It remains to consider quantifiers (over infinite universe). We first prove (1). For $\varphi(\boldsymbol{x}) = \exists y\, \eta(x, y)$, assume that $\pi[\![\varphi(\boldsymbol{a})]\!] = \bigsqcup_{b \in A} \pi[\![\eta(\boldsymbol{a}, b)]\!] > t$. Then there is some $b \in A$ with $\pi[\![\eta(\boldsymbol{a}, b)]\!] > t$ and thus $\pi_3[\![\varphi(\boldsymbol{a})]\!] \geq \pi_3[\![\eta(\boldsymbol{a}, b)]\!] = 1$ by induction. For $\varphi(\boldsymbol{x}) = \forall y\, \eta(x, y)$, assume $\pi[\![\varphi(\boldsymbol{a})]\!] = \bigsqcap_{b \in A} \pi[\![\eta(\boldsymbol{a}, b)]\!] > t$. Then $\pi[\![\eta(\boldsymbol{a}, b)]\!] > t$ for all b and the claim again follows by induction on η.

The proof for (2) is analogous. For $\varphi(\boldsymbol{x}) = \exists y\, \eta(x, y)$, assume $\bigsqcup_{b \in A} \pi[\![\eta(\boldsymbol{a}, b)]\!] < t$. Then $\pi[\![\eta(\boldsymbol{a}, b)]\!] < t$ for all b and the claim follows by induction on η. For $\varphi(\boldsymbol{x}) = \forall y\, \eta(x, y)$, assume $\bigsqcap_{b \in A} \pi[\![\eta(\boldsymbol{a}, b)]\!] < t$. There must be $b \in A$ with $\pi[\![\eta(\boldsymbol{a}, b)]\!] < t$. Then $\pi_3[\![\eta(\boldsymbol{a}, b)]\!] \leq \varepsilon$ by induction and the claim follows.

Using the claim, $\pi[\![\psi]\!] > t > \pi[\![\vartheta]\!]$ implies $\pi_3[\![\psi]\!] = 1 > \varepsilon \geq \pi_3[\![\vartheta]\!]$ as intended.

Case 2: no t with $\pi[\![\psi]\!] > t > \pi[\![\vartheta]\!]$ exists. That is, $\pi[\![\psi]\!]$ is the discrete successor of $\pi[\![\vartheta]\!]$. We define $\pi_3 := h_{\geq \pi[\![\psi]\!]} \circ \pi$ and claim that for all $\varphi(\boldsymbol{x}) \in \mathrm{FO}$ and tuples $\boldsymbol{a} \subseteq A$,

(1) if $\pi[\![\varphi(\boldsymbol{a})]\!] \geq \pi[\![\psi]\!]$, then $\pi_3[\![\varphi(\boldsymbol{a})]\!] = 1$;

(2) if $\pi[\![\varphi(\boldsymbol{a})]\!] \leq \pi[\![\vartheta]\!]$, then $\pi_3[\![\varphi(\boldsymbol{a})]\!] \leq \varepsilon$.

The proof is again by induction on φ, with the only interesting cases being the quantifiers. We first prove (1). For $\varphi(\boldsymbol{x}) = \exists y\, \eta(\boldsymbol{x}, y)$, assume $\bigsqcup_{b \in A} \pi[\![\eta(\boldsymbol{a}, b)]\!] \geq \pi[\![\psi]\!] > \pi[\![\vartheta]\!]$. Then there is b with $\pi[\![\eta(\boldsymbol{a}, b)]\!] > \pi[\![\vartheta]\!]$. But then $\pi[\![\eta(\boldsymbol{a}, b)]\!] \geq \pi[\![\psi]\!]$ and the claim follows by induction. For $\varphi(\boldsymbol{x}) = \forall y\, \eta(\boldsymbol{x}, y)$, assume $\bigsqcap_{b \in A} \pi[\![\eta(\boldsymbol{a}, b)]\!] \geq \pi[\![\psi]\!]$. Then $\pi[\![\eta(\boldsymbol{a}, b)]\!] \geq \pi[\![\psi]\!]$ for all b and the claim follows by induction.

The proof for (2) is completely analogous. Applying the claim to ψ and ϑ gives $\pi_3[\![\psi]\!] = 1 > \varepsilon \geq \pi_3[\![\vartheta]\!]$ as intended. □

It remains an open question whether strong compactness holds for the larger class of lattice semirings, or even for fully-idempotent semirings. We already know that separating homomorphisms exist for lattice semirings, but we face the same two issues as for min-max semirings. We refer to [Mrk24] for more details on (continuous) separating homomorphisms for lattice semirings.

7.1.3 Outlook

We have seen that several fundamental results of (finite) model theory can be generalised to semiring semantics, to some extent: they remain true for certain classes of simple semirings, but often fail when arbitrary semirings are considered. While the details depend on the model-theoretic property in question, it seems that much of the classical theory carries over at least to lattice semirings. This is supported by the existence of separating homomorphisms from lattice semirings into the Boolean semiring, which can be used (under some additional assumptions) to lift Boolean equivalences to equivalences over lattice semirings.

The results summarised here are first significant steps towards a model theory for semiring semantics, but of course many questions remain open. For most of the properties, there are counterexamples that establish some boundary for a generalisation to semiring semantics, but there is still a gap to be closed. For instance, Gaifman's theorem holds for lattice semirings and fails in the tropical semiring (cf. [BGN23a]), but the case of fully-idempotent semirings is open; the situation for strong compactness is similar. A different example are the 0-1 laws we study in this chapter, for which no counterexamples are known.

There are also further properties to be investigated. One area are infinite interpretations, which have only recently been studied in more detail. The requirements for infinitary operations lead to challenging algebraic and order-theoretic issues, and while a proposal for an appropriate notion of infinitary semirings has been made in [BGMN24] (cf. Section 3.4), this is not set in stone and deserves further study. Another area are fixed-point logics. We have considered semiring semantics for LFP (Section 5.2), including algorithmic questions for its evaluation (Chapter 4) and the relationship to model-checking games (Section 6.2). What remains open is an analysis of its expressive power, and also of variants with inflationary or partial fixed-point operators. We have hinted at differences in the alternation hierarchy compared to classical semantics (see Section 5.2.5), so this may be an interesting topic for future research. A starting point might be an analysis of the Stage comparison theorem [Kre04] for semiring semantics. A related area is descriptive complexity theory. As suggested in [BGM24, Conclusion], it may by possible to separate FO and LFP with semiring semantics from PTIME by adapting the well-known CFI construction to semirings.

Finally, one can also study modifications to the semantics. We have argued that a weakness of semiring semantics, from a model-theoretic point of view, is its inability to talk about semiring values. Adding this ability would of course lead to significant changes in the expressive power. A less invasive extension would be to only allow comparison with 0, say by writing $\varphi = 0$ or $\varphi \neq 0$. These and similar extensions have been considered in [BHK+23] for team semantics over semirings (cf. Section 5.4), and such additions would alleviate the lack of compositionality of semiring semantics.

7.2 Zero-One Law for Finite Lattice Semirings

7.2.1 Random Semiring Interpretations

Towards the 0-1 laws for first-order and infinitary logic in semiring semantics, we first generalise the concepts of random interpretations and extension properties to semiring semantics. Here we only consider interpretations over finite semirings; an appropriate definition for infinite semirings, such as the Viterbi semiring or the min-max semiring over the real interval $[0, 1]$, is discussed in Section 7.3.

Finite Semirings

For a universe A, a relational vocabulary τ and a commutative semiring K, we denote by $K\text{-Int}_A[\tau]$ the set of K-interpretations $\pi\colon \mathrm{Lit}_A(\tau) \to K$ on universe A. Given a probability measure μ on $K\text{-Int}_A[\tau]$, a sentence $\psi \in L^{\omega}_{\infty\omega}(\tau)$, and a set of semiring values $J \subseteq K$ let

$$\mu\big[\pi[\![\psi]\!] \in J\big] \;:=\; \mu\big\{\pi \in K\text{-Int}_A[\tau] \;\big|\; \pi[\![\psi]\!] \in J\big\}.$$

The probability measures $\mu_{n,p}$ we are interested in are obtained by choosing semiring values for the literals over the universe $A = [n]$ independently and at random, keeping in mind that for complementary literals α and $\neg\alpha$ precisely one should get the value 0, and the other one an arbitrary nonzero value. Given a probability distribution p on $K^+ := K \setminus \{0\}$, a random K-interpretation π thus makes, independently for each relational atom $\alpha = R\boldsymbol{a}$, a random choice with probability $\frac{1}{2}$ whether α or $\neg\alpha$ shall be true; if α is true, then set

$\pi(\neg\alpha) = 0$ and select for $\pi(\alpha)$ a random value from K^+ according to p; analogously, if α is false, then we set $\pi(\alpha) = 0$ and select $\pi(\neg\alpha) \in K^+$ at random. Every K-interpretation π chosen in this way is model-defining.

For finite semirings K, the most natural probability distribution on K^+ is the uniform one, so that the probability that $\pi(\alpha) = j$ is $\frac{1}{2}(|K| - 1)$ for any $j \neq 0$. But our results hold for all measures $\mu_{n,p}$ as long as the choices whether α or $\neg\alpha$ should receive a positive value are done with a constant probability (not necessarily $\frac{1}{2}$) and all semiring values occur with positive probability, i.e., $p\colon K^+ \to (0, 1]$. For fixed p, ψ and $j \in K$ we then consider the sequence $(\mu_{n,p}[\pi[\![\psi]\!] = j])_{n<\omega}$ of probabilities that ψ evaluates to the semiring value j in a random K-interpretation on universe $[n]$ (with positive semiring values chosen according to the probability distribution p). This setting can also be applied to countably infinite semirings K.

Definition 7.2.1. We say that a *0-1 law* holds for a class of sentences $\Phi \subseteq L^{\omega}_{\infty\omega}$, a finite or countable semiring K, and a probability distribution p on K^+, if for each sentence $\psi \in \Phi$ and each value $j \in K$ the sequence $(\mu_{n,p}[\pi[\![\psi]\!] = j])_{n<\omega}$ converges to either 0 or 1, as n goes to infinity.

In that case, we denote by $\mathrm{ASV}_{K,p}(\psi)$ the *almost sure valuation* of ψ for K and p, defined as the unique value $j \in K$ such that $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] = j] = 1$. We further write $\mathrm{ASV}_{K,p}(\Phi) := \{\mathrm{ASV}_{K,p}(\psi) \mid \psi \in \Phi\}$ for the set of almost sure valuations of sentences in Φ.

Extension Properties

Similar to the Boolean case, we study configurations of k points (which we always assume to be pairwise distinct) and whether they can be extended to $k + 1$ points. Recall that $\mathrm{Lit}_A(\tau)$ denotes the set of *instantiated* literals over the universe A, i.e., $R\boldsymbol{a}$ and $\neg R\boldsymbol{a}$ for $R \in \tau$ and $\boldsymbol{a} \subseteq A$. When we work with the finite-variable fragments $\mathrm{FO}^k(\tau)$ and $L^k_{\infty\omega}(\tau)$, we always assume that variables are named $x_1, x_2, \dots, x_k$, and we write $\mathrm{Lit}_k(\tau)$ for the set of (uninstantiated) literals $R\boldsymbol{z}$, $\neg R\boldsymbol{z}$ with $R \in \tau$ and $\boldsymbol{z} \subseteq \{x_1, \dots, x_k\}$.

Definition 7.2.2. An (atomic) *k-type* (of vocabulary τ in the semiring K) is a consistent valuation $\rho\colon \mathrm{Lit}_k(\tau) \to K$, consistent in the sense that for every τ-atom β, precisely one of $\rho(\beta)$, $\rho(\neg\beta)$ is 0.

Given a K-interpretation $\pi\colon \mathrm{Lit}_A(\tau) \to K$, each tuple $\boldsymbol{a} = (a_1, \dots, a_k)$ of pairwise distinct elements induces the k-type $\rho^\pi_{\boldsymbol{a}}\colon \mathrm{Lit}_k(\tau) \to K$ that maps all literals $\beta \in \mathrm{Lit}_k(\tau)$ to $\pi(\beta[\boldsymbol{a}])$, where $\beta[\boldsymbol{a}]$ results from β by instantiating each variable x_i by a_i.

For two k-types $\rho, \rho'\colon \mathrm{Lit}_k(\tau) \to K$, we write $\rho =_{\mathbb{B}} \rho'$ if ρ and ρ' map precisely the same literals to 0 (thus inducing the same Boolean type when identifying all nonzero values). Further let $\rho \leq \rho'$ if $\rho(\beta) \leq \rho(\beta')$ for all $\beta \in \mathrm{Lit}_k(\tau)$. Notice that this can only be the case if $\rho =_{\mathbb{B}} \rho'$; indeed if $0 = \rho(\beta) < \rho'(\beta) \neq 0$ then $0 \neq \rho(\neg\beta) \not\leq \rho'(\neg\beta) = 0$.

We say that a tuple $\boldsymbol{a}$ of pairwise distinct elements *realises the atomic k-type* ρ in π, if $\rho^\pi_{\boldsymbol{a}} = \rho$. For $k > m$ an atomic k-type ρ^+ *extends* the atomic m-type ρ if $\rho^+ \restriction \mathrm{Lit}_m(\tau) = \rho$. In that case, every realisation $(a_1, \dots, a_k)$ of ρ^+ in π restricts to a realisation $(a_1, \dots, a_m)$ of ρ. On the other side, it is not clear whether a tuple that realises ρ can be extended to a realisation of ρ^+. We formulate extension properties that guarantee the existence of such extensions, which play a central role in the proof of 0-1 laws. Given an atomic m-type ρ, let $\mathrm{ext}(\rho)$ be the set of atomic $(m+1)$-types that extend ρ.

Definition 7.2.3. A K-interpretation π has the *k-extension property* if for every $m < k$, every atomic m-type ρ and every extension $\rho^+ \in \mathrm{ext}(\rho)$, the following holds: every tuple $\boldsymbol{a}$ that realises ρ in π can be extended to a realisation $(\boldsymbol{a}, b)$ of ρ^+, for some $b \in A \setminus \boldsymbol{a}$.

We remark that, unlike in Fagin's proof of the classical 0-1 law, it does not seem possible to express the k-extension property as a first-order formula. Instead, our proof only needs the following observation.

Lemma 7.2.4. *Fix a finite relational vocabulary τ and let K be a countable semiring with a probability distribution $p\colon K^+ \to (0,1]$. For every atomic m-type ρ and every extension $\rho^+ \in \mathrm{ext}(\rho)$,*

$$\lim_{n\to\infty} \mu_{n,p}[\text{ every realisation of } \rho \text{ in } \pi \text{ can be extended to a realisation of } \rho^+] = 1,$$

and the convergence to this limit is exponentially fast. For any finite semiring, we moreover have, again with exponential convergence, that random K-interpretations almost surely have the k-extension property (for any fixed k).

Proof. We first calculate, for any given $(m+1)$-type ρ^+ and its restriction $\rho = \rho^+ \upharpoonright \mathrm{Lit}_m(\tau)$, a bound for the probability that a random K-interpretation on n elements has some realisation $\boldsymbol{a}$ of ρ that can *not* be extended to a realisation $(\boldsymbol{a}, b)$ of ρ^+. There is a fixed collection $\alpha_1, \dots, \alpha_q$ of relational atoms in $\mathrm{Lit}_{m+1}(\tau)$ in which the variable x_{m+1} occurs; hence there is a fixed collection $(s_1, r_1), \dots, (s_q, r_q)$ of elements of $\{\bot, \top\} \times K^+$, where

$$(s_i, r_i) = \begin{cases} (\top, j), & \text{if } \rho^+(\alpha_i) = j \text{ and } \rho^+(\neg\alpha) = 0, \\ (\bot, j), & \text{if } \rho^+(\alpha_i) = 0 \text{ and } \rho^+(\neg\alpha) = j. \end{cases}$$

Thus, the probability that all values chosen by a random K-interpretations coincide with those required by ρ^+ is

$$f(\rho^+) := 2^{-q} \prod_{i=1}^{q} p(r_i) > 0.$$

Thus, for any given realisation $\boldsymbol{a} = (a_1, \dots, a_m)$ of ρ, the probability that a fixed $b \in [n] \setminus \{a_1, \dots, a_m\}$ does *not* provide a realisation $(\boldsymbol{a}, b)$ of ρ^+ is $(1 - f(\rho^+))$. It follows that

$$\mu_{n,p}[\text{ some realisation of } \rho \text{ does not extend to a realisation of } \rho^+] \le n^m (1 - f(\rho^+))^{n-m}$$

which for growing n converges to 0 exponentially fast.

Over an infinite semiring there exist infinitely many atomic k-types for any $k \ge 1$, so we cannot realise all of them on a finite universe. Thus $\mu_{n,p}[\pi \text{ has the } k\text{-extension property }] = 0$ for all n. However, over a finite semiring, each k admits only a bounded number of atomic k-types, and we conclude that $\lim_{n\to\infty} \mu_{n,p}[\pi \text{ has the } k\text{-extension property }] = 1$. ☐

7.2.2 Polynomials for Almost Sure Valuations

The evaluation of a first-order sentence in semiring semantics on a finite universe A can be described by a polynomial in the most general provenance semiring $\mathbb{N}[X]$, for a suitably large set X (see Section 5.1 and [GT17a]). More precisely, each sentence $\psi \in \mathrm{FO}(\tau)$ induces a polynomial $f_\psi^A \in \mathbb{N}[\boldsymbol{X}^{(A)}]$, where $\boldsymbol{X}^{(A)}$ contains two indeterminates $X_{R\boldsymbol{a}}$ and $X_{\neg R\boldsymbol{a}}$ for each instantiated atom $R\boldsymbol{a} \in \mathrm{Lit}_A(\tau)$. This polynomial describes the evaluation of ψ in the following sense: for any K-interpretation π, we have $\pi[\![\psi]\!] = f_\psi^A[\pi]$, where $f_\psi^A[\pi]$ results from f_ψ^A by substituting X_β by $\pi(\beta)$, for every literal $\beta \in \mathrm{Lit}_A(\tau)$.

Clearly, the set $\boldsymbol{X}^{(A)}$, and hence the polynomial f_ψ^A, depends on A. We shall prove that for semiring interpretations with the k-extension property, we can do better, by constructing polynomials f_ψ that depend only on ψ. Since these describe the evaluation of ψ on any finite interpretation with the k-extension property, they also describe the asymptotic almost sure valuations and thus allow us to derive a 0-1 law.

Excluding Quantifiers

For the following proofs, some minor headaches in the form of necessary case distinctions can be caused by equalities and inequalities. To simplify our reasoning, we rewrite formulae by means of the *excluding quantifiers* $\exists^{\neq}$ and $\forall^{\neq}$, with the equivalences (in Boolean as well as semiring semantics) that for any formula $\varphi(\boldsymbol{x}, y)$, with $\boldsymbol{x} = (x_1, \dots, x_k)$,

$$\exists^{\neq} y\, \varphi(\boldsymbol{x}, y) \equiv \exists y (\bigwedge_{i=1}^{k} y \neq x_i \wedge \varphi(\boldsymbol{x}, y)), \qquad \forall^{\neq} y\, \varphi(\boldsymbol{x}, y) \equiv \forall y (\bigvee_{i=1}^{k} y = x_i \vee \varphi(\boldsymbol{x}, y)).$$

Clearly, the classical quantifiers $\exists$ and $\forall$ can be expressed (again in Boolean as well as semiring semantics) by $\exists y\, \varphi(\boldsymbol{x}, y) \equiv \bigvee_{i=1}^{k} \varphi(\boldsymbol{x}, x_i) \vee \exists^{\neq} y\, \varphi(\boldsymbol{x}, y)$ and $\forall y\, \varphi(\boldsymbol{x}, y) \equiv \bigwedge_{i=1}^{k} \varphi(\boldsymbol{x}, x_i) \wedge \forall^{\neq} y\, \varphi(\boldsymbol{x}, y)$. We thus define polynomials only for formulae with these excluding quantifiers.

The Polynomials f_ψ

For any natural number i, let $\boldsymbol{X}^{(i)}$ be the set of indeterminates X_α and $X_{\neg\alpha}$ for τ-atoms $\alpha = R\boldsymbol{z} \in \mathrm{Lit}_i(\tau)$ using only variables $x_1, \dots, x_i$. Notice that $\boldsymbol{X}^{(i)}$ depends only on i and τ, but not on the universe. Further, let $E = (\{0, e, 1\}, \max, \min, 0, 1)$ be the min-max semiring with the order $0 < e < 1$. We describe any formula $\psi(x_1, \dots, x_i)$ with $i \leq k$ by a formal polynomial $f_\psi \in \mathrm{AbsPoly}(E, \mathbb{B}, \boldsymbol{X}^{(i)})$, independent of the size of the universe on which we evaluate ψ. Recall that $\mathrm{AbsPoly}(E, \mathbb{B}, \boldsymbol{X}^{(i)})$ is the finite semiring of absorptive polynomials with coefficients in E, exponents in $\mathbb{B}$ (i.e., without exponents), and with indeterminates in $\boldsymbol{X}^{(i)}$ (see Section 3.7). As usual, we write such a polynomial as $f_\psi = c_1 m_1 + \dots + c_l m_l$, with coefficients $c_i \in \{0, e, 1\}$ and monomials $m_i \colon \boldsymbol{X}^{(i)} \to \mathbb{B}$.

Definition 7.2.5. Let $\psi(x_1, \dots, x_i) \in L^\omega_{\infty\omega}(\tau)$ for a finite relational vocabulary τ. Recall that we assume that ψ is in negation normal form and written with the excluding quantifiers $\exists^{\neq}$ and $\forall^{\neq}$. The associated polynomial $f_\psi(\boldsymbol{X}^{(i)})$ is defined by induction, as follows.

- If ψ is an equality $x_j = x_\ell$ then $f_\psi := 1$ if $j = \ell$ and $f_\psi := 0$ if $j \neq \ell$. Similarly, if ψ is an inequality $x_j \neq x_\ell$ then $f_\psi := 1$ if $j \neq \ell$ and $f_\psi := 0$ if $j = \ell$.
- If ψ is a (positive or negative) literal β, then $f_\psi := X_\beta$.

- For $\psi = \bigvee \Phi$, we set $f_\psi := \widehat{\sum}_{\varphi \in \Phi} f_\varphi$, and analogously $f_{\bigwedge \Phi} := \widehat{\prod}_{\varphi \in \Phi} f_\varphi$.
- Consider $\psi(\boldsymbol{x}) = \exists^{\neq} y\, \varphi(\boldsymbol{x}, y)$ and assume w.l.o.g. that $y = x_{i+1}$. For the inner formula, we have a polynomial f_φ with indeterminates in $\boldsymbol{X}^{(i+1)}$ which we write as $f_\varphi(\boldsymbol{X}^{(i)}, \boldsymbol{Y}^{(i)})$, where $\boldsymbol{Y}^{(i)} = \boldsymbol{X}^{(i+1)} \setminus \boldsymbol{X}^{(i)}$. Let S be the set of all consistent selector functions $s \colon \boldsymbol{Y}^{(i)} \to \{0, 1\}$, consistent in the sense that precisely one of X_α, $X_{\neg\alpha}$ is mapped to 0, for all τ-atoms α. Now set $f_\psi(\boldsymbol{X}^{(i)}) := \sum_{s \in S} f_\varphi(\boldsymbol{X}^{(i)}, s(\boldsymbol{Y}^{(i)}))$.
- Finally consider $\psi(\boldsymbol{x}) = \forall^{\neq} y\, \varphi(\boldsymbol{x}, y)$ with $y = x_{i+1}$ and again write $f_\varphi(\boldsymbol{X}^{(i)}, \boldsymbol{Y}^{(i)})$ as above. Let S be the set of all consistent selector functions $s \colon \boldsymbol{Y}^{(i)} \to \{0, e\}$ and set $f_\psi(\boldsymbol{X}^{(i)}) := \prod_{s \in S} f_\varphi(\boldsymbol{X}^{(i)}, s(\boldsymbol{Y}^{(i)}))$.

Remark 7.2.6. Infinitary sums and products in $\mathrm{AbsPoly}(E, \mathbb{B}, \boldsymbol{X}^{(i)})$ are simply suprema and infima, which are trivially well-defined since $\mathrm{AbsPoly}(E, \mathbb{B}, \boldsymbol{X}^{(i)})$ is a finite lattice semiring. We use the notation with sums and products to distinguish between the polynomials and the operations of the lattice semiring K. Since the infinitary operations are idempotent, we use the simpler notation $\bigvee \Phi$ instead of $\bigvee_{i \in I} \varphi_i$ (and analogously for $\bigwedge \Phi$).

Remark 7.2.7. In the original proof of the 0-1 law for first-order logic in semiring semantics [GHNW22a], we used polynomials in $E[\boldsymbol{X}^{(i)}]$. Here, absorptive polynomials $\mathrm{AbsPoly}(E, \mathbb{B}, \boldsymbol{X}^{(i)})$ are used instead, for two reasons. Since we want to evaluate the polynomials in lattice semirings, it suffices to work with this simpler, in particular finite, semiring. Moreover, extending the proof to infinitary logic would require more involved arguments with power series instead of polynomials $E[\boldsymbol{X}^{(i)}]$, but no modification is necessary when working with $\mathrm{AbsPoly}(E, \mathbb{B}, \boldsymbol{X}^{(i)})$.

Example 7.2.8. Consider $\psi := \exists^{\neq} x(\neg Exx \wedge \forall^{\neq} y(Exy \vee (\neg Exy \wedge \exists^{\neq} z(Exz \wedge Ezy))))$, an FO^3-sentence defining the directed graphs that contain some centre from which all nodes are reachable in one or two steps. For ease of notation we abbreviate the indeterminates associated with the atoms as $X := X_{Exx}$, $Y := X_{Exy}$, $Z := X_{Exz}$ and $U := X_{Ezy}$, as well as $\overline{X}, \overline{Y}, \overline{Z}$ and $\overline{U}$ for the corresponding negated atoms. The following table describes the polynomials f_φ for the subformulae of ψ.

φ	f_φ
$Exz \wedge Ezy$	ZU
$\exists^{\neq} z(Exz \wedge Ezy)$	1
$\neg Exy \wedge \exists^{\neq} z(Exz \wedge Ezy)$	$\overline{Y}$
$Exy \vee (\neg Exy \wedge \exists^{\neq} z(Exz \wedge Ezy))$	$Y + \overline{Y}$
$\forall^{\neq} y(Exy \vee (\neg Exy \wedge \exists^{\neq} z(Exz \wedge Ezy))))$	e
$\neg Exx \wedge \forall^{\neq} y(Exy \vee (\neg Exy \wedge \exists^{\neq} z(Exz \wedge Ezy))))$	$e\overline{X}$
ψ	e

We remark that the classically equivalent sentence ψ' obtained from ψ by omitting the literal $\neg Exy$ is instead described by $f_{\psi'} = 1$. As we shall prove, this means that the almost sure valuation of ψ in a min-max semiring is the smallest positive value ε, whereas the almost sure valuation of ψ' is the largest value 1.

ψ	f_ψ
$x_i = x_j$	1 or 0 (depending on $i = j$)
$R\boldsymbol{z}$, $\neg R\boldsymbol{z}$	$X_{R\boldsymbol{z}}$, $X_{\neg R\boldsymbol{z}}$
$\varphi \vee \vartheta$, $\varphi \wedge \vartheta$	$f_\varphi + f_\vartheta$, $f_\varphi \cdot f_\vartheta$
$\bigvee \Phi$, $\bigwedge \Phi$	$\widehat{\sum}_{\varphi\in\Phi} f_\varphi$, $\widehat{\prod}_{\varphi\in\Phi} f_\varphi$
$\exists^{\neq} y\ \varphi(\boldsymbol{x}, y)$	$\sum_{s\in S} f_\varphi(\boldsymbol{X}^{(i)}, s(\boldsymbol{Y}^{(i)}))$, with consistent assignments $s\colon \boldsymbol{Y}^{(i)} \to \{0, 1\}$
$\forall^{\neq} y\ \varphi(\boldsymbol{x}, y)$	$\prod_{s\in S} f_\varphi(\boldsymbol{X}^{(i)}, s(\boldsymbol{Y}^{(i)}))$, with consistent assignments $s\colon \boldsymbol{Y}^{(i)} \to \{0, e\}$

Figure 7.3: Construction of the polynomial f_ψ (Definition 7.2.5).

We next observe that polynomials $f \in \mathrm{AbsPoly}(E, \mathbb{B}, \boldsymbol{X}^{(k)})$, with indeterminates X_β for literals $\beta \in \mathrm{Lit}_k(\tau)$, are evaluated to semiring values $f[\rho] \in K$ by atomic k-types $\rho\colon \mathrm{Lit}_k(\tau) \to K$, for any semiring K with a distinguished element ε (which we will define later). Indeed, ρ defines a unique homomorphism $h^\varepsilon_\rho\colon \mathrm{AbsPoly}(E, \mathbb{B}, \boldsymbol{X}^{(k)}) \to K$, induced by $h^\varepsilon_\rho(e) := \varepsilon$ and $h^\varepsilon_\rho(X_\beta) := \rho(\beta)$. We put $f[\rho] := h^\varepsilon_\rho(f)$ and remark that by monotonicity of absorptive polynomials, we have $f[\rho] \le f[\rho']$ whenever $\rho \le \rho'$.

7.2.3 Proof of the Zero-One Law

We now use the polynomials f_ψ to obtain a first 0-1 law for finite lattice semirings. In such semirings, we define $\varepsilon_K := \bigsqcap\{j \in K \mid j \neq 0\}$ as the smallest positive element, if such an element exists (otherwise $\varepsilon_K = 0$). In finite min-max semirings, we always have $\varepsilon_K > 0$. The key step is to prove that for K-interpretations into finite lattice semirings with the k-extension property, the polynomials f_ψ constructed in Definition 7.2.5 provide a concise and adequate description of any infinitary formula $\psi(\boldsymbol{x}) \in L^k_{\infty\omega}$.

Theorem 7.2.9. *Let $(K, \sqcup, \sqcap, 0, 1)$ be a finite lattice semiring, τ a finite relational vocabulary and $k < \omega$. Then, for every K-interpretation $\pi\colon \mathrm{Lit}_A(\tau) \to K$ with the k-extension property, every formula $\psi(x_1, \dots, x_i) \in L^k_{\infty\omega}(\tau)$ and every tuple $\boldsymbol{a} = (a_1, \dots, a_i)$ of pairwise distinct elements of A, we have that $\pi[\![\psi(\boldsymbol{a})]\!] = f_\psi[\rho^\pi_{\boldsymbol{a}}]$.*

Proof. We proceed by induction on ψ. If ψ is a literal, the claim is immediate from the definition of f_ψ. For infinite conjunctions and disjunctions, notice that infinite suprema and infima are well-defined in K due to its finiteness. Moreover, recall that the evaluation of a polynomial in $\mathrm{AbsPoly}(E, \mathbb{B}, \boldsymbol{X}^{(i)})$ for a k-type ρ commutes with arbitrary suprema and infima (by Corollary 3.7.14). So for $\psi = \bigvee \Phi$, we have

$$\pi[\![\psi(\boldsymbol{a})]\!] = \bigsqcup_{\varphi\in\Phi} \pi[\![\varphi(\boldsymbol{a})]\!] = \bigsqcup_{\varphi\in\Phi} f_\varphi[\rho^\pi_{\boldsymbol{a}}] = \Big(\widehat{\sum_{\varphi\in\Phi}} f_\varphi\Big)[\rho^\pi_{\boldsymbol{a}}] = f_\psi[\rho^\pi_{\boldsymbol{a}}]$$

by induction, and analogously for disjunctions.

Let now $\psi(\boldsymbol{x}) = \exists^{\neq} y\, \varphi(\boldsymbol{x}, y)$ and w.l.o.g. $y = x_{i+1}$. Recall that $f_\psi(\boldsymbol{X}^{(i)})$ is defined as $\sum_{s\in S} f_\varphi(\boldsymbol{X}^{(i)}, s(\boldsymbol{Y}^{(i)}))$, where $\boldsymbol{Y}^{(i)} = \boldsymbol{X}^{(i+1)} \setminus \boldsymbol{X}^{(i)}$ and S is the set of consistent selector functions $s\colon \boldsymbol{Y}^{(i)} \to \{0,1\}$. By induction,

$$\pi[\![\psi(\boldsymbol{a})]\!] = \bigsqcup_{b\in A\setminus \boldsymbol{a}} \pi[\![\varphi(\boldsymbol{a},b)]\!] = \bigsqcup_{b\in A\setminus \boldsymbol{a}} f_\varphi[\rho^\pi_{\boldsymbol{a},b}]. \qquad (*)$$

We first prove that $f_\psi[\rho^\pi_{\boldsymbol{a}}]$ is an upper bound for $\pi[\![\psi(\boldsymbol{a})]\!]$. For every $b \in A \setminus \boldsymbol{a}$, define the selector function s_b by $s_b(X_\beta) = 1$ if $\rho^\pi_{\boldsymbol{a},b}(\beta) \neq 0$ (and $s_b(X_\beta) = 0$ otherwise), for every literal $\beta \in \mathrm{Lit}_{i+1}(\tau) \setminus \mathrm{Lit}_i(\tau)$. We simplify notation and write $f_\varphi[\rho^\pi_{\boldsymbol{a}}, s_b(\boldsymbol{Y}^{(i)})]$ instead of $f_\varphi(\boldsymbol{X}^{(i)}, s_b(\boldsymbol{Y}^{(i)}))[\rho_{\boldsymbol{a}}]$. Since 1 is the largest semiring value, we have $f_\varphi[\rho^\pi_{\boldsymbol{a},b}] \leq f_\varphi[\rho^\pi_{\boldsymbol{a}}, s_b(\boldsymbol{Y}^{(i)})]$ by monotonicity. Hence $\pi[\![\psi(\boldsymbol{a})]\!] \leq \bigsqcup_{s\in S} f_\psi[\rho^\pi_{\boldsymbol{a}}, s(\boldsymbol{Y}^{(i)})] = f_\psi[\rho^\pi_{\boldsymbol{a}}]$ by $(*)$.

The other direction holds by the extension property. Every selector function $s \in S$ induces an extension $\rho_s \in \mathrm{ext}(\rho^\pi_{\boldsymbol{a}})$ with $\rho_s(\beta) = s(X_\beta)$ for $\beta \in \mathrm{Lit}_{i+1}(\tau) \setminus \mathrm{Lit}_i(\tau)$. Since π has the k-extension property, there is $b_s \in A \setminus \boldsymbol{a}$ with $\rho^\pi_{\boldsymbol{a},b_s} = \rho_s$, hence $f_\varphi[\rho^\pi_{\boldsymbol{a}}, s(\boldsymbol{Y}^{(i)})] = f_\varphi[\rho^\pi_{\boldsymbol{a},b_s}]$. As this holds for all s, we have $f_\psi[\rho^\pi_{\boldsymbol{a}}] \leq \pi[\![\psi(\boldsymbol{a})]\!]$ by $(*)$ and thus equality.

Finally let $\psi(\boldsymbol{x}) = \forall^{\neq} y\, \varphi(\boldsymbol{x}, y)$. Recall that $f_\psi(\boldsymbol{X}^{(i)})$ is defined as $\prod_{s\in S} f_\varphi(\boldsymbol{X}^{(i)}, s(\boldsymbol{Y}^{(i)}))$, where this time we consider selector functions $s\colon \boldsymbol{Y}^{(i)} \to \{0, e\}$ instead of $\{0,1\}$. We again have $\pi[\![\psi(\boldsymbol{a})]\!] = \bigsqcap_{b\in A\setminus\boldsymbol{a}} f_\varphi[\rho^\pi_{\boldsymbol{a},b}]$ by induction. Since ε_K is the smallest positive semiring value (or 0), we first observe that, completely analogous to the previous case, $f_\psi[\rho^\pi_{\boldsymbol{a}}]$ is a lower bound for $\pi[\![\psi(\boldsymbol{a})]\!]$. If $\varepsilon_K > 0$, then the other direction is analogous as well: for each $s \in S$, define ρ_s by $\rho_s(\beta) = \varepsilon_K$ if $s(X_\beta) = e$, and $\rho_s(\beta) = 0$ if $s(X_\beta) = 0$ (recall that e becomes ε_K when evaluating f_ψ); this extension is realised by the extension property.

It remains to prove $f_\psi[\rho^\pi_{\boldsymbol{a}}] \geq \pi[\![\psi(\boldsymbol{a})]\!]$ in the case $\varepsilon_K = 0$ (defining ρ_s by setting $\rho_s(\beta) = \varepsilon_k$ or $\rho_s(\beta) = 0$ would not be consistent). Recall that min-max semirings always have $\varepsilon_K > 0$, so this case only happens for lattice semirings where the underlying order is not total. Let $\min(K)$ be the set of minimal nonzero elements of K. Observe that $|\min(K)| \geq 2$ and $\bigsqcap \min(K) = 0$, as K is finite and $\varepsilon_K = 0$. Let R be the set of extensions $\rho \in \mathrm{ext}(\rho^\pi_{\boldsymbol{a}})$ such that $\rho(\beta) = 0$ or $\rho(\beta) \in \min(K)$, for all $\beta \in \mathrm{Lit}_{i+1}(\tau) \setminus \mathrm{Lit}_i(\tau)$. By the k-extension property, all $\rho \in R$ are realised by some $b \in A \setminus \boldsymbol{a}$, hence $\bigsqcap_{\rho\in R} f_\varphi[\rho] \geq \bigsqcap_{b\in A\setminus\boldsymbol{a}} f_\varphi[\rho^\pi_{\boldsymbol{a},b}] = \pi[\![\psi(\boldsymbol{a})]\!]$.

Now consider f_ψ. As we evaluate e to $\varepsilon_K = 0$, the selector function s does not matter and we have $f_\psi[\rho^\pi_{\boldsymbol{a}}] = f_\varphi[\rho^\pi_{\boldsymbol{a}}, 0]$ (that is, we map all variables $X_\alpha, X_{\neg\alpha} \in \boldsymbol{Y}^{(i)}$ to 0, ignoring the usual consistency requirement). We claim that $f_\varphi[\rho^\pi_{\boldsymbol{a}}, 0] = \bigsqcap_{\rho\in R} f_\varphi[\rho]$. To see this, we write $f_\varphi = g + h$ or, more precisely, $f_\varphi(\boldsymbol{X}^{(i)}, \boldsymbol{Y}^{(i)}) = g(\boldsymbol{X}^{(i)}, \boldsymbol{Y}^{(i)}) + h(\boldsymbol{X}^{(i)})$, where g contains all the monomials of f_φ that contain any $X_\beta \in \boldsymbol{Y}^{(i)}$, and h the remaining ones. Recall that when we evaluate $f_\psi = g + h$, we interpret addition by the semiring operation $\sqcup$. Since lattice semirings are distributive and R finite, we have

$$\bigsqcap_{\rho\in R} f_\varphi[\rho] = \bigsqcap_{\rho\in R} (g[\rho] \sqcup h[\rho^\pi_{\boldsymbol{a}}]) = \Big(\bigsqcap_{\rho\in R} g[\rho]\Big) \sqcup h[\rho^\pi_{\boldsymbol{a}}].$$

Now consider any minimal element $\bot \in \min(K)$ and some type $\rho \in R$ with $\rho(\beta) \in \{0, \bot\}$ for all $\beta \in \mathrm{Lit}_{i+1}(\tau) \setminus \mathrm{Lit}_i(\tau)$. By definition, each monomial m of g contains an indeterminate X_β for some $\beta \in \mathrm{Lit}_{i+1}(\tau) \setminus \mathrm{Lit}_i(\tau)$, so $m[\rho] \leq \bot$ (recall that multiplication is $\sqcap$). Hence $g[\rho] \leq \bot$. As this holds for all $\bot \in \min(K)$, we have shown $\bigsqcap_{\rho\in R} g[\rho] \leq \bigsqcap \min(K) = 0$. It follows that $\bigsqcap_{\rho\in R} f_\varphi[\rho] = h[\rho^\pi_{\boldsymbol{a}}] = f_\varphi[\rho^\pi_{\boldsymbol{a}}, 0]$ as claimed. $\square$

Applying this result to sentences implies a 0-1 law for $L^{\omega}_{\infty\omega}$ on finite lattice semirings, and also for LFP via the embedding into $L^{\omega}_{\infty\omega}$.

Corollary 7.2.10 (0-1 law on finite lattice semirings)**.** *Let K be a finite lattice semiring with a probability distribution $p\colon K^{+} \to (0,1]$, and let τ be a finite relational vocabulary.*

For every sentence $\psi \in L^{\omega}_{\infty\omega}(\tau) \cup \mathrm{LFP}(\tau)$ and $j \in K$, the sequence $(\mu_{n,p}[\pi[\![\psi]\!] = j])_{n<\omega}$ converges exponentially fast to either 0 or 1, as n goes to infinity. The only possible almost sure valuations of sentences are $\mathrm{ASV}_{K,p}(L^{\omega}_{\infty\omega}(\tau) \cup \mathrm{LFP}(\tau)) = \{0, 1, \varepsilon_K\}$.

Proof. Notice that K is absorptive and trivially fully continuous. Due to the embedding of LFP in Theorem 5.3.5, it suffices to consider the case $\psi \in L^{\omega}_{\infty\omega}(\tau)$.

Fix k such that $\psi \in L^{k}_{\infty\omega}(\tau)$. By Lemma 7.2.4, the probability that a random K-interpretation π on $[n]$ has the k-extension property converges to 1 exponentially fast, as n goes to infinity. But on K-interpretations with the k-extension property, ψ is described by a polynomial f_ψ. Since ψ has no free variables, we have that either $f_\psi = 0$, or $f_\psi = 1$, or $f_\psi = e$, and the atomic type to consider is the trivial empty type $\emptyset$, which implies that $f_\psi[\emptyset]$ is either 0, or 1, or ε_K. By applying Theorem 7.2.9, we conclude that the probabilities $\mu_{n,p}\big[\pi[\![\psi]\!] = f_\psi[\emptyset]\big]$ converge to 1 exponentially fast. □

Notice that, as in the Boolean case, the 0-1 law does not extend to arbitrary formulae with free variables. Indeed for an atomic formula, say Exy, any value $j \in K^{+}$ and any fixed pair of constants $k, \ell \in \mathbb{N}$, we have that $\lim_{n\to\infty} \mu_{n,p}[E(k,\ell) = j] = \frac{1}{2}p(j)$, which is in general not 0 or 1. Nevertheless we can extend the 0-1 law to formulae $\psi(\boldsymbol{x})$ with free variables, with the additional constraint that every relational atom contains a quantified variable; this implies that f_ψ is either 0, 1 or e.

Corollary 7.2.11. *Let K, p, τ be as in Corollary 7.2.10. Let Φ be the set of fully instantiated first-order sentences $\psi(\boldsymbol{a})$ where $\psi(x_1, \dots, x_i)$ is a formula in $L^{\omega}_{\infty\omega}(\tau)$ with free variables $x_1, \dots, x_i$, in which every relational atom contains a quantified variable, and $\boldsymbol{a} = (a_1, \dots, a_i)$ is a tuple of distinct natural numbers, i.e. of elements of all universes $[n]$ for large enough n. Then the 0-1 law holds for $K, p,$ and Φ, and $\mathrm{ASV}_{K,p}(\Phi) = \{0, 1, \varepsilon_K\}$.*

Corollary 7.2.10 splits the relational sentences into three classes, according to whether their valuations in finite lattice semirings are almost surely 0, 1, or ε_K. This split is the same for all finite lattice semirings, since it just depends on the associated polynomial f_ψ. The classical 0-1 law for infinitary (and first-order) logic, saying that every relational sentence $\psi \in L^{\omega}_{\infty\omega}$ is asymptotically either almost surely true, or almost surely false, can thus be seen as a special case of Corollary 7.2.10 for the Boolean semiring (where $\varepsilon_K = 1$). In particular, the almost sure valuations ε_K and 1 in any finite lattice semiring K occur precisely for the formulae which are almost surely true in the Boolean case.

Example 7.2.12 (Secret facts)**.** Semiring semantics in finite min-max semirings can be used to model access restrictions to atomic facts. For example, we can use the min-max semiring with elements $0 < \mathsf{S} < \mathsf{C} < \mathsf{P}$, which stand for "inaccessible" (or false), "secret", "confidential", and "public". An interpretation π into this semiring labels atomic facts by access restrictions and the associated valuation $\pi[\![\psi]\!]$ of a sentence ψ then describes the clearance level that is

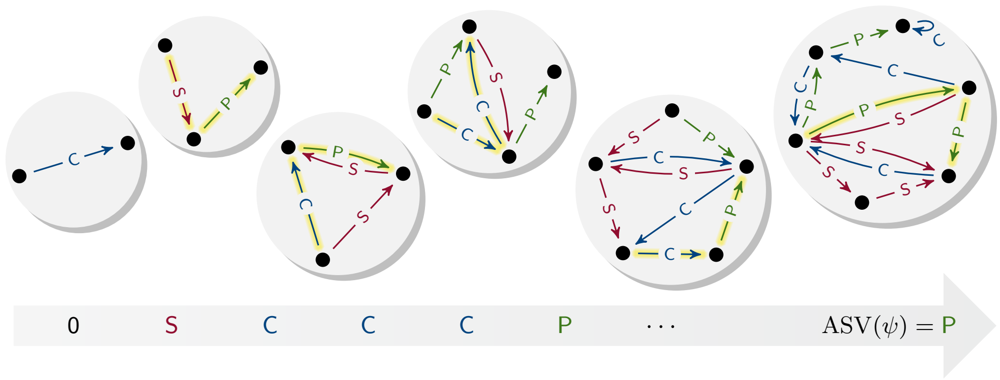

Figure 7.4: Illustration of the intuition for quantifier elimination in the access control semiring (cf. Example 7.2.12). The values of $\psi = \exists x\, \exists y\, \exists z\, (Exy \wedge Eyz)$ in the (randomly chosen) interpretations are listed at the bottom, approaching the largest value P.

necessary to verify the truth of ψ under these restrictions (see Fig. 7.4 for an example with annotated graphs). The 0-1 law in Corollary 7.2.10 implies that under a random assignment of access restrictions (assuming positive probabilities of all security levels), any sentence can almost surely either be checked with publicly available information, cannot be checked at all, or requires clearance for secret information (but never for confidential information).

Figure 7.4 illustrates the intuition for this behaviour: for increasingly large interpretations, it becomes more and more likely (or, if the appropriate extension property holds, certain) that whenever an existential quantification holds in the underlying graph, then there is also a witness which yields the largest possible truth value. Similarly, universal quantifiers are interpreted as minimum, and with increasing size it is almost surely the case that the smallest positive value (or 0) is assumed.

7.2.4 Complexity

We now study the complexity of computing the almost sure valuation of a given first-order sentence ψ in finite lattice semirings (we do not consider infinitary sentences here, due to their infinite representations). As shown above, this amounts to the computation of the associated polynomial f_ψ. While f_ψ is either 0, 1, or e for a sentence ψ, the polynomials $f_\varphi(\boldsymbol{X}^{(k)})$ associated with formulae $\varphi(x_1, \dots, x_k)$ are much more complicated and can have exponential length. Rather than computing these intermediate polynomials explicitly, we present a recursive procedure for computing the values $f_\varphi[\rho]$ for any formula $\varphi(x_1, \dots, x_k) \in \mathrm{FO}(\tau)$ and any atomic k-type $\rho\colon \mathrm{Lit}_k(\tau) \to K$ with values in a finite min-max semiring.

We remark that the polynomial f_ψ is the same for all finite lattice semirings. For determining the almost sure valuations of first-order sentences it would therefore suffice to define the procedure just for the three-element semiring E. However, we can solve, with moderate additional effort, the more general problem of computing valuations $\pi[\![\psi(\boldsymbol{a})]\!]$ of formulae with free variables not just for E, but for any finite min-max semiring K, and any finite K-interpretation $\pi\colon \mathrm{Lit}_A(\tau) \to K$ with the k-extension property. Indeed, by Theorem 7.2.9 we know that $\pi[\![\psi(\boldsymbol{a})]\!] = f_\psi[\rho^\pi_{\boldsymbol{a}}]$.

We first prove that this evaluation problem can be solved in PSPACE, for any fixed finite min-max semiring K. Using the well-known fact that PSPACE coincides with alternating polynomial time, we present the evaluation algorithm as an alternating procedure $\mathbf{Eval}(\psi, \rho, c)$ which, given $\psi(x_1, \ldots, x_k) \in \mathrm{FO}(\tau)$, an atomic k-type $\rho\colon \mathrm{Lit}_k(\tau) \to K$, and a value $c \in K$ determines whether $f_\psi[\rho] = c$ (avoiding an explicit construction of f_ψ). We assume that the reader is familiar with the notion of an alternating algorithm and its presentation as a game between an existential and a universal player (see, e.g., [BDG90]).

For a complexity analysis, it is appropriate to assume that formulae are written with the standard quantifiers $\exists$ and $\forall$, rather than $\exists^{\neq}$ and $\forall^{\neq}$, since the elimination of standard quantifiers by excluding ones can increase the length of formulae exponentially. As a consequence, when treating quantifiers, the evaluation procedure will have to deal with potential equalities between different variables. Accordingly, for a formula $\varphi = \exists x_{k+1} \vartheta(x_1, \ldots, x_k, x_{k+1})$ we have the polynomial $f_\varphi := \sum_{i=1}^{k} f_{\vartheta(x_1, \ldots x_k, x_i)} + f_{\exists^{\neq} x_{k+1} \vartheta}$ and analogously for universal quantifiers.

The idea of the evaluation procedure is that, at any step where it has to be verified whether $f_\varphi[\rho] = c$ for some triple (φ, ρ, c), the existential player guesses values c_i for the immediate subformulae φ_i of φ which, if correct, would imply that indeed $f_\varphi[\rho] = c$. The universal player then challenges one of these claims. For formulae of the form $\exists x_{i+1}\, \vartheta$ or $\forall x_{i+1}\, \vartheta$, this involves (existential and/or universal) choices of selector functions $s\colon \boldsymbol{Y}^{(i)} \to \{0, 1\}$ or $s\colon \boldsymbol{Y}^{(i)} \to \{0, e\}$ and the modification of $\rho\colon \mathrm{Lit}_i(\tau) \to K$ to the extended type $\rho s\colon \mathrm{Lit}_{i+1}(\tau) \to K$ defined by

$$(\rho s)(\alpha) = \begin{cases} \rho(\beta) & \text{if } \beta \in \mathrm{Lit}_i(\tau), \\ s(X_\beta) & \text{if } \beta \in \mathrm{Lit}_{i+1}(\tau) \setminus \mathrm{Lit}_i(\tau). \end{cases}$$

The procedure ends at triples (φ, ρ, c) where φ is atomic, at which point the algorithm just checks whether $\rho(\varphi) = c$. A detailed description of the algorithm for any relational vocabulary τ and any min-max semiring $(K, \max, \min, 0, 1)$ is given in Fig. 7.5.

It is obvious that the algorithm runs in alternating polynomial time, but it remains to prove that it is correct; we proceed by induction on ψ. Given a triple (ψ, ρ, c) such that, indeed, $f_\psi[\rho] = c$, the algorithm accepts by making the following existential choices. At a disjunction or conjunction, the existential player guesses the correct values of the immediate subformula. For a formula $\exists x_{k+1}\, \vartheta(x_1, \ldots, x_k, x_{k+1})$ the existential player guesses the values $c_i = f_{\vartheta(x_1, \ldots, x_k, x_i)}[\rho]$ and $c_{k+1} = f_{\exists^{\neq} x_{k+1} \vartheta}[\rho]$. If the universal player challenges the value for some $i \leq k$, the existential player wins the remaining game from the triple $(\varphi(x_1, \ldots, x_k, x_i), \rho, c_i)$ by induction hypothesis. If instead c_{k+1} is challenged, then the existential player guesses some selector function $s\colon \boldsymbol{Y}^{(k)} \to \{0, 1\}$ such that $c_{k+1} = f_{\exists^{\neq} x_{k+1} \vartheta}[\rho] = f_\vartheta[\rho s]$. The universal player challenges this by choosing also a function $s'\colon \boldsymbol{Y}^{(k)} \to \{0, 1\}$. If $s' = s$ this corresponds to the challenge to prove that, indeed, $f_\vartheta[\rho s] = c_{k+1}$; since this is the case, and by induction hypothesis, the existential player wins the remaining game. If $s' \neq s$ this corresponds to the challenge to prove that $f_\vartheta[\rho s'] \leq c_{k+1}$. The existential player answers this by guessing the correct value $c' := f_\vartheta[\rho s']$ and, again by the hypothesis, wins the remaining game. For formulae $\forall x_{k+1}\, \vartheta(x_1, \ldots, x_k, x_{k+1})$, the reasoning is analogous.

Consider now a triple (φ, ρ, c) such that $f_\psi[\rho] \neq c$. Then the existential player must make incorrect guesses, and the universal player can make sure that such incorrect triples are propagated through the play, and are then detected at the end, when an atomic formula is evaluated. Consider again the case of a formula $\varphi = \exists x_{k+1}\, \vartheta(x_1, \ldots, x_k, x_{k+1})$. From an

Eval(ψ, ρ, c)
input:
  a formula $\psi(x_1, \dots, x_k) \in \mathrm{FO}(\tau)$ in nnf
  an atomic type $\rho\colon \mathrm{Lit}_k(\tau) \to K$
  an element $c \in K$

if ψ is an atom or negated atom **then**
  accept if $\rho(\varphi) = c$, **else reject**

if $\psi = \varphi_1 \vee \varphi_2$ **then**
  guess $c_1, c_2 \in K$ with $\max(c_1, c_2) = c$
  universally choose $i \in \{1, 2\}$
  Eval(φ_i, ρ, c_i)

if $\psi = \varphi_1 \wedge \varphi_2$ **then**
  guess $c_1, c_2 \in K$ with $\min(c_1, c_2) = c$
  universally choose $i \in \{1, 2\}$
  Eval(φ_i, ρ, c_i)

if $\psi = \exists x_{k+1}\, \varphi$ **then**
  guess $c_1, \dots c_{k+1}$ s.t. $\max(c_1, \dots, c_{k+1}) = c$
  universally choose $i \in \{1, \dots, k+1\}$
  if $i \le k$ **then**
    set $\vartheta(x_1, \dots, x_k) \coloneqq \varphi(x_1, \dots, x_k, x_i)$
    Eval(ϑ, ρ, c_i)
  if $i = k + 1$ **then**
    guess $s\colon \boldsymbol{Y}^{(k)} \to \{0, 1\}$
    universally choose $s'\colon \boldsymbol{Y}^{(k)} \to \{0, 1\}$
    if $s' = s$ **then Eval**$(\varphi, \rho s, c_{k+1})$
    else guess $c' \le c_{k+1}$ and **Eval**$(\varphi, \rho s', c')$

if $\psi = \forall x\, \varphi$ **then**
  guess $c_1, \dots c_{k+1}$ s.t. $\min(c_1, \dots, c_{k+1}) = c$
  universally choose $i \in \{1, \dots, k+1\}$
  if $i \le k$ **then**
    set $\vartheta(x_1, \dots, x_k) \coloneqq \varphi(x_1, \dots, x_k, x_i)$
    Eval(ϑ, ρ, c_i)
  if $i = k + 1$ **then**
    guess $s\colon \boldsymbol{Y}^{(k)} \to \{0, \varepsilon_K\}$
    universally choose $s'\colon \boldsymbol{Y}^{(k)} \to \{0, \varepsilon_K\}$
    if $s' = s$ **then Eval**$(\varphi, \rho s, c_{k+1})$
    else guess $c' \ge c_{k+1}$ and **Eval**$(\varphi, \rho s', c')$

Figure 7.5: Alternating procedure **Eval**(ψ,ρ,c) to decide $f_\psi[\rho] = c$ in min-max semirings.

incorrect triple (φ, ρ, c), the existential player guesses $c_1, \dots, c_{k+1}$ with $\max(c_1, \dots, c_{k+1}) = c$. Hence either $(\vartheta(x_1, \dots, x_k, x_i), \rho, c_i)$ is incorrect for some $i \le k$, in which case the universal players chooses such an i and wins by induction hypothesis, or the triple $(\exists^{\neq} x_{k+1} \vartheta, \rho, c_{k+1})$ is incorrect. In that case, for any function $s\colon \boldsymbol{Y}^{(k)} \to \{0, 1\}$ that the existential player might guess, it is either the case that $f_\vartheta[\rho s] \ne c_{k+1}$, in which case the universal players wins by choosing $s' = s$, or that there exists another function $s'\colon \boldsymbol{Y}^{(k)} \to \{0, 1\}$ with the property that $f_\vartheta[\rho s'] > c_{k+1}$. Whatever element $c' \le c_{k+1}$ the existential player then guesses, the universal player will then win the remaining game from the incorrect triple $(\vartheta, \rho s', c')$. Again, the reasoning for universally quantified formulae is completely analogous.

We thus have established the following result, for any finite min-max semiring K and any relational vocabulary τ.

Theorem 7.2.13. *Let K be a fixed finite min-max semiring. Given $\psi(x_1, \dots, x_k) \in \mathrm{FO}(\tau)$ and an atomic k-type $\rho\colon \mathrm{Lit}_k(\tau) \to K$, the value $f_\psi[\rho]$ can be computed in* Pspace.

If we are only interested in the case where ψ is a sentence, we can work over the min-max semiring E and thus determine in Pspace whether f_ψ is 0,1, or e. On the other side, is has been proved by Grandjean [Gra83] that, in classical Boolean semantics, the problem whether a given first-order sentence is almost surely true or almost surely false is Pspace-complete.

Corollary 7.2.14. *For any finite lattice semiring K, verifying the almost sure valuation of first-order sentences in K is a* Pspace*-complete problem.*

Grandjean's result readily implies that, for any semiring K, deciding whether or not the almost sure K-valuation of a first-order sentence is 0, is Pspace-complete as well. It might still be the case that if it is known that ψ is almost surely true in the Boolean sense, then the problem whether its almost sure valuation in a finite lattice semiring is 1 or ε_K could be solved more efficiently. However, this is not the case.

Proposition 7.2.15. *The problem to decide whether a given almost surely true first order sentence evaluates in finite lattice semirings with at least three elements almost surely to 1, or to ε_K, is* PSPACE*-complete.*

Proof. It remains to show PSPACE-hardness. For any fixed finite structure $\mathfrak{A}$ with at least two elements, the problem of evaluating a given first-order sentence on $\mathfrak{A}$ is PSPACE-complete. In particular this holds if $\mathfrak{A}$ is just a two-element set without any relations, i.e. $\mathfrak{A} = \{0, 1\}$. Given a sentence $\psi \in \mathrm{FO}(\emptyset)$, we consider $\psi^* := \exists 0 \exists 1 (0 \neq 1 \wedge \psi')$ where ψ' is obtained by relativising all quantifiers to $\{0, 1\}$, i.e. by replacing subformulae $\exists x \varphi$ by $\exists x((x = 0 \vee x = 1) \wedge \varphi)$ and $\forall x \varphi$ by $\forall x((x = 0 \vee x = 1) \rightarrow \varphi)$. Clearly if $\{0, 1\} \models \psi$ then ψ^* almost surely evaluates to 1 (on any semiring) and if $\{0, 1\} \not\models \psi$ then ψ^* almost surely evaluates to 0.

Let now P be a unary relation symbol and consider the reduction that maps any sentence $\psi \in \mathrm{FO}(\emptyset)$ to $\psi^* \vee \forall x(Px \vee \neg Px) \in \mathrm{FO}(\{P\})$. Notice that such a sentence is almost surely true in the Boolean sense, and that the almost sure valuation of $\forall x(Px \vee \neg Px)$ is ε in any finite lattice semiring. Hence the almost sure valuation of $\psi^* \vee \forall x(Px \vee \neg Px)$ is 1 if $\{0, 1\} \models \psi$, and ε, otherwise. This proves that deciding whether an almost surely true sentence evaluates to 1 or to ε in a lattice semiring with at least three elements is PSPACE-hard. □

7.3 Zero-One Law for Infinite Lattice Semirings

We now move to infinite lattice semirings $(K, \sqcup, \sqcap, 0, 1)$, in particular to semirings defined over the real numbers. Notice that the semantics of $L^{\omega}_{\infty\omega}$ on finite interpretations is well-defined in such semirings, even if they are not complete lattices (see Remark 5.3.3). However, the probability measures have to be adapted to work with infinite semirings.

7.3.1 Random Interpretations Revisited

In the case that K is countable, we can define probability measures on K-interpretations as in Section 7.2.1. In the general case, we assume that we have a probability space $(K^+, \mathcal{F}, p)$ whose underlying σ-algebra $\mathcal{F}$ contains all intervals $[a, b] = \{x \in K \mid a \leq x \leq b\}$ for $a, b \in K$ (notice that $[a, b]$ is a sublattice). We thus get probabilities $p[x \in J]$ for all closed, open, and half-open intervals $J \subseteq K^+$. We further assume that $p[x = 1] > 0$, i.e. we have a positive probability that a randomly chosen value coincides with the maximal semiring value.[1]

The measures $\mu_{n,p}$ for random K-interpretations with universe $[n] = \{0, \ldots, n-1\}$ are induced by p as in the finite case: Again, we consider the probabilistic process which, for each instantiated atom $R\boldsymbol{a}$ over $[n]$ first makes a random choice whether $R\boldsymbol{a}$ or $\neg R\boldsymbol{a}$ is true, each with probability $\frac{1}{2}$ (this is an arbitrary choice, any fixed probability would work), and

[1]This is a natural assumption in our context of random semiring interpretations, but it is not really essential; large values can instead be treated in an analogous way as we do for small positive ones.

then assigns to the true literal a positive semiring value according to p, so that we have a probability that $\pi(R\boldsymbol{a}) \in J$ for every interval[2] $J \subseteq K$. We consider two cases concerning the probabilities of small positive semiring values.

Definition 7.3.1. We say that the probability measure p is *ε-bounded* on small semiring values, for $\varepsilon \in K$, if one of the following cases applies.

(1) *p is weakly ε-bounded* if $p[x = \varepsilon] > 0$ and $p[0 < x \leq \delta] = 0$ for all $\delta \in K^+$ with $\varepsilon \not\leq \delta$.

In particular, the smallest possible positive value of a literal is ε.

(2) *p is strictly ε-bounded* if p is not weakly ε'-bounded (for any ε') and further $p[0 < x \leq \varepsilon] = 0$ and $p[0 < x \leq \delta] > 0$ for all $\delta > \varepsilon$.

That is, p only admits positive values greater than ε. We include the case $\varepsilon = 0$.

To avoid going through case distinctions in the proofs to follow, we say that a semiring value $\delta \in K^+$ is *p-relevant*, if either $\delta > \varepsilon$, or if $p[x = \varepsilon] > 0$ and $\delta = \varepsilon$. Moreover, we write $\varepsilon \ll \delta$ if there is a $\gamma \in K$ with $\varepsilon < \gamma < \delta$.

In the remainder of this section we consider infinite lattice semirings $(K, \sqcup, \sqcap, 0, 1)$ together with a probability measure p on K^+ assigning probabilities to all intervals, such that p is ε-bounded with $\varepsilon \ll 1$. We remark that we make the assumption $\varepsilon \ll 1$ only to simplify the presentation, but this is not an actual restriction (one can easily verify that Corollaries 7.3.10 and 7.3.11 also holds in the few special cases with $\varepsilon \not\ll 1$).

Definition 7.3.2. Let $\Phi \subseteq L^{\omega}_{\infty\omega}(\tau)$. We say that a *0-1 law* holds for Φ, K, and p if for each sentence $\psi \in \Phi$ and each interval $J \subseteq K$ the sequence $(\mu_{n,p}[\pi[\![\psi]\!] \in J])_{n<\omega}$ converges to either 0 or 1, as n goes to infinity.

We further say that j is the *almost sure valuation* of ψ (for K and p), denoted $\mathrm{ASV}_{K,p}(\psi) = j$, if there is a decreasing sequence $(J_i)_{i<\omega}$ of intervals $J_i \subseteq K$ with $\bigcap_{i<\omega} J_i = \{j\}$ such that $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] \in J_i] = 1$ for all $i < \omega$.

We also have to define the extension properties a bit differently, as we cannot realise all possible extensions over an infinite semiring in a finite structure.

Definition 7.3.3. Given an atomic m-type ρ, we say that $\rho^+ \in \mathrm{ext}(\rho)$ is a *maximal extension* of ρ if $\rho^+(\beta) \in \{0, 1\}$, for every literal $\beta \in \mathrm{Lit}_{m+1}(\tau) \setminus \mathrm{Lit}_m(\tau)$. Further we say that $\rho^- \in \mathrm{ext}(\rho)$ is a *δ-small extension* of ρ, if $\rho^-(\beta) \leq \delta$ for every $\beta \in \mathrm{Lit}_{m+1}(\tau) \setminus \mathrm{Lit}_m(\tau)$.

We remark that, by definition of atomic types, a δ-small extension ρ^- maps out of each pair $\alpha, \neg\alpha$ of complementary literals that contain the variable x_{m+1} precisely one to 0 and the other one into the interval $(0, \delta]$.

[2] If $0 \in J$, then $\mu_{n,p}[\pi(R\boldsymbol{a}) \in J] = \frac{1}{2} + \frac{1}{2}p[x \in J \setminus \{0\}]$, otherwise $\mu_{n,p}[\pi(R\boldsymbol{a}) \in J] = \frac{1}{2}p[x \in J \setminus \{0\}]$.

Definition 7.3.4. A semiring interpretation $\pi\colon \mathrm{Lit}_A(\tau) \to K$ has the (k,δ)-*extension property*, where $\delta \in K^+$, if for every $m < k$, every tuple $\boldsymbol{a} \in A^m$, and every maximal extension $\rho^+ \in \mathrm{ext}(\rho^\pi_{\boldsymbol{a}})$, there exists

(1) an element $b \in A \setminus \boldsymbol{a}$ such that $\rho^\pi_{\boldsymbol{a},b} = \rho^+$, and

(2) an element $c \in A \setminus \boldsymbol{a}$ such that $\rho^\pi_{\boldsymbol{a},c} \leq \rho^+$ and $\rho^\pi_{\boldsymbol{a},c}$ is a δ-small extension of $\rho^\pi_{\boldsymbol{a}}$.

That is, if π has the (k,δ)-extension property, then every realisation of an atomic m-type in π can be extended to realisations of all its maximal extensions, but also to realisations of δ-small extensions (with the same underlying Boolean types as the maximal extensions).

Lemma 7.3.5. *Let K be an infinite lattice semiring with an ε-bounded probability measure p. For every fixed k, every finite relational vocabulary τ, and every p-relevant δ,*

$$\lim_{n\to\infty} \mu_{n,p}[\pi \text{ has the } (k,\delta)\text{-extension property }] = 1,$$

and the convergence to this limit is exponentially fast.

Proof. For a given probability measure $\mu_{n,p}$ we first calculate a bound for the probability that a given realisation $\boldsymbol{a}$ of an atomic m-type ρ (with $m < k$) cannot be extended to a realisation $\boldsymbol{a}, b$ of a given *maximal extension* ρ^+ of ρ. This is analogous to the argument in Lemma 7.2.4. For any pair $\alpha, \neg\alpha$ of complementary literals in $\mathrm{Lit}_{m+1}(\tau) \setminus \mathrm{Lit}_m(\tau)$, the probability that randomly chosen values according to p for α and $\neg\alpha$ are 1 and 0, as prescribed by ρ^+, is $p[x = 1]/2$. There is a fixed number q of pairs of such literals, so the probability that all chosen values coincide with those required by ρ^+ is a fixed number $\gamma \coloneqq (p([x = 1]/2)^q$. It follows that

$$\mu_{n,p}[\text{ some realisation of } \rho \text{ does not extend to a realisation of } \rho^+] \leq n^m(1-\gamma)^{n-m}$$

which for growing n converges to 0 exponentially fast.

Let us now consider extensions with small truth values. Fix ρ and some maximal extension $\rho^+ \in \mathrm{ext}(\rho)$. For each p-relevant δ there exists a number $g(\delta) > 0$ such that $p[0 < x \leq \delta] = g(\delta)$. Hence the probability that values for complementary literals α and $\neg\alpha$ with the variable x_{m+1}, chosen according to p, define a δ-small extension $\rho^- \leq \rho^+$ is $\gamma \coloneqq (g(\delta)/2)^q$, and with precisely the same calculation as above, we conclude that

$$\mu_{n,p}[\text{ some realisation of } \rho \text{ does not extend to a realisation of some } \delta\text{-small } \rho^- \leq \rho^+]$$

converges to 0 exponentially fast. □

7.3.2 Proof for Infinite Lattice Semirings

We again use the polynomials f_ψ of Definition 7.2.5 to represent formulae $\psi(\boldsymbol{x}) \in L^k_{\infty\omega}$. However, the evaluation of these polynomials must be more flexible, taking into account different parameters for small positive values. Specifically, given $\delta > 0$ and an atomic k-type ρ, we evaluate a polynomial $f \in \mathrm{AbsPoly}(E, \mathbb{B}, \boldsymbol{X}^{(k)})$ to a semiring value $f^\delta[\rho] \in K$, via the homomorphism $h^\delta_\rho\colon \mathrm{AbsPoly}(E, \mathbb{B}, \boldsymbol{X}^{(k)}) \to K$ induced by $h^\delta_\rho(e) \coloneqq \delta$ and $h^\delta_\rho(X_\beta) \coloneqq \rho(\beta)$, for literals $\beta \in \mathrm{Lit}_k(\tau)$. We can now formulate an analogue of Theorem 7.2.9, requiring only a mild assumption on the lattice structure:

Definition 7.3.6. A lattice semiring $(K, \sqcup, \sqcap, 0, 1)$ has *no divisors of* 0 if $a \sqcap b = 0$ implies $a = 0$ or $b = 0$. It is *0-1-irreducible* if it has no divisors of 0 and additionally, $a \sqcup b = 1$ implies $a = 1$ or $b = 1$.

Both properties are always satisfied in min-max semirings (as the natural order is total).

Theorem 7.3.7. *Let $(K, \sqcup, \sqcap, 0, 1)$ be a (possibly infinite) lattice semiring without divisors of 0. Let $\delta > 0$ and let $\pi\colon \mathrm{Lit}_A(\tau) \to K$ be a finite K-interpretation with the (k, δ)-extension property. Then, for every formula $\psi(x_1, \dots, x_i) \in L^k_{\infty\omega}(\tau)$ and every tuple $\boldsymbol{a} \in A^i$, either*

- $f^\delta_\psi[\rho^\pi_{\boldsymbol{a}}] = \pi[\![\psi(\boldsymbol{a})]\!] = 0$, *or* **(0)**
- $f^\delta_\psi[\rho^\pi_{\boldsymbol{a}}], \pi[\![\psi(\boldsymbol{a})]\!] \neq 0$ and $f^\delta_\psi[\rho^\pi_{\boldsymbol{a}}] \leq \pi[\![\psi(\boldsymbol{a})]\!] \sqcup \delta \leq f^\delta_\psi[\rho^\pi_{\boldsymbol{a}}] \sqcup \delta$. **(δ)**

Proof. The proof is by induction over ψ along the lines of the proof of Theorem 7.2.9. Since δ is fixed, we drop the annotation δ in f^δ_ψ to ease notation. For literals, we always have $f_\psi[\rho^\pi_{\boldsymbol{a}}] = \pi[\![\psi(\boldsymbol{a})]\!]$ and one of the two cases **(0)** or **(δ)** holds.

Let $\psi = \bigvee \Phi$. If **(0)** applies to all $\varphi \in \Phi$, then it also applies to ψ. Otherwise, **(δ)** applies to some $\varphi \in \Phi$ and hence $f_\psi[\rho^\pi_{\boldsymbol{a}}] \geq f_\varphi[\rho^\pi_{\boldsymbol{a}}] \neq 0$ and $\pi[\![\psi(\boldsymbol{a})]\!] \geq \pi[\![\varphi(\boldsymbol{a})]\!] \neq 0$. We show that **(δ)** applies to ψ:

$$f_\psi[\rho^\pi_{\boldsymbol{a}}] = \bigsqcup_{\varphi \in \Phi} f_\varphi[\rho^\pi_{\boldsymbol{a}}] \leq \bigsqcup_{\varphi \in \Phi} \pi[\![\varphi(\boldsymbol{a})]\!] \sqcup \delta = \pi[\![\psi(\boldsymbol{a})]\!] \sqcup \delta,$$

and similarly for the second inequality.

For $\psi = \bigwedge \Phi$, case **(0)** applies whenever **(0)** holds for any $\varphi \in \Phi$. So assume that **(δ)** applies to all $\varphi \in \Phi$. Recall that the infimum ranges over finitely many values (since $\mathrm{AbsPoly}(E, \mathbb{B}, \boldsymbol{X}^{(k)})$ is finite and by Remark 5.3.3). Hence $f_\psi[\rho^\pi_{\boldsymbol{a}}] \neq 0$ and $\pi[\![\psi(\boldsymbol{a})]\!] \neq 0$, as K has no divisors of 0. The two inequalities for **(δ)** follow by induction and (the dual of) distributivity:

$$f_\psi[\rho^\pi_{\boldsymbol{a}}] = \bigsqcap_{\varphi \in \Phi} f_\varphi[\rho^\pi_{\boldsymbol{a}}] \leq \bigsqcap_{\varphi \in \Phi} (\pi[\![\varphi(\boldsymbol{a})]\!] \sqcup \delta) = \Big(\bigsqcap_{\varphi \in \Phi} \pi[\![\varphi(\boldsymbol{a})]\!]\Big) \sqcup \delta = \pi[\![\psi(\boldsymbol{a})]\!] \sqcup \delta,$$

and

$$f_\psi[\rho^\pi_{\boldsymbol{a}}] \sqcup \delta = \Big(\bigsqcap_{\varphi \in \Phi} f_\varphi[\rho^\pi_{\boldsymbol{a}}]\Big) \sqcup \delta = \bigsqcap_{\varphi \in \Phi} (f_\varphi[\rho^\pi_{\boldsymbol{a}}] \sqcup \delta) \geq \bigsqcap_{\varphi \in \Phi} (\pi[\![\varphi(\boldsymbol{a})]\!] \sqcup \delta) = \pi[\![\psi(\boldsymbol{a})]\!] \sqcup \delta.$$

Let now $\psi(\boldsymbol{x}) = \exists^{\neq} y\, \varphi(\boldsymbol{x}, y)$. Recall that

$$\pi[\![\psi(\boldsymbol{a})]\!] = \bigsqcup_{b \in A \setminus \boldsymbol{a}} \pi[\![\vartheta(\boldsymbol{a})]\!] \quad \text{and} \quad f_\psi(\boldsymbol{X}^{(i)}) = \bigsqcup_{s \in S} f_\varphi(\boldsymbol{X}^{(i)}, s(\boldsymbol{Y}^{(i)})),$$

where $\boldsymbol{Y}^{(i)} = \boldsymbol{X}^{(i+1)} \setminus \boldsymbol{X}^{(i)}$ and S is the set of all consistent selector functions $s\colon \boldsymbol{Y}^{(i)} \to \{0, 1\}$.

We first prove that $f_\psi[\rho^\pi_{\boldsymbol{a}}] = \bigsqcup_{b \in A \setminus \boldsymbol{a}} f_\varphi[\rho^\pi_{\boldsymbol{a},b}]$. Recall that each selector function s induces the maximal extension ρ_s with $\rho_s(\beta) = s(\beta)$ for the new literals β. By the (k, δ)-extension property, there is an element b with $\rho_{\boldsymbol{a},b} = \rho_s$ and we then have $f_\varphi[\rho^\pi_{\boldsymbol{a}}, s(\boldsymbol{Y}^{(i)})] = f_\varphi[\rho^\pi_{\boldsymbol{a},b}]$. Hence $f_\psi[\rho^\pi_{\boldsymbol{a}}] \leq \bigsqcup_{b\in A\setminus\boldsymbol{a}} f_\varphi[\rho^\pi_{\boldsymbol{a},b}]$. Conversely, let $b \in A \setminus \boldsymbol{a}$ and consider the type $\rho^\pi_{\boldsymbol{a},b}$. Let ρ^+ be the maximal extension induced by $\rho^\pi_{\boldsymbol{a},b}$ (i.e., with the same underlying Boolean type). By the (k,δ)-extension property, there is an element b^+ with $\rho^+ = \rho^\pi_{\boldsymbol{a},b^+}$. Then $\rho^\pi_{\boldsymbol{a},b} \leq \rho^\pi_{\boldsymbol{a},b^+}$ and hence $f_\varphi[\rho^\pi_{\boldsymbol{a},b}] \leq f_\varphi[\rho^\pi_{\boldsymbol{a},b^+}]$ by monotonicity. Setting $s(\beta) = \rho^+_b(\beta)$, we have $f_\varphi[\rho^\pi_{\boldsymbol{a},b^+}] = f_\varphi[\rho^\pi_{\boldsymbol{a}}, s(\boldsymbol{Y}^{(i)})]$ and hence $\bigsqcup_{b\in A\setminus\boldsymbol{a}} f_\varphi[\rho^\pi_{\boldsymbol{a},b}] \leq f_\psi[\rho^\pi_{\boldsymbol{a}}]$.

To prove that either **(0)** or **($\boldsymbol{\delta}$)** holds for ψ, we proceed by case distinction for each b. If **(0)** holds for all $\varphi(\boldsymbol{a}, b)$, then $\pi[\![\psi(\boldsymbol{a})]\!] = 0$ and also $\bigsqcup_{b\in A\setminus\boldsymbol{a}} f_\varphi[\rho^\pi_{\boldsymbol{a},b}] = 0$, so **(0)** holds for ψ. Otherwise, there is at least one b with $f_\varphi[\rho^\pi_{\boldsymbol{a},b}], \pi[\![\varphi(\boldsymbol{a}, b)]\!] \neq 0$, hence $f_\psi[\rho^\pi_{\boldsymbol{a}}], \pi[\![\psi(\boldsymbol{a})]\!] \neq 0$ as well. We ignore all b for which **(0)** holds, as they do not affect the supremum. Then **($\boldsymbol{\delta}$)** holds for ψ:

$$\begin{aligned} f_\psi[\rho^\pi_{\boldsymbol{a}}] = \bigsqcup_{b\in A\setminus\boldsymbol{a}} f_\varphi[\rho^\pi_{\boldsymbol{a},b}] &\leq \bigsqcup_{b\in A\setminus\boldsymbol{a}} (\pi[\![\varphi(\boldsymbol{a}, b)]\!] \sqcup \delta) = \pi[\![\psi(\boldsymbol{a})]\!] \sqcup \delta \\ &\leq \bigsqcup_{b\in A\setminus\boldsymbol{a}} (f_\varphi[\rho^\pi_{\boldsymbol{a},b}] \sqcup \delta) = f_\psi[\rho^\pi_{\boldsymbol{a}}] \sqcup \delta. \end{aligned}$$

Finally, let $\psi(\boldsymbol{x}) = \forall^{\neq} y\, \varphi(\boldsymbol{x}, y)$. Recall that

$$\pi[\![\psi(\boldsymbol{a})]\!] = \bigsqcap_{b\in A\setminus\boldsymbol{a}} \pi[\![\vartheta(\boldsymbol{a})]\!] \quad \text{and} \quad f_\psi(\boldsymbol{X}^{(i)}) = \bigsqcap_{s\in S} f_\varphi(\boldsymbol{X}^{(i)}, s(\boldsymbol{Y}^{(i)})),$$

where now we consider selector functions $s\colon \boldsymbol{Y}^{(i)} \to \{0, \delta\}$.

As for existential quantification, we first relate the selector functions s to the elements $b \in A \setminus \boldsymbol{a}$. Since π only guarantees δ-small extensions, we relax the equality by δ:

$$\Big(\bigsqcap_{b\in A\setminus\boldsymbol{a}} f_\varphi[\rho^\pi_{\boldsymbol{a},b}]\Big) \sqcup \delta \;\geq\; f_\psi[\rho^\pi_{\boldsymbol{a}}] \;\geq\; \bigsqcap_{b\in A\setminus\boldsymbol{a}} f_\varphi[\rho^\pi_{\boldsymbol{a},b}].$$

The second inequality is easy: For each selector function s, consider the maximal extension ρ^+_s of $\rho^\pi_{\boldsymbol{a}}$ induced by s (i.e., with the same underlying Boolean type). By the (k,δ)-extension property, there is an element c such that $\rho^\pi_{\boldsymbol{a},c}$ is a δ-small extension with $\rho^\pi_{\boldsymbol{a},c} \leq \rho^+_s$. Then also $\rho^\pi_{\boldsymbol{a},c}(\beta) \leq s(X_\beta)$ for all new literals β by definition of s and δ-small, hence $f_\varphi[\rho^\pi_{\boldsymbol{a}}, s(\boldsymbol{Y}^{(i)})] \geq f_\varphi[\rho^\pi_{\boldsymbol{a},c}]$ by monotonicity and the inequality follows.

For the first inequality, we consider the monomials of $f_\varphi(\boldsymbol{X}^{(i)}, \boldsymbol{Y}^{(i)})$ and split the polynomial into $f_\varphi(\boldsymbol{X}^{(i)}, \boldsymbol{Y}^{(i)}) = g(\boldsymbol{X}^{(i)}) + h(\boldsymbol{X}^{(i)}, \boldsymbol{Y}^{(i)})$, where g contains precisely those monomials that contain no indeterminates in $\boldsymbol{Y}^{(i)}$. We clearly have $f_\varphi[\rho^\pi_{\boldsymbol{a},b}] \geq g[\boldsymbol{a}]$. Recall that the universe A is finite, so we can apply distributivity and obtain:

$$\Big(\bigsqcap_{b\in A\setminus\boldsymbol{a}} f_\varphi[\rho^\pi_{\boldsymbol{a},b}]\Big) \sqcup \delta = \bigsqcap_{b\in A\setminus\boldsymbol{a}} (f_\varphi[\rho^\pi_{\boldsymbol{a},b}] \sqcup \delta) \geq g[\boldsymbol{a}] \sqcup \delta \geq \bigsqcap_{s\in S} (g[\boldsymbol{a}] \sqcup h(\boldsymbol{a}, s(\boldsymbol{Y}^{(i)}))) = f_\psi[\rho^\pi_{\boldsymbol{a}}].$$

For the last inequality, we use the fact that $s(X_\beta) \in \{0, \delta\}$ for all $X_\beta \in \boldsymbol{Y}^{(i)}$ and hence $m[\boldsymbol{a}, s(\boldsymbol{Y}^{(i)})] \leq \delta$ for all monomials m of h by construction.

To prove that **(0)** or **($\boldsymbol{\delta}$)** holds for ψ, we again proceed by case distinction for each b. First assume that **(0)** holds for some $\varphi(\boldsymbol{a}, b)$, so $\pi[\![\varphi(\boldsymbol{a}, b)]\!] = f_\varphi[\rho^\pi_{\boldsymbol{a},b}] = 0$. Then also $\pi[\![\psi(\boldsymbol{a})]\!] = 0$ and it remains to prove $f_\psi[\rho^\pi_{\boldsymbol{a}}] = 0$. We again consider the monomials of $f_\varphi(\boldsymbol{X}^{(i)}, \boldsymbol{Y}^{(i)})$ and split the polynomial into $f_\varphi(\boldsymbol{X}^{(i)}, \boldsymbol{Y}^{(i)}) = g(\boldsymbol{X}^{(i)}) + h(\boldsymbol{X}^{(i)}, \boldsymbol{Y}^{(i)})$ as above. By $f_\varphi[\rho^\pi_{\boldsymbol{a},b}] = 0$, we must have $g[\boldsymbol{a}] = 0$ and $h[\boldsymbol{a}, b] = 0$. As there are no divisors of 0, this means that every monomial in $h[\boldsymbol{a}, b]$ must contain a literal $X_\beta \in \boldsymbol{X}^{(i)} \cup \boldsymbol{Y}^{(i)}$ such that $\rho^\pi_{\boldsymbol{a},b}(\beta) = 0$. Let s be any selector function such that whenever $\rho^\pi_{\boldsymbol{a},b}(X_\beta) = 0$ for $X_\beta \in \boldsymbol{Y}^{(i)}$, also $s(X_\beta) = 0$ (the other values can be chosen arbitrarily). Observe that such a selector function exists in S since $\rho^\pi_{\boldsymbol{a},b}$ is a type (i.e., consistent on opposing literals). Then $h[\boldsymbol{a}, b] = h[\boldsymbol{a}, s(\boldsymbol{Y}^{(i)})] = 0$ by construction of s and it follows that $f_\psi[\rho^\pi_{\boldsymbol{a}}] = 0$.

Lastly, assume that **($\boldsymbol{\delta}$)** holds for all b. Then $f_\varphi[\rho^\pi_{\boldsymbol{a},b}], \pi[\![\varphi(\boldsymbol{a}, b)]\!] \neq 0$ for all b and thus $f_\psi[\rho^\pi_{\boldsymbol{a}}], \pi[\![\psi(\boldsymbol{a})]\!] \neq 0$, since there are no divisors of 0 (recall that the infimum is over a finite universe or set S). Using the relaxed equality and distributivity, we obtain:

$$\begin{aligned} f_\psi[\rho^\pi_{\boldsymbol{a}}] \leq \bigsqcap_{b \in A \setminus \boldsymbol{a}} (f_\varphi[\rho^\pi_{\boldsymbol{a},b}] \sqcup \delta) &\leq \bigsqcap_{b \in A \setminus \boldsymbol{a}} ((\pi[\![\varphi(\boldsymbol{a}, b)]\!] \sqcup \delta) \sqcup \delta) = \pi[\![\psi(\boldsymbol{a})]\!] \sqcup \delta \\ &\leq \bigsqcap_{b \in A \setminus \boldsymbol{a}} ((f_\varphi[\rho^\pi_{\boldsymbol{a},b}] \sqcup \delta) \sqcup \delta) = f_\psi[\rho^\pi_{\boldsymbol{a},b}] \sqcup \delta. \end{aligned}$$ □

The above theorem essentially establishes a relaxed version of the equality $f^\delta_\psi[\rho^\pi_{\boldsymbol{a}}] = \pi[\![\psi(\boldsymbol{a})]\!]$ that holds in finite lattice semirings. The reason is the (k, δ)-extension property, which does not guarantee that the value δ is assumed by extensions, but only makes the weaker guarantee that some values in $(0, \delta]$ are assumed. For sentences, this relaxed equality reduces to the following three cases.

Corollary 7.3.8. *Let $(K, \sqcup, \sqcap, 0, 1)$ be a (possibly infinite) 0-1-irreducible lattice semiring. Let $0 < \delta < 1$ and let $\pi\colon \mathrm{Lit}_A(\tau) \to K$ be a K-interpretation with the (k, δ)-extension property. Then, for every sentence $\psi \in L^\omega_{\infty\omega}(\tau)$,*

- *if $f_\psi = 0$, then also $\pi[\![\psi]\!] = 0$;*
- *if $f_\psi = 1$, then also $\pi[\![\psi]\!] = 1$;*
- *if $f_\psi = e$, then $0 < \pi[\![\psi]\!] \leq \delta$.*

Proof. The first statement is immediate by Theorem 7.3.7. For the second statement, Theorem 7.3.7 implies $\pi[\![\psi]\!] \sqcup \delta = 1$. By assumption on K and $\delta < 1$, this implies $\pi[\![\psi]\!] = 1$. For the last statement, recall that we have $f^\delta_\psi[\emptyset] = \delta$ for $f_\psi = e$. Theorem 7.3.7 states $\pi[\![\psi]\!] \neq 0$ as well as $\pi[\![\psi]\!] \sqcup \delta = \delta$ which implies $\pi[\![\psi]\!] \leq \delta$. □

To determine which intervals occur almost surely in the case $f_\psi = e$, we need the following simple observation.

Lemma 7.3.9. *Let $J \subseteq K$ be a directed interval, i.e., $x, y \in J$ implies $x \sqcap y \in J$ and $x \sqcup y \in J$. If $\pi\colon \mathrm{Lit}_A(\tau) \to J \cup \{0, 1\}$ is a K-interpretation that maps all literals into J (or to 0 or 1), then also $\pi[\![\psi]\!] \in J \cup \{0, 1\}$, for every sentence $\psi \in L^\omega_{\infty\omega}(\tau)$.*

Proof. Straightforward induction on ψ. Recall that we assume the universe A to be finite, so all logical operators are evaluated as finite $\sqcap$ or $\sqcup$ (including infinitary conjunctions and disjunctions, cf. Remark 5.3.3) and the value thus remains in $J \cup \{0, 1\}$. □

Corollary 7.3.10 (0-1 law on infinite lattice semirings)**.** *Let $(K, \sqcup, \sqcap, 0, 1)$ be a 0-1-irreducible lattice semiring with ε-bounded probability measure, where $\varepsilon \ll 1$. Then, for every sentence $\psi \in L^{\omega}_{\infty\omega}(\tau) \cup \mathrm{LFP}(\tau)$ over relational vocabulary τ and every interval $J \subseteq K$, the sequence $(\mu_{n,p}[\pi[\![\psi]\!] \in J])_{n<\omega}$ converges exponentially fast to either 0 or 1, as n goes to infinity.*

Further, the only intervals J for which $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] \in J] = 1$ is possible are those where either $0 \in J$, $1 \in J$, $\varepsilon \in J$, or $(\varepsilon, \delta) \subseteq J$ for some $\delta > \varepsilon$.

Proof. First consider $\psi \in \mathrm{LFP}$. By similar reasoning as in Remark 5.3.3, the evaluation $\pi[\![\psi]\!]$ for an interpretation π with finite universe takes places in a finite sublattice K_0 of K and is thus well-defined (regardless of the continuity properties of K). By Theorem 5.3.5, we obtain a sentence $\psi' \in L^{\omega}_{\infty\omega}$ that is equivalent to ψ in all finite lattice semirings K_0, and hence also $\psi \equiv_K \psi'$ (recall that $\equiv_K$ is equivalence on finite universes), so it suffices to consider sentences in $L^{\omega}_{\infty\omega}$.

Now let $\psi \in L^{k}_{\infty\omega}(\tau)$ for some k, and consider the associated polynomial f_ψ. Since ψ is a sentence, we have $f_\psi \in \{0, 1, e\}$. For every p-relevant δ, the sequence

$$\mu_{n,p}[\pi \text{ has the } (k,\delta)\text{-extension property }]$$

converges to 1 exponentially fast. We first consider the case that $f_\psi = 0$ or $f_\psi = 1$. Since $\varepsilon \ll 1$, there is a p-relevant $\delta < 1$ and Corollary 7.3.8 thus implies that $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] \in J] = 1$ if $f_\psi \in J$ and $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] \in J] = 0$ otherwise.

Now consider the case that $f_\psi = e$. For every p-relevant δ and any K-interpretation π with the (k, δ)-extension property, Corollary 7.3.8 implies $0 < \pi[\![\psi]\!] \le \delta$. Since the (k, δ)-extension property is asymptotically almost surely satisfied, it follows that

$$\lim_{n\to\infty} \mu_{n,p}[0 < \pi[\![\psi]\!] \le \delta] = 1 \tag{$*$}$$

with exponential convergence, for every p-relevant δ. We get back to the two cases concerning the parameter ε of p:

(1) p is weakly ε-bounded: $p[x = \varepsilon] > 0$ and $p[0 < x \le \delta] = 0$ for all $\delta \in K^+$ with $\varepsilon \not\le \delta$.

Since p only admits values in the closed (and hence directed) interval $J_\varepsilon = [\varepsilon, 1]$, Lemma 7.3.9 implies $\mu_{n,p}[\pi[\![\psi]\!] \in J_\varepsilon] = 1$ for all n. Conversely, ε is p-relevant, so together with $(*)$,

$$\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] \in J] = \begin{cases} 1 & \text{if } \varepsilon \in J, \\ 0 & \text{otherwise.} \end{cases}$$

(2) p is strictly ε-bounded: $p[0 < x \le \varepsilon] = 0$ and $p[0 < x \le \delta] > 0$ for all $\delta > \varepsilon$.

First assume that there are $\delta, \delta' > \varepsilon$ with $\delta \sqcap \delta' = \varepsilon$ (which implies $\delta, \delta' < 1$). If π has both the (k, δ)- and the (k, δ')-extension property, then Corollary 7.3.8 implies $\pi[\![\psi]\!] \le \delta$ and $\pi[\![\psi]\!] \le \delta'$, so $\pi[\![\psi]\!] \le \delta \sqcap \delta' = \varepsilon$. Since δ and δ' are p-relevant, both extension properties almost surely hold. We further have $\pi[\![\psi]\!] \ge \varepsilon$ by Lemma 7.3.9, since $p[\varepsilon \le x \le 1] = 1$ and the interval $[\varepsilon, 1]$ is closed (and hence directed). Combining both bounds yields

$$\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] \in J] = \begin{cases} 1 & \text{if } \varepsilon \in J, \\ 0 & \text{otherwise.} \end{cases}$$

If no such δ, δ' exist, then the interval $J_\varepsilon = (\varepsilon, 1]$ is directed. Since p only admits values in J_ε, Lemma 7.3.9 implies $\mu_{n,p}[\pi[\![\psi]\!] \in J_\varepsilon] = 1$ for all n. Together with $(*)$, we get

$$\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] \in J] = \begin{cases} 1 & \text{if } (\varepsilon, \delta) \subseteq J \text{ for some } \delta > \varepsilon \\ 0 & \text{otherwise.} \end{cases}$$

We remark that the intervals (ε, δ) are nonempty (see the proof of Corollary 7.3.11). □

Corollary 7.3.11 (Almost sure valuations)**.** *For every infinite lattice semiring K with ε-bounded probability measure p, where $\varepsilon \ll 1$, and every relational vocabulary τ, the only possible almost sure valuations are* $\mathrm{ASV}_{K,p}(L^\omega_{\infty\omega}(\tau) \cup \mathrm{LFP}(\tau)) = \{0, 1, \varepsilon\}$.

Proof. In the cases where $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] \in J] = 1$ holds whenever $0 \in J$, $1 \in J$, or $\varepsilon \in J$, we clearly have $\mathrm{ASV}_{K,p}(\psi) = 0$, 1, or ε, respectively.

In the only remaining case, p is strictly ε-bounded and we have $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] \in J] = 1$ exactly if $(\varepsilon, \delta) \subseteq J$ for some $\delta > \varepsilon$. We claim that $\varepsilon \ll \gamma$ for every $\gamma > \varepsilon$ (in particular, (ε, δ) is nonempty). This is true by assumption for $\gamma = 1$, so we only consider $\gamma < 1$. If there would be a smallest $\gamma > \varepsilon$, then p would be weakly γ bounded, a contradiction. If there would be two minimal $\gamma, \gamma' > \varepsilon$, then $\gamma \sqcap \gamma' = \varepsilon$ and we would be in the case where $\varepsilon \in J$ (see the proof of Corollary 7.3.10). Hence the claim holds.

Assume that the almost sure valuation exists, so $\mathrm{ASV}_{K,p}(\psi) = j$ for some $j \in K$. Then there is a sequence of intervals with $\bigcap_{i<\omega} J_i = \{j\}$ and $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] \in J_i] = 1$ for all i. Clearly $j \geq \varepsilon$, as every J_i must contain a nonempty interval (ε, δ) for some $\delta > \varepsilon$. Assume towards a contradiction that $j > \varepsilon$. By the claim, there is some $j > \gamma' > \varepsilon$. But then there must be an i such that $\gamma' \notin J_i$, as otherwise $\gamma' \in \bigcap_{i<\omega} J_i$. Since J_i must contain $j > \gamma'$, this means that J_i cannot intersect (ε, γ'). This leads to a contradiction, since $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] \in (\varepsilon, \gamma')] = 1$ and hence $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] \in J_i] = 0$. □

We remark that the almost sure valuations of sentences can be different in *finite* and *infinite* min-max semirings. The polynomials are the same, and hence $f_\psi \in \{0, 1\}$ implies $\mathrm{ASV}_{K,p}(\psi) = f_\psi$ in both cases, but the values can differ in case of $f_\psi = e$ if p admits arbitrarily small positive values.

Example 7.3.12. Consider the semiring $K = ([0,1], \max, \min, 0, 1)$ over real numbers. We define a discrete probability distribution $p\colon K^+ \to [0,1]$ by $p(\frac{1}{2^n}) = \frac{1}{2^{n+1}}$ for all $n \in \mathbb{N}$, and $p(x) = 0$ otherwise. Then $p(1) = \frac{1}{2} > 0$ and p is ε-bounded for $\varepsilon = 0$, as the values $\frac{1}{2^n}$ with positive probability become arbitrarily small.

The sentence $\psi = \forall x(Px \vee \neg Px)$ induces $f_\psi = e$ and is clearly (almost surely) true in the Boolean semiring. However, since p is 0-bounded, we have $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] \in J_i] = 1$ for the intervals $J_i = [0, \frac{1}{2^i}]$ and hence $\mathrm{ASV}_{K,p}(\psi) = 0$.

7.4 Zero-One Laws Beyond Lattice Semirings

7.4.1 Absorptive Semirings

We now generalize our results beyond min-max and lattice semirings to the more general class of absorptive semirings $(K, +, \cdot, 0, 1)$, including among others the Viterbi and tropical semirings $\mathbb{V}$ and $\mathbb{T}$, the Łukasiewicz semiring $\mathbb{L}$, and generalised absorptive polynomials $\mathbb{S}^\infty(X)$. Recall that absorptive semirings are always naturally ordered, by $a \leq_K b \Leftrightarrow a+b=b$, and that addition coincides with the supremum of $\leq_K$. In contrast, multiplication can be different from the infimum $\sqcap$, but is guaranteed to be decreasing due to absorption, i.e., $ab \leq a$ and hence $ab \leq a \sqcap b$. For a well-defined semantics of infinitary logic (cf. Section 5.3), we additionally require K to be fully continuous, so that infinitary disjunctions and conjunctions can be defined through the infinitary operations in Section 3.4.2.

In order to lift our results from lattice semirings to absorptive semirings, we associate with every absorptive semiring $(K, +, \cdot, 0, 1)$ the *underlying lattice semiring* $K_{\text{inf}} = (K, +, \sqcap, 0, 1)$ over the same domain that replaces multiplication with the infimum-operation of the natural order. Let $\pi\colon \text{Lit}_A(\tau) \to K$ be a K-interpretation into an absorptive semiring K. Since K and K_{inf} have the same domain, we can view it also as an interpretation π_{inf} into K_{inf}, with $\pi_{\text{inf}}(\beta) = \pi(\beta)$ for all $\beta \in \text{Lit}_A(\tau)$. Since also the natural order $\leq_K$ is the same for K and K_{inf}, we can compare the semiring values $\pi[\![\psi]\!]$ and $\pi_{\text{inf}}[\![\psi]\!]$.

Proposition 7.4.1. *Let K be an absorptive, fully-continuous semiring, and let π be a K-interpretation over finite universe A. For every formula $\psi(\boldsymbol{x}) \in L^\omega_{\infty\omega}$ and every tuple $\boldsymbol{a} \subseteq A$, we have*

- $\pi[\![\psi(\boldsymbol{a})]\!] \leq_K \pi_{\text{inf}}[\![\psi(\boldsymbol{a})]\!]$, *and*
- $\pi_{\text{inf}}[\![\psi(\boldsymbol{a})]\!] = 1$ *if, and only if,* $\pi[\![\psi(\boldsymbol{a})]\!] = 1$,

Proof. The first statement readily follows by induction. Literals are mapped to the same values in π and π_{inf}. For disjunctions and existential quantification the induction step is trivial, since these are interpreted by the $+$ operation (supremum) in both K and K_{inf}. For conjunctions $\psi = \bigwedge_{i\in I} \varphi_i$, we have $\pi[\![\bigwedge_i \varphi_i]\!] = \widehat{\prod_i} \pi[\![\varphi_i]\!] \leq \bigsqcap_i \pi[\![\varphi_i]\!] \leq \bigsqcap_i \pi_{\text{inf}}[\![\varphi_i]\!] = \pi_{\text{inf}}[\![\bigwedge_i \varphi_i]\!]$ by absorption, and analogously for universal quantification.

For the second statement, it remains to prove that $\pi_{\text{inf}}[\![\psi(\boldsymbol{a})]\!] = 1$ implies $\pi[\![\psi(\boldsymbol{a})]\!] = 1$. This again follows by induction. For conjunctions, observe that $\pi_{\text{inf}}[\![\bigwedge_i \varphi_i]\!] = 1$ implies $\pi_{\text{inf}}[\![\varphi_i]\!] = 1$ for all i, and hence also $\pi[\![\bigwedge_i \varphi_i]\!] = 1$ by induction. The same argument applies for universal quantification, the other cases are trivial. □

We remark that the equivalence that we have for the value 1 also holds for 0 if the semiring has no divisors of 0, but not in general. For instance, an interpretation into the Łukasiewicz semiring $\mathbb{L}$ that interprets two literals α and β by values in the open interval $(0, \frac{1}{2})$, interprets the conjunction $\alpha \wedge \beta$ by 0, whereas the associated interpretation into $\mathbb{L}_{\text{inf}}$ picks the smaller of the two values.

Proposition 7.4.1 implies that every sentence ψ whose almost sure valuation $\text{ASV}_{K_{\text{inf}},p}(\psi)$ is 0 or 1, has the same almost sure valuation in K. However, if $\text{ASV}_{K_{\text{inf}},p}(\psi) = \varepsilon$, with $\varepsilon > 0$, the situation is more complicated, since K need not be mulitiplicatively idempotent. Hence, even if ε is the smallest positive value that may appear for valuations $\pi(\alpha)$ of literals, more complicated formulae may get smaller valuations.

Example 7.4.2 (Values $< \varepsilon$)**.** For a simple example, consider a universal sentence $\psi = \forall y(Py \vee \neg Py)$ and a random interpretation in the Viterbi semiring $\mathbb{V} = ([0,1]_{\mathbb{R}}, \max, \cdot, 0, 1)$, with a probability distribution p satisfying $p[0 < x < \frac{1}{2}] = 0$, and $p[x = \frac{1}{2}] > 0$. Here, $\varepsilon = \frac{1}{2}$ and in the associated min-max semiring on $[0,1]$, we clearly have that ψ asymptotically almost surely evaluates to ε. But for random interpretations in the Viterbi semiring the valuation of ψ asymptotically gets arbitrary small: for every $\delta > 0$ we have that $\lim_{n\to\infty} \mu_{n,p}[0 < \pi[\![\psi]\!] < \delta] = 1$ and hence $\mathrm{ASV}_{\mathbb{V},p}(\psi) = 0$.

Example 7.4.3 (Absorptive polynomials)**.** A perhaps more unusual, but also more interesting example is obtained by evaluating the same formula ψ in the semiring $\mathbb{S}^\infty(x)$ of generalised absorptive polynomials with just one indeterminate (in this case, the natural order is total: $1 > x > x^2 > \cdots > x^\infty > 0$), under the probability distribution that assigns to each pair of complementary literals $Pj, \neg Pj$ with equal probability $\frac{1}{4}$ a pair of values from $\{(1,0), (x,0), (0,1), (0,x)\}$. In other words, we first decide for each atom, independently and with uniform probability $\frac{1}{2}$, whether it is true or false, and then, again independently and with uniform probability $\frac{1}{2}$, whether or not we want to track the effect of this decision for the valuation of the formulae we consider.

The valuation $\pi[\![\varphi]\!] \in \mathbb{S}^\infty(x)$ then is either 0, 1, or a monomial x^k, for $k \in \mathbb{N}^\infty$, which tells us how many tracked literals are needed for establishing the truth of φ. For the given probability distribution p, we have that $\varepsilon = x$, and indeed, for any natural number n and every element $j \in [n]$, the probability measures $\mu_{n,p}$ evaluate the formula $P(j) \vee \neg P(j)$ either to 1 or to x, each with probability $\frac{1}{2}$. As a consequence, we have for $\psi = \forall y(Py \vee \neg Py)$ that $\lim_{n\to\infty} \mu_{n,p}[0 < \pi[\![\psi]\!] < x^k] = 1$, for all $k \in \mathbb{N}$. Hence the almost sure valuation of ψ is $\mathrm{ASV}_{\mathbb{S}^\infty(x),p}(\psi) = x^\infty$.

We can nevertheless show that almost sure valuations of ε transfer from K_{inf} to K, under the assumption that ε is idempotent ($\varepsilon \cdot \varepsilon = \varepsilon$). This applies, for instance, to the smallest nonzero element x^∞ of $\mathbb{S}^\infty(x)$ (and also to the multivariate case, say $x^\infty y^\infty$ in $\mathbb{S}^\infty(x,y)$).

Proposition 7.4.4. *Let K be an absorptive, fully-continuous semiring, and let π be a K-interpretation over finite universe A. Let further $\varepsilon \in K$ such that $\varepsilon \cdot \varepsilon = \varepsilon$. Then for every formula $\psi(\boldsymbol{x}) \in L^\omega_{\infty\omega}$ and every tuple $\boldsymbol{a} \subseteq A$, we have that $\pi_{\text{inf}}[\![\psi(\boldsymbol{a})]\!] = \varepsilon$ implies $\pi[\![\psi(\boldsymbol{a})]\!] = \varepsilon$.*

Proof. By Proposition 7.4.1, it suffices to prove by induction on ψ that $\pi_{\text{inf}}[\![\psi(\boldsymbol{a})]\!] \geq \varepsilon$ implies $\pi[\![\psi(\boldsymbol{a})]\!] \geq \varepsilon$. For literals, disjunctions and existential quantification, this is trivial. For conjunctions, observe that $\pi_{\text{inf}}[\![\bigwedge_{i\in I} \varphi_i]\!] \geq \varepsilon$ implies $\pi_{\text{inf}}[\![\varphi_i]\!] \geq \varepsilon$ for all i. Since infinite products in K preserve idempotent elements (Theorem 3.4.2), we then have $\pi[\![\bigwedge_{i\in I} \varphi_i]\!] \geq \widehat{\prod}_{i\in I} \varepsilon = \varepsilon$ by induction and monotonicity. Analogously for universal quantification. □

With this assumption, we can lift the 0-1 laws for finite and infinite lattice semirings to absorptive semirings, leading to the following result about the almost sure valuations.

Corollary 7.4.5. *Let K be an absorptive, fully-continuous semiring, let τ be a finite relational vocabulary, and let $\psi \in L^{\omega}_{\infty\omega}(\tau) \cup \mathrm{LFP}(\tau)$ be a sentence. If $\mathrm{ASV}_{K_{\mathrm{inf}},p}(\psi) \in \{0,1\}$ or $\mathrm{ASV}_{K_{\mathrm{inf}},p}(\psi) \in \{0,1,\varepsilon\}$ with $\varepsilon \cdot \varepsilon = \varepsilon$ (in K), then $\mathrm{ASV}_{K,p}(\psi) = \mathrm{ASV}_{K_{\mathrm{inf}},p}(\psi)$.*

Remark 7.4.6. The assumption of full continuity is only required for infinitary formulae and LFP. For first-order formulae, all of the results in this section hold for arbitrary absorptive semirings K, without requiring continuity (cf. [GHNW22a]).

7.4.2 The Natural Semiring

We now discuss the natural semiring $(\mathbb{N}, +, \cdot, 0, 1)$, which is important for bag semantics in databases, and its fully-continuous completion $(\mathbb{N}^\infty, +, \cdot, 0, 1)$. We mostly work in the semiring $\mathbb{N}^\infty$ which allows natural definitions of infinitary operations (see Section 3.4.3). However, we make no assumptions on the probability of the value ∞, and hence our results also apply to $\mathbb{N}$ when restricted to first-order logic.

The most important technical difference to the previously considered semirings is that multiplication is now increasing rather than decreasing with respect to the natural order, which leads to a different asymptotic behaviour of universal quantification. To adapt our proof, we first define a different extension property which is both stronger and weaker compared to the extension properties for lattice semirings: stronger, since it guarantees not just one, but many realisations of extension types, but weaker, since it does not guarantee realisations of every type (which would be infinitely many), but only that every underlying Boolean type has realisations with sufficiently large values.

Definition 7.4.7. Given an atomic m-type ρ with values in $\mathbb{N}^\infty$ we call an extension $\rho' \in \mathrm{ext}(\rho)$ *large* if out of any pair $\alpha, \neg\alpha$ of complementary literals containing the variable x_{m+1}, it maps one of them to 0, and to other to some number ≥ 2.

Recall that $\rho =_{\mathbb{B}} \rho'$ holds if ρ and ρ' induce the same Boolean type. Let $\gamma > 0$ be some constant. We say that an $\mathbb{N}^\infty$-interpretation $\pi\colon \mathrm{Lit}_A(\tau) \to \mathbb{N}^\infty$ has the *strong (k,γ)-extension property* if for every $m < k$, every tuple $\boldsymbol{a} \in A^m$ and every extension $\rho^+ \in \mathrm{ext}(\rho^\pi_{\boldsymbol{a}})$,

$$\big|\{b \in A \setminus \boldsymbol{a} \mid \rho^\pi_{\boldsymbol{a},b} =_{\mathbb{B}} \rho^+ \text{ and } \rho^\pi_{\boldsymbol{a},b} \text{ is large}\}\big| \;\geq\; \gamma|A|.$$

We consider probability distributions $p\colon \mathbb{N}^\infty_{>0} \to [0,1]$ with the property that $p[x \geq 2] > 0$ and the associated measures $\mu_{n,p}$ on $\mathbb{N}^\infty$-interpretations of τ-structures with universe n. Associated random $\mathbb{N}^\infty$-interpretations almost surely have strong extension properties.

Proposition 7.4.8. *Let $p\colon \mathbb{N}^\infty_{>0} \to [0,1]$ be a probability distribution with $p[x \geq 2] > 0$. For every $k < \omega$ there exists a real number $\gamma > 0$ such that*

$$\lim_{n\to\infty} \mu_{n,p}\big[\pi \text{ has the strong } (k,\gamma)\text{-extension property}\big] = 1.$$

This follows by general results of probability theory that have, for instance, been used also by Blass and Gurevich [BG03] to prove strong extension properties of random graphs. Specifically, we can apply the following fact, see [BG03, Lemma 6.2].

Lemma 7.4.9. *Let X be the number of successes in n trials, each having at least probability δ of success. Then, for each $\alpha \in (0,1)$ there exists some $\beta \in (0,1)$ such that for all natural numbers n, $Prob[X \leq \alpha\delta n] \leq \beta^n$.*

Proof of Proposition 7.4.8. Let $q = p[x \geq 2] > 0$. For every tuple $\boldsymbol{a} \in [n]^m$, and every new element $b \in [n] \setminus \boldsymbol{a}$, the probability that $\rho^\pi_{\boldsymbol{a},b}$ is large and $\rho^\pi_{\boldsymbol{a},b} =_{\mathbb{B}} \rho^+$ is at least $\delta = (q/2)^\ell$ where ℓ is the number of relational atoms containing the variable x_{m+1}. Fix any $\alpha \in (0,1)$ and choose γ with $0 < \gamma < \alpha\delta$. For large enough n, the probabilty that $\rho^\pi_{\boldsymbol{a}}$ does not have at at least γn extensions to a type $\rho^\pi_{\boldsymbol{a},b} \geq \rho^+$ is then bounded by β^{n-m}, for some $\beta < 1$. There are n^m tuples $\boldsymbol{a}$ to consider and $2^\ell =_{\mathbb{B}}$ equivalence classes of extensions ρ^+. Thus the probability that the strong k-extension property fails for π is bounded by $n^m 2^\ell \beta^{n-m}$ which converges to 0 exponentially fast. □

We again want to represent formulae $\psi(x_1, \dots, x_i) \in L^k_{\infty\omega}$ by algebraic expressions $g(\boldsymbol{X}^{(i)})$ with indeterminates X_α and $X_{\neg\alpha}$, for each τ-atom in variables from $x_1, \dots, x_i$. However, instead of the absorptive polynomials in the case of lattice semirings, we need here a slightly different definition to include ∞ as a coefficient and exponent:

Definition 7.4.10. An ∞-*expression* (over $\boldsymbol{X}^{(i)}$) is a formal arithmetic expression (of infinite size) $g(\boldsymbol{X}^{(i)})$ consisting of indeterminates $X_\alpha, X_{\neg\alpha} \in \boldsymbol{X}^{(i)}$, constants $0, 1, \infty$, infinitary operations $\sum, \prod$ and the unitary operation $^\infty$ (infinite power).

Given a mapping $\sigma\colon \boldsymbol{X}^{(i)} \to \mathbb{N}^\infty$, the ∞-expression $g(\boldsymbol{X}^{(i)})$ evaluates to $g[\sigma] \in \mathbb{N}^\infty$ with the usual rules, extended by $0^\infty = 0$, $1^\infty = 1$, and $n^\infty = \infty$ for $n \geq 2$. Two ∞-expressions $g(\boldsymbol{X}^{(i)})$ and $g'(\boldsymbol{X}^{(i)})$ are *equivalent*, denoted $g \equiv g'$, if $g[\sigma] = g'[\sigma]$ for every consistent mapping $\sigma\colon \boldsymbol{X}^{(i)} \to \mathbb{N}^\infty$ (that is, out of any pair $X_\alpha, X_{\neg\alpha}$, one is mapped to 0 and the other one to a nonzero value).

Remark 7.4.11. It may seem tempting to use power series $\mathrm{Poly}(\mathbb{N}^\infty, \mathbb{N}^\infty, \boldsymbol{X}^{(i)})$ instead of the more complicated ∞-expressions. The main reason for working with ∞-expressions is the exponent ∞. Indeed, it is not clear to what power series an expression such as $(x+y)^\infty$ should evaluate. In absorptive semirings we have $(x+y)^\infty = x^\infty + y^\infty$, but this does not work here, because $(1+1)^\infty = \infty \neq 2 = 1^\infty + 1^\infty$. Instead, we keep the expression $(x+y)^\infty$ in its unsimplified form.

Lemma 7.4.12. *Let $\sigma, \sigma'\colon \boldsymbol{X}^{(i)} \to \mathbb{N}^\infty$ and let $g(\boldsymbol{X}^{(i)})$ be an ∞-expression. Then,*

- *if $\sigma =_{\mathbb{B}} \sigma'$, then $g[\sigma] = 0$ if and only if $g[\sigma'] = 0$,*
- *if $\sigma \leq \sigma'$, then also $g[\sigma] \leq g[\sigma']$.*
- *if g is not constant on atomic types, then it assumes arbitrarily large values: for every $n \in \mathbb{N}$ there exists a type σ_n with $g[\sigma_n] \geq n$.*

Proof. The first two claims follow by a straightforward induction on g, since the operations $\sum, \prod, ^\infty$ are monotone. For the third claim, assume that there exist atomic types σ, σ' with $g[\sigma] > g[\sigma']$. We can then, without loss of generality, choose σ and σ' so that they differ on precisely one pair $\alpha, \neg\alpha$ of complementary literals, i.e. $\sigma(\beta) = \sigma'(\beta)$ for all $\beta \notin \{\alpha, \neg\alpha\}$.

Further we assume that $\sigma(\alpha) > \sigma'(\alpha)$. For each n, we then consider the type σ_n such that $\sigma_n(\beta) = \sigma(\beta)$ for all $\beta \neq \alpha$ and $\sigma_n(\alpha) = \max(n, \sigma(\alpha))$. We claim that, for every ∞-expression f with $f[\sigma] > f[\sigma']$, we have that $f[\sigma_n] \geq n$. We prove the claim by induction on f.

- The only atomic expression f with $f[\sigma] > f[\sigma']$ is $f = X_\alpha$, for which $f[\sigma_n] \geq n$.
- If $f = \sum_{i \in I} f_i$ then we must have $f_i[\sigma] > f_i[\sigma']$ for some i. Then $f_i[\sigma_n] \geq n$ by induction, and hence also $f[\sigma_n] \geq n$.
- If $f = \prod_{i \in I} f_i$ then we must again have $f_i[\sigma] > f_i[\sigma']$ for some i, so $f_i[\sigma_n] \geq n$. Moreover, we must have $f_j[\sigma] > 0$ for all $j \neq i$, and since $\sigma \leq \sigma_n$ this implies $f_j[\sigma_n] > 0$. Altogether, this implies $f[\sigma_n] \geq n$.
- Finally, for $f = \infty \cdot h$ or $f = h^\infty$ we have that $f[\sigma] > f[\sigma']$ implies that $h[\sigma] > h[\sigma']$ and hence $h[\sigma_n] \geq n$, which implies that $f[\sigma_n] = \infty$. □

Given a mapping $\rho\colon \boldsymbol{X}^{(i)} \to \mathbb{N}^\infty$ and a selector function $s\colon \boldsymbol{Y}^{(i)} \to \mathbb{N}^\infty$, we write ρs for the combined mapping $\rho s\colon \boldsymbol{X}^{(i+1)} \to \mathbb{N}^\infty$ that behaves like ρ on $\boldsymbol{X}^{(i)}$ and like s on $\boldsymbol{Y}^{(i)} = \boldsymbol{X}^{(i+1)} \setminus \boldsymbol{X}^{(i)}$. We also need the following lemma on ∞-expressions. Notice that this applies in particular to $\sigma = \rho^\pi_{\boldsymbol{a}} s$ (the type $\rho^\pi_{\boldsymbol{a}}$ extended by a selector function s) and $\sigma' = \rho^\pi_{\boldsymbol{a},b}$. Indeed, if $\rho^\pi_{\boldsymbol{a},b} =_{\mathbb{B}} \rho^\pi_{\boldsymbol{a}} s$ and $\rho^\pi_{\boldsymbol{a},b}$ is a large extension of $\rho^\pi_{\boldsymbol{a}}$, the condition in the lemma is satisfied.

Lemma 7.4.13. *Let $g(\boldsymbol{X}^{(i)})$ be an ∞-expression and consider $\sigma, \sigma'\colon \boldsymbol{X}^{(i)} \to \mathbb{N}^\infty$ such that $\sigma =_{\mathbb{B}} \sigma'$ and for all $X_\beta \in \boldsymbol{X}^{(i)}$, either $\sigma(X_\beta) = \sigma'(X_\beta)$ or $\sigma'(X_\beta) \geq 2$. Then $g[\sigma] \geq 2$ implies $g[\sigma'] \geq 2$.*

Proof. By induction on g. The claim is trivial for constants.

- If $g = X_\beta$, then either $g[\sigma] = g[\sigma']$ or $g[\sigma'] \geq 2$, so the claim holds.
- If $g = \prod_{i \in I} g_i$ and $g[\sigma] \geq 2$, then we must have $g_i[\sigma] \geq 2$ for some i and $g_j[\sigma] \neq 0$ for all $j \neq i$. By induction and $\sigma =_{\mathbb{B}} \sigma'$, the same holds for $g_i[\sigma']$ and all $g_j[\sigma']$, so we have $g[\sigma'] \geq 2$ as claimed.
- If $g = h^\infty$ and $g[\sigma] \geq 2$, then also $h[\sigma] \geq 2$ and the claim follows by induction.
- If $g = \sum_{i \in I} g_i$ and $g[\sigma] \geq 2$, we distinguish two cases. First assume that there are $i_1, i_2 \in I$ with $i_1 \neq i_2$ and $g_{i_1}[\sigma] = g_{i_2}[\sigma] = 1$. Since $\sigma =_{\mathbb{B}} \sigma'$, it follows that $g_{i_1}[\sigma'], g_{i_2}[\sigma'] \neq 0$ and hence $g[\sigma'] \geq 2$. Otherwise, there must be an $i \in I$ with $g_i[\sigma] \geq 2$ and the claim follows by induction. □

The definition of the arithmetic expressions $g_\psi(\boldsymbol{X}^{(i)})$ is to some extent analogous to the polynomials $f_\psi(\boldsymbol{X}^{(i)})$ for lattice semirings, but the algebraic operations are no longer idempotent and the rules for the quantifiers are different and use the constant ∞.

Definition 7.4.14. Given $\psi(x_1, \dots, x_i) \in L^\omega_{\infty\omega}(\tau)$ in negation normal form and written with the excluding quantifiers $\exists^{\neq}$ and $\forall^{\neq}$, we define the associated ∞-expression $g_\psi(\boldsymbol{X}^{(i)})$ inductively as follows.

- For (in)equalities, literals, disjunctions and conjunctions, the definition is identical to Definition 7.2.5. That is, we set $g_\psi \in \{0,1\}$ for (in)equalities, as well as $g_\psi := X_\alpha$ and $g_\psi := X_{\neg\alpha}$ for (negated) atoms. For $\psi = \bigvee_{i \in I} \varphi_i$ and $\psi' = \bigwedge_{i \in I} \varphi'_i$, we further set $g_\psi := \sum_{i \in I} g_{\varphi_i}$ and $g_{\psi'} := \prod_{i \in I} g_{\varphi'_i}$.
- For $\psi(\boldsymbol{x}) = \exists^{\neq} y\, \varphi(\boldsymbol{x}, y)$ and $\psi'(\boldsymbol{x}) = \forall^{\neq} y\, \varphi'(\boldsymbol{x}, y)$, let $\boldsymbol{Y}^{(i)} = \boldsymbol{X}^{(i+1)} \setminus \boldsymbol{X}^{(i)}$. Let further S be the set of all consistent selector functions $s\colon \boldsymbol{Y}^{(i)} \to \{0, \infty\}$ (that is, one of X_α, $X_{\neg\alpha}$ is mapped to 0, the other one to ∞, for every atom α). Now set

$$g_\psi(\boldsymbol{X}^{(i)}) := \infty \cdot \left(\sum_{s \in S} g_\varphi(\boldsymbol{X}^{(i)}, s(\boldsymbol{Y}^{(i)})) \right), \quad g_{\psi'}(\boldsymbol{X}^{(i)}) := \left(\prod_{s \in S} g_{\varphi'}(\boldsymbol{X}^{(i)}, s(\boldsymbol{Y}^{(i)})) \right)^{\infty}.$$

Notice that for quantifiers, a single positive value in the sum or a single value ≥ 2 in the product will result in the value ∞. This is justified by the (k, γ)-extension property, which guarantees that every selector function has not just one, but many large realisations which, as n grows, leading to arbitrarily large values of ψ or ψ'.

Example 7.4.15. Recall $\psi = \exists^{\neq} x(\neg Exx \wedge \forall^{\neq} y(Exy \vee (\neg Exy \wedge \exists^{\neq} z(Exz \wedge Ezy))))$ of Example 7.2.8. Using the same notation, we obtain the following ∞-expressions (we always simplify expressions without indeterminates, e.g. $\infty \cdot (\infty + 0) = \infty$ in the first step).

φ	g_φ
$Exz \wedge Ezy$	ZU
$\exists^{\neq} z(Exz \wedge Ezy)$	∞
$\neg Exy \wedge \exists^{\neq} z(Exz \wedge Ezy)$	$\overline{Y} \cdot \infty$
$Exy \vee (\neg Exy \wedge \exists^{\neq} z(Exz \wedge Ezy))$	$Y + \overline{Y} \cdot \infty$
$\forall^{\neq} y(Exy \vee (\neg Exy \wedge \exists^{\neq} z(Exz \wedge Ezy)))$	∞
$\neg Exx \wedge \forall^{\neq} y(Exy \vee (\neg Exy \wedge \exists^{\neq} z(Exz \wedge Ezy)))$	$\overline{X} \cdot \infty$
ψ	∞

It should come as no surprise that the resulting value is positive, since ψ is asymptotically almost surely true in Boolean semantics. But notice that Example 7.2.8 resulted in e, representing the smallest positive value, whereas we obtain the largest value ∞ in the natural semiring.

Example 7.4.16. For an example with more complicated ∞-expressions, consider $\psi = \exists^{\neq} x(Px \wedge \forall^{\neq} y(x \neq y \vee \neg Px))$. Using the indeterminate X for Px, we obtain:

φ	g_φ
$x \neq y \vee \neg Px$	$1 + \overline{X}$
$\forall^{\neq} y(x \neq y \vee \neg Px)$	$(1 + \overline{X})^\infty$
$Px \wedge \forall^{\neq} y(x \neq y \vee \neg Px)$	$X \cdot (1 + \overline{X})^\infty$
ψ	∞

For an example resulting in 0, replace the subformula $x \neq y$ by $x = y$ to obtain $X \cdot (0 + \overline{X})^\infty$ in the third and thus 0 in the last row.

The main technical result of this section is the following theorem, similar to Theorems 7.2.9 and 7.3.7. We again prove that the expressions g_ψ provide an adequate description of formulae $\psi(\boldsymbol{x})$, now including the special case where the value of ψ becomes arbitrarily large.

Theorem 7.4.17. *Let $\gamma > 0$ and let $\pi\colon \mathrm{Lit}_A(\tau) \to \mathbb{N}^\infty$ be an $\mathbb{N}^\infty$-interpretation on universe $A = [n]$ with the strong (k,γ)-extension property. Further assume that $\gamma n \geq 2$. Then, for every formula $\psi(x_1,\dots,x_i) \in L^k_{\infty\omega}(\tau)$ with the associated ∞-expression g_ψ and every tuple $\boldsymbol{a}$ of distinct elements from A either*

- $\pi[\![\psi(\boldsymbol{a})]\!] = g_\psi[\rho^\pi_{\boldsymbol{a}}] \in \mathbb{N}$, *or* **(1)**
- $g_\psi[\rho^\pi_{\boldsymbol{a}}] = \infty$ and $\pi[\![\psi(\boldsymbol{a})]\!] \geq \gamma n$. **(2)**

Proof. We proceed by induction on ψ. If ψ is a literal (or an (in)equality), it is immediate from the definition of g_ψ that **(1)** holds if $\pi[\![\psi(\boldsymbol{a})]\!] < \infty$, and **(2)** holds if $\pi[\![\psi(\boldsymbol{a})]\!] = \infty$.

For $\psi = \bigvee_{i\in I} \varphi_i$, we have $g_\psi = \sum_i g_{\varphi_i}$. If **(1)** holds for all i, then also $g_\psi[\rho^\pi_{\boldsymbol{a}}] = \pi[\![\psi(\boldsymbol{a})]\!]$ and one of the two cases applies. Otherwise, **(2)** applies to some φ_i, and then also to ψ.

For $\psi = \bigwedge_{i\in I} \varphi_i$, we have $g_\psi = \prod_{i\in I} g_{\varphi_i}$. If $\pi[\![\varphi_i(\boldsymbol{a})]\!] = 0$ for some i, then **(1)** must hold for φ_i (recall that $\gamma n > 0$) and hence also $g_\psi[\rho^\pi_{\boldsymbol{a}}] = \pi[\![\psi(\boldsymbol{a})]\!] = 0$. Otherwise, the argument is identical to the one for disjunctions.

Let now $\psi(\boldsymbol{x}) = \exists^{\neq} y\, \varphi(\boldsymbol{x}, y)$. We will show that case **(1)** applies if $\pi[\![\psi(\boldsymbol{a})]\!] = g_\psi[\rho^\pi_{\boldsymbol{a}}] = 0$, otherwise case **(2)** applies. We recall:

$$\pi[\![\psi(\boldsymbol{a})]\!] = \sum_{b\in A\setminus\boldsymbol{a}} \pi[\![\varphi(\boldsymbol{a}, b)]\!], \qquad g_\psi[\rho^\pi_{\boldsymbol{a}}] = \infty \cdot \left(\sum_{s\in S} g_\varphi[\rho^\pi_{\boldsymbol{a}}, s(\boldsymbol{Y}^{(i)})]\right).$$

We begin with a general observation that is used throughout the proof: For every element $b \in A \setminus \boldsymbol{a}$, we can define a selector function $s_b \in S$ with $\rho^\pi_{\boldsymbol{a},b} =_{\mathbb{B}} \rho^\pi_{\boldsymbol{a}} s_b$ by simply setting $s_b(X_\beta) = \infty$ precisely if $\rho^\pi_{\boldsymbol{a},b}(\beta) > 0$. Conversely, every selector function $s \in S$ is consistent, so there is a type $\rho^+ \in \mathrm{ext}(\rho^\pi_{\boldsymbol{a}})$ with $\rho^+ =_{\mathbb{B}} \rho^\pi_{\boldsymbol{a}} s$ (we may set $\rho^+(\beta) = 1$ whenever $s(X_\beta) = \infty$). The (k,γ)-extension property then guarantees at least γn many elements $b \in A \setminus \boldsymbol{a}$ such that $\rho^\pi_{\boldsymbol{a},b} =_{\mathbb{B}} \rho^\pi_{\boldsymbol{a}} s$ and $\rho^\pi_{\boldsymbol{a},b}$ is large.

First assume that there is a selector function $s \in S$ with $g_\varphi[\rho^\pi_{\boldsymbol{a}}, s(\boldsymbol{Y}^{(i)})] > 0$. Then obviously $g_\psi[\rho^\pi_{\boldsymbol{a}}] = \infty$. There are at least γn many elements b with $\rho^\pi_{\boldsymbol{a},b} =_{\mathbb{B}} \rho^\pi_{\boldsymbol{a}} s$ as descried above. By Lemma 7.4.12, this implies $g_\varphi[\rho^\pi_{\boldsymbol{a},b}] > 0$. Then also $\pi[\![\varphi(\boldsymbol{a}, b)]\!] > 0$ by induction (using either case **(1)** or **(2)**) and hence $\pi[\![\psi(\boldsymbol{a})]\!] \geq \gamma n$ for the sum, so case **(2)** holds.

Now assume that no such selector function exists, hence $g_\psi[\rho^\pi_{\boldsymbol{a}}] = 0$. If there was a b with $\pi[\![\varphi(\boldsymbol{a}, b)]\!] > 0$, then $g_\varphi[\rho^\pi_{\boldsymbol{a},b}] > 0$ by induction and the selector function s_b contradicts our assumption. Hence no such b exists and we have $\pi[\![\psi(\boldsymbol{a})]\!] = 0$, so case **(1)** applies.

Finally, let $\psi(\boldsymbol{x}) = \forall^{\neq} y\, \varphi(\boldsymbol{x}, y)$. Similarly to the previous case, we will show that case **(1)** only applies for the values $0, 1$, otherwise case **(2)** applies. We recall:

$$\pi[\![\psi(\boldsymbol{a})]\!] = \prod_{b\in A\setminus\boldsymbol{a}} \pi[\![\varphi(\boldsymbol{a}, b)]\!], \qquad g_\psi[\rho^\pi_{\boldsymbol{a}}] = \left(\prod_{s\in S} g_\varphi[\rho^\pi_{\boldsymbol{a}}, s(\boldsymbol{Y}^{(i)})]\right)^\infty.$$

We proceed with a similar case distinction, but additionally account for one of the factors being 0. To this end, assume there is $s \in S$ with $g_\varphi[\rho^\pi_{\boldsymbol{a}}, s(\boldsymbol{Y}^{(i)})] = 0$, hence also $g_\psi[\rho^\pi_{\boldsymbol{a}}] = 0$. We obtain an element b with $\rho^\pi_{\boldsymbol{a},b} =_{\mathbb{B}} \rho^\pi_{\boldsymbol{a}} s$ by the extension property. By Lemma 7.4.12 and case **(1)**, this yields $0 = g_\varphi[\rho^\pi_{\boldsymbol{a},b}] = \pi[\![\varphi(\boldsymbol{a}, b)]\!] = \pi[\![\varphi(\boldsymbol{a})]\!]$, so case **(1)** applies. Similarly, if $\pi[\![\varphi(\boldsymbol{a}, b)]\!] = 0$ for some b, then $g_\varphi[\rho^\pi_{\boldsymbol{a},b}] = 0$ by case **(1)** and $g_\varphi[\rho^\pi_{\boldsymbol{a}}, s_b(\boldsymbol{Y}^{(i)})] = 0$ for the induced selector function.

From now on, we thus have $\pi[\![\varphi(\boldsymbol{a}, b)]\!], g_\varphi[\rho^\pi_{\boldsymbol{a}}, s(\boldsymbol{Y}^{(i)})] > 0$ for all $s \in S$. First assume that there is a selector function s with $g_\varphi[\rho^\pi_{\boldsymbol{a}}, s(\boldsymbol{Y}^{(i)})] \geq 2$. Then $g_\psi[\boldsymbol{a}] \geq 2^\infty = \infty$. The extension property guarantees γn many elements b such that $\rho^\pi_{\boldsymbol{a},b}$ is large and $\rho^\pi_{\boldsymbol{a},b} =_{\mathbb{B}} \rho^\pi_{\boldsymbol{a}} s$. All of these elements satisfy $g_\varphi[\rho^\pi_{\boldsymbol{a},b}] \geq 2$ by Lemma 7.4.13 and thus $\pi[\![\varphi(\boldsymbol{a}, b)]\!] \geq 2$ by induction (using either case **(1)** or case **(2)**). We thus have $\pi[\![\psi(\boldsymbol{a})]\!] \geq 2^{\gamma n} \geq \gamma n$ and case **(2)** applies.

Lastly, the only remaining case is that $g_\varphi[\rho^\pi_{\boldsymbol{a}}, s(\boldsymbol{Y}^{(i)})] = 1$ for all selector functions. Then also $g_\psi[\boldsymbol{a}] = 1^\infty = 1$. If there was a b with $\pi[\![\varphi(\boldsymbol{a}, b)]\!] \geq 2$, then also $g_\varphi[\rho^\pi_{\boldsymbol{a},b}] \geq 2$ (using either case **(1)** or **(2)**). By Lemma 7.4.13, the induced selector function s_b contradicts our assumption: $g_\varphi[\rho^\pi_{\boldsymbol{a}}, s(\boldsymbol{Y}^{(i)})] \geq 2$. Hence $\pi[\![\varphi(\boldsymbol{a}, b)]\!] = 1$ for all b and thus $\pi[\![\psi(\boldsymbol{a})]\!] = 1$ as well. □

Corollary 7.4.18 (0-1 law for $L^\omega_{\infty\omega}$ on $\mathbb{N}^\infty$)**.** *Let $p\colon \mathbb{N}^\infty_{>0} \to [0, 1]$ be a probability distribution with $p[x \geq 2] > 0$, and let τ be a finite relational vocabulary.*

Then for every sentence $\psi \in L^\omega_{\infty\omega}(\tau)$ there either is a value $j \in \mathbb{N}$ such that we have $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] = j] = 1$, or $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] > j] = 1$ for all $j \in \mathbb{N}$.

Proof. For every sentence $\psi \in L^k_{\infty\omega}(\tau)$ the associated ∞-expression g_ψ contains no indeterminates and thus evaluates to a value $g_\psi[\emptyset] \in \mathbb{N}^\infty$. Now let $\gamma > 0$ so that random $\mathbb{N}^\infty$-interpretations almost surely have the strong (k, γ)-extension property. It follows that $\lim_{n\to\infty} \mu_{n,p}\big[\pi[\![\psi]\!] = g_\psi[\emptyset]\big] = 1$ in case $g_\psi[\emptyset] \in \mathbb{N}$, or $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] > j] = 1$ for all $j \in \mathbb{N}$, in case $g_\psi[\emptyset] = \infty$. □

We remark that the case $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] > j] = 1$ for all $j \in \mathbb{N}$ also covers the case where $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] = \infty] = 1$. A trivial example of such a formula is $\psi = \bigvee_{i<\omega} \forall x(x = x)$, where the infinite disjunction (over the same subformula) results in the value ∞ even on finite interpretations.

All of our arguments also hold for $\mathbb{N}$ instead of $\mathbb{N}^\infty$, so for first-order sentences (where the semantics is well-defined in $\mathbb{N}$), we immediately obtain the following 0-1 law.

Corollary 7.4.19 (0-1 law for FO on $\mathbb{N}$)**.** *Let $p\colon \mathbb{N}_{>0} \to [0, 1]$ be a probability distribution with $p[x \geq 2] > 0$, and let τ be a finite relational vocabulary.*

Then for every sentence $\psi \in \mathrm{FO}(\tau)$ there either is a value $j \in \mathbb{N}$ such that we have $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] = j] = 1$, or $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] > j] = 1$ for all $j \in \mathbb{N}$.

These 0-1 laws induce a partition of the sentences $L^\omega_{\infty\omega}(\tau)$ into classes $(\Phi_j)_{j\in\mathbb{N}^\infty}$. For $j \in \mathbb{N}$, the class Φ_j contains those sentences ψ for which $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] = j] = 1$ (and hence $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] > j] = 0$). The additional class Φ_∞ contains the sentences ψ which almost surely evaluate to unboundedly large values (including the value ∞), so $\lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] > j] = 1$ for all $j \in \mathbb{N}$. We observe that, as for lattice semirings, these classes align with the Boolean case.

Lemma 7.4.20. *For every sentence $\psi \in L^{\omega}_{\infty\omega}(\tau)$, we have $\psi \in \Phi_0$ if, and only if, ψ is asymptotically almost surely false under Boolean semantics.*

Proof. Recall that semiring semantics of $L_{\infty\omega}$ in $\mathbb{N}^\infty$ (or of FO in $\mathbb{N}$) is compatible with Boolean semantics in the sense that for any model-defining $\mathbb{N}^\infty$-interpretation π and any sentence ψ, we have $\pi[\![\psi]\!] > 0$ if, and only if, the induced Boolean structure satisfies $\mathfrak{A}_\pi \models \psi$ (see Corollary 5.1.6 and Proposition 5.3.4).

The measure $\mu_{n,p}$ over the semiring $\mathbb{N}^\infty$ is defined by first randomly deciding for each relational atom whether α or $\neg\alpha$ shall be true (with some constant probability, say $\frac{1}{2}$), and then assigning a random positive value from $\mathbb{N}^\infty_{>0}$ to the true literal. It thus makes no difference whether we consider random Boolean structures or the Boolean structures induced by random $\mathbb{N}^\infty$-interpretations. That is, $\mu_{n,\frac{1}{2}}(\psi) = \mu_{n,p}[\pi[\![\psi]\!] > 0]$ (recall that $\mu_{n,\frac{1}{2}}(\psi)$ denotes the probability that ψ holds in a random Boolean structure from $\mathrm{Str}_{n,\frac{1}{2}}(\tau)$). Hence also $\lim_{n\to\infty} \mu_{n,\frac{1}{2}}(\psi) = \lim_{n\to\infty} \mu_{n,p}[\pi[\![\psi]\!] > 0]$ and the claim follows. □

Moreover, there are trivial examples showing that all classes Φ_j are nonempty. Trivially false sentences such as $\exists x(x \neq x)$ are in Φ_0, whereas $\exists x(x = x) \in \Phi_\infty$, and for any $j \in \mathbb{N} > 0$, we have that $\bigvee_{1\leq i\leq j} \forall x_i(x_i = x_i) \in \Phi_j$.

Remark 7.4.21. Instead of the almost sure valuation of a first-order sentence, one might also consider the *asymptotic expected valuation* $E_p[\pi[\![\psi]\!]] \coloneqq \lim_{n\to\infty} \sum_{j\in\mathbb{N}} j \cdot \mu_{n,p}[\pi[\![\psi]\!] = j]$, and analogously for $L^{\omega}_{\infty\omega}$ and $\mathbb{N}^\infty$. However, due to the possibility of extremely large values of particular events with very low probability, we lose the correspondence to Boolean semantics. As an example consider the sentence $\Delta \coloneqq \forall^{\neq} x\, \forall^{\neq} y\, \forall^{\neq} z(Exy \wedge Eyz \wedge Ezx)$ on random graphs, saying that any three distinct nodes form a triangle. Clearly this sentence is almost surely false on finite graphs since it only evaluates to a positive value on cliques. However, for any probability distribution p with $p[x \geq 2] = q > 0$, we have that with probability $q^{n(n-1)/2}$ a random $\mathbb{N}$-valued graph with vertex set $[n]$ is a clique where each edge has a value ≥ 2. On such a clique, the value of Δ is at least $8^{n(n-1)(n-2)}$. Hence, although Δ evaluates to 0 on all non-cliques, we have that $E_p[\pi[\![\Delta]\!]] \geq \lim_{n\to\infty} q^{n(n-1)/2} \cdot 8^{n(n-1)(n-2)} = \infty$.

7.5 The Random Countable K-Interpretation

A classical fact about 0-1 laws in Boolean semantics is the ω-categoricity of the theory T of extension axioms, and thus the existence of a unique countable τ-structure that satisfies all of them. For finite semirings, this fact extends in a straightforward way to our setting.

Theorem 7.5.1. *For every finite semiring K and every finite relational vocabulary τ there exists a countable K-interpretation $\pi_{\mathcal{R}} \colon \mathrm{Lit}_A(\tau) \to K$ that has the k-extension property for all natural numbers k. Moreover $\pi_{\mathcal{R}}$ is unique up to isomorphism.*

Proof. For any fixed finite semiring K, we can represent K-interpretations $\pi \colon \mathrm{Lit}_A(\tau) \to K$ as classical relational structures $\mathfrak{A}^\pi$ (not to be confused with the induced τ-structure $\mathfrak{A}_\pi$) with universe A over a vocabulary τ^K consisting of relations R^+_j and R^-_j, for $R \in \tau$ and $j \in K$, where $R^+_j = \{\boldsymbol{a} \mid \pi(R\boldsymbol{a}) = j\}$ and $R^-_j = \{\boldsymbol{a} \mid \pi(\neg R\boldsymbol{a}) = j\}$. It is then not difficult to axiomatise the extension properties of K-interpretations in $\mathrm{FO}(\tau^K)$:

- There is a sentence $\mathsf{Int}_K \in \mathrm{FO}(\tau^K)$ such that $\mathfrak{A} \models \mathsf{Int}_K$ if, and only if, $\mathfrak{A} \cong \mathfrak{A}^\pi$ for a K-interpretation π (we have to assert that the relations R_j^+ and R_j^- induce a well-defined and model-defining mapping π).
- For any atomic k-type ρ there is a formula $\mathsf{type}_\rho(\boldsymbol{x})$ such that $\mathfrak{A}^\pi \models \mathsf{type}_\rho(\boldsymbol{a})$ if, and only if, $\rho^\pi_{\boldsymbol{a}} = \rho$ (by asserting membership in R_j^+ and R_j^- according to ρ).
- Hence π has the k-extension property if, and only if,

$$\mathfrak{A}^\pi \models \forall \boldsymbol{x}(\mathsf{type}_\rho(\boldsymbol{x}) \rightarrow \exists y\ \mathsf{type}_{\rho^+}(\boldsymbol{x}, y))$$

for every $i \le k$, every atomic i-type ρ and every extension $\rho^+ \in \mathrm{ext}(\rho)$.

Let now $T(K)$ be the collection of all these extension axioms together with the sentence Int_K. Obviously, every finite subset of $T(K)$ is satisfiable, so by compactness and the Löwenheim-Skolem Theorem, there exists a countable model $\mathcal{R} \models T(K)$. It follows that there exists a countable K-interpretation $\pi_\mathcal{R} \colon \mathrm{Lit}_A(\tau) \to K$, the one represented by $\mathcal{R}$, which has the k-extension property for all natural numbers k.

Further, the standard back-and-forth argument shows that any two such interpretations must be isomorphic. Specifically suppose that $\pi_A \colon \mathrm{Lit}_A(\tau) \to K$ and $\pi_B \colon \mathrm{Lit}_B(\tau) \to K$, with countable universes A and B, both have the k-extension property for all $k < \omega$. Fix enumerations of the universes A and B, and construct partial isomorphisms $p_n = \{(a_1, b_1), \dots, (a_n, b_n)\} \subseteq A \times B$ by induction as follows. Let $p_0 = \emptyset$. If p_n is already defined, let ρ_n be the n-type realised by $\boldsymbol{a} = (a_1, \dots, a_n)$ in π_A, and also by $\boldsymbol{b} = (b_1, \dots, b_n)$ in π_B (given that p_n is a partial isomorphism). For even n, let c be the first element in the enumeration of A that does not appear in p_n, and let $\rho^+ \in \mathrm{ext}(\rho_n)$ be the type realised by $(\boldsymbol{a}, c)$ in π_A. Since π_B has the $(n+1)$-extension property it follows that there exist some $d \in B$ such that also $(\boldsymbol{b}, d)$ realises ρ^+ in π_B. Select the smallest such d in the enumeration of B and set $p_{n+1} = p_n \cup \{(c, d)\}$. For odd n, we proceed analogously, starting with the first d in the enumeration of B that does not occur in p_n. In this way we get an increasing sequence $(p_n)_{n<\omega}$ of partial isomorphisms whose union $p := \bigcup_{n<\omega} p_n$ covers all elements of A and B and thus is an isomorphism between π_A and π_B. □

For finite semirings K in which infinite sums and products are well-defined, the countable random K-interpretation provides evaluations $\pi_\mathcal{R}[\![\psi]\!]$ for arbitrary first-order and infinitary sentences. In particular, this is the case for finite lattice semirings. Notice that Theorem 7.2.9 does not depend on the universe being finite, which implies that for any finite lattice semiring and every sentence $\psi \in L^\omega_{\infty\omega}(\tau)$, we have that $\pi_\mathcal{R}[\![\psi]\!] = f_\psi$. But this coincides with the almost sure valuation of ψ on random finite K-interpretations.

Corollary 7.5.2. *Let K be a finite lattice semiring with a probability distribution $p \colon K^+ \to (0,1]$, and let τ be a finite relational vocabulary. For every sentence $\psi \in L^\omega_{\infty\omega}(\tau) \cup \mathrm{LFP}(\tau)$, the valuation of ψ by the random countable K-interpretation $\pi_\mathcal{R}$ coincides with the almost sure valuation of ψ by finite k-interpretations: $\pi_\mathcal{R}[\![\psi]\!] = \mathrm{ASV}_{K,p}(\psi)$.*

7.6 Conclusion and Future Work

We have seen that the most fundamental result about logic on random structures, the 0-1 law for first-order logic, can be extended to semiring semantics for first-order logic, and also to finite-variable infinitary and fixed-point logic. The specific results and details of our proofs depend on the underlying semiring, but generally follow the same pattern. The first step is the observation that the cornerstone of classical 0-1 laws, the *extension axioms*, generalise to *extension properties* of random semiring interpretations. A new ingredient is the algebraic representation of logical formulae by *polynomials* f_ψ (or similar algebraic expressions) which, due to the extension properties, require only a constant supply of indeterminates that is independent of the size of the universe. This representation facilitates an inductive proof which, at its core, is a *quantifier elimination* argument: we replace quantification by a sum or product over finitely many selector functions, which assign either the value 0, or the minimal or maximal positive semiring value.

This not only proves the 0-1 law, saying that the asymptotic probabilities of statements $\pi[\![\psi]\!] = j$ converge to 0 or 1, but also tells us which semiring values j can actually appear as almost sure valuations. In finite or infinite lattice semirings, these are just three values: 0, 1, and the smallest possible positive value $j > 0$ (if it exists). The complexity of distinguishing these values for a given first-order sentence remains a Pspace-complete problem, as in classical semantics with only two truth values. In the natural semiring, the fact that we interpret true equalities as 1 permits trivial constructions showing that every number $j \in \mathbb{N}$ can occur as almost sure valuations.

The results presented here are a first fundamental step towards understanding the power of semiring semantics for random interpretations. Next steps may include generalisations of our results in several directions; we outline some possibilities that appear promising for future work.

Uniform proof techniques. While our results provide a good understanding of 0-1 laws in lattice semirings (and partial results also for absorptive semirings), our analysis of the natural semirings shows that both the extension properties and the algebraic representations of formulae have to be adapted to the semiring. This raises the question whether there is a more uniform argument that works for a large class of semirings, perhaps for all finite semirings, or if the 0-1 law actually fails for certain semirings. A related question is whether the countable random interpretation also exists for infinite semirings, since our reduction to classical structures is only applicable to finite semirings.

More general probability distributions. We have considered here the case of random interpretations induced by a fixed probability distribution p on the semiring, corresponding to the Erdős-Rényi $G_{n,p}$-model of random graph theory (where p is a constant) and to the classical 0-1 law of Glebskii et al. and Fagin. The study of logic (in classical semantics) on random structures is much richer and has, in particular, investigated models where the probabilities are given as a function of n, the size of the universe, often involving sparse structures (see, e.g., [Spe93]). This leads to more involved calculations of probabilities and requires more sophisticated mathematical machinery, so a generalisation of this work to semiring semantics poses an interesting challenge.

More expressive logics. A different direction for generalisations is to consider stronger logics. We have already seen that the techniques we developed to prove the 0-1 law of first-order logic in semiring semantics (cf. [GHNW22a]) readily extend to infinitary logic $L^{\omega}_{\infty\omega}$, and thus also to LFP over absorptive, fully-continuous semirings. Recently, Badia, Caicedo, and Noguera [BCN23] have extended our results for first-order logic over lattice semirings also to the more general setting of many-valued logics over lattice algebras (but using quite different arguments).

It is known that finite-variable infinitary logic $L^{\omega}_{\infty\omega}$ marks the boundary of 0-1 laws, as it is easy to find counterexamples in $L_{\omega_1\omega}$ with infinitely many variables. Classical model theory has further studied 0-1 laws for existential second order, leading to a classification of the quantifier prefix classes for which a 0-1 law holds (see [KV00] for a survey). The proofs of these 0-1 laws use compactness arguments, which prevent a straightforward generalisation to semiring semantics. However, Lacoste has also found elegant finitistic proofs [Lac96, Lac97], and it seems likely that these proofs can be extended to semiring semantics.

Applications. While classical 0-1 laws give a simple high-level argument for the inexpressibility of Boolean properties for which the 0-1 law fails (such as even cardinality), our results may pave the way towards inexpressibility results for numerical parameters in semiring semantics by showing that their probabilistic behaviour is different from those of logical sentences, for instance in the natural semiring. A further interesting aspect is the study of *certain answers* for queries over incompletely specified databases. Libkin [Lib18] proposes a probabilistic approach that measures how close an answer is to certainty, based on the observation that for the standard model of missing data, the classical 0-1 law holds. Semiring semantics, for instance via its connection to bag semantics, confidence scores and cost analysis, provides an interesting possibility to extend such approaches to a more general setting.

List of Figures

Bibliography

[AN01] A. Arnold and D. Niwiński. *Rudiments of μ-Calculus*. Studies in Logic and the Foundations of Mathematics. Elsevier, 2001. `doi:10.1016/S0049-237X(01)X8001-4`. (Cited on p. 104.)

[ANP21] A. Arnold, D. Niwiński, and P. Parys. A quasi-polynomial black-box algorithm for fixed point evaluation. In *29th EACSL Annual Conference on Computer Science Logic (CSL 2021)*, volume 183 of *Leibniz International Proceedings in Informatics*, pages 9:1–9:23. Dagstuhl, 2021. `doi:10.4230/LIPIcs.CSL.2021.9`. (Cited on pp. 14, 105, 106, and 167.)

[BBPT22] C. Bourgaux, P. Bourhis, L. Peterfreund, and M. Thomazo. Revisiting semiring provenance for datalog. In *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning (KR 2022)*, pages 91–101. IJCAI Organization, 2022. `doi:10.24963/kr.2022/10`. (Cited on p. 125.)

[BCN23] G. Badia, X. Caicedo, and C. Noguera. Asymptotic truth-value laws in many-valued logics, 2023. `arXiv:2306.13904`. (Cited on pp. 178 and 217.)

[BCTV08] P. Buneman, J. Cheney, W. C. Tan, and S. Vansummeren. Curated databases. In *PODS'08: Proceedings of the twenty-seventh ACM SIGMOD-SIGACT-SIGART symposium on principles of database systems*, pages 1–12. ACM, 2008. `doi:10.1145/1376916.1376918`. (Cited on p. 29.)

[BDG90] J. L. Balcázar, J. Díaz, and J. Gabarró. *Structural Complexity II*. EATCS Monographs on Theoretical Computer Science. Springer, 1990. `doi:10.1007/978-3-642-75357-2`. (Cited on p. 196.)

[BG03] A. Blass and Y. Gurevich. Strong extension axioms and Shelah's zero-one law for choiceless polynomial time. *Journal of Symbolic Logic*, 68(1):65–131, 2003. `doi:10.2178/jsl/1045861507`. (Cited on p. 208.)

[BGM24] S. Brinke, E. Grädel, and L. Mrkonjić. Ehrenfeucht-fraïssé games in semiring semantics. In *32nd EACSL Annual Conference on Computer Science Logic (CSL 2024)*, volume 288 of *Leibniz International Proceedings in Informatics*, pages 19:1–19:22. Dagstuhl, 2024. `doi:10.4230/LIPICS.CSL.2024.19`. (Cited on pp. 181, 183, and 187.)

[BGMN24] S. Brinke, E. Grädel, L. Mrkonjić, and M. Naaf. Semiring provenance in the infinite. In *The Provenance of Elegance in Computation - Essays Dedicated to*

Val Tannen, volume 119 of *Open Access Series in Informatics*, pages 3:1–3:26. Dagstuhl, 2024. doi:10.4230/OASIcs.Tannen.3. (Cited on pp. 11, 38, 40, 44, 45, 50, 109, 183, and 187.)

[BGN23a] C. Bizière, E. Grädel, and M. Naaf. Locality theorems in semiring semantics. In *48th International Symposium on Mathematical Foundations of Computer Science (MFCS 2023)*, volume 272 of *Leibniz International Proceedings in Informatics*, pages 20:1–20:15. Dagstuhl, 2023. doi:10.4230/LIPICS.MFCS.2023.20. (Cited on pp. 9, 17, 182, and 186.)

[BGN23b] C. Bizière, E. Grädel, and M. Naaf. Locality theorems in semiring semantics, 2023. Full version of [BGN23a]. arXiv:2303.12627. (Cited on p. 17.)

[BHK+23] T. Barlag, M. Hannula, J. Kontinen, N. Pardal, and J. Virtema. Unified foundations of team semantics via semirings. In *Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning (KR 2023)*, pages 75–85. IJCAI Organization, 2023. doi:10.24963/KR.2023/8. (Cited on pp. 128 and 187.)

[Bir67] G. Birkhoff. *Lattice Theory.* AMS, 1967. (Cited on p. 21.)

[BKLM23] L. E. Bertossi, B. Kimelfeld, E. Livshits, and M. Monet. The shapley value in database management. *SIGMOD Record*, 52(2):6–17, 2023. doi:10.1145/3615952.3615954. (Cited on p. 126.)

[BKT01] P. Buneman, S. Khanna, and W. C. Tan. Why and where: A characterization of data provenance. In *Database Theory - ICDT 2001*, volume 1973 of *Lecture Notes in Computer Science*, pages 316–330. Springer, 2001. doi:10.1007/3-540-44503-X_20. (Cited on p. 125.)

[Blu24] A. Blumensath. *Logic, Algebra and Geometry.* 2024. Last accessed on November 11, 2024. URL: https://www.fi.muni.cz/~blumens. (Cited on pp. 19 and 20.)

[BMR97] S. Bistarelli, U. Montanari, and F. Rossi. Semiring-based constraint satisfaction and optimization. *Journal of the ACM*, 44(2):201–236, 1997. doi:10.1145/256303.256306. (Cited on pp. 3 and 32.)

[BOPP20] C. Bourgaux, A. Ozaki, R. Peñaloza, and L. Predoiu. Provenance for the description logic elhr. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI 2020)*, pages 1862–1869. IJCAI Organization, 2020. doi:10.24963/IJCAI.2020/258. (Cited on p. 128.)

[BS81] S. Burris and H. P. Sankappanavar. *A Course in Universal Algebra*, volume 78 of *Graduate Texts in Mathematics.* Springer, 1981. (Cited on p. 29.)

[BS07] J. Bradfield and C. Stirling. Modal mu-calculi. In *Handbook of Modal Logic*, volume 3 of *Studies in logic and practical reasoning*, pages 721–756. North-Holland, 2007. doi:10.1016/S1570-2464(07)80015-2. (Cited on p. 120.)

[BW18] J. Bradfield and I. Walukiewicz. The mu-calculus and model checking. In *Handbook of Model Checking*, pages 871–919. Springer, 2018. doi:10.1007/978-3-319-10575-8_26. (Cited on p. 146.)

[CCT09] J. Cheney, L. Chiticariu, and W. C. Tan. Provenance in databases: Why, how, and where. *Foundations and Trends in Databases*, 1(4):379–474, 2009. doi:10.1561/1900000006. (Cited on p. 125.)

[CDF+19] W. Czerwinski, L. Daviaud, N. Fijalkow, M. Jurdzinski, R. Lazic, and P. Parys. Universal trees grow inside separating automata: Quasi-polynomial lower bounds for parity games. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2019)*, pages 2333–2349. SIAM, 2019. doi:10.1137/1.9781611975482.142. (Cited on p. 167.)

[CGLP15] F. Canavoi, E. Grädel, S. Lessenich, and W. Pakusa. Defining winning strategies in fixed-point logic. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2015)*, pages 366–377. IEEE, 2015. doi:10.1109/LICS.2015.42. (Cited on pp. 2, 8, 154, 167, and 168.)

[CJ09] A. Chapman and H. V. Jagadish. Why not? In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD 2009)*, pages 523–534. ACM, 2009. doi:10.1145/1559845.1559901. (Cited on p. 111.)

[CJK+17] C. S. Calude, S. Jain, B. Khoussainov, W. Li, and F. Stephan. Deciding parity games in quasipolynomial time. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2017)*, pages 252–263. ACM, 2017. doi:10.1145/3055399.3055409. (Cited on p. 167.)

[Coh00] E. Cohen. Separation and reduction. In *Mathematics of Program Construction*, volume 1837 of *Lecture Notes in Computer Science*, pages 45–59. Springer, 2000. doi:10.1007/10722010_4. (Cited on p. 36.)

[Com89] K. J. Compton. 0-1 laws in logic and combinatorics. In *Algorithms and Order*, NATO Science Series C, pages 353–383. Springer, 1989. doi:10.1007/978-94-009-2639-4_10. (Cited on pp. 176 and 177.)

[DAA13] C. V. Damásio, A. Analyti, and G. Antoniou. Justifications for logic programming. In *Logic Programming and Nonmonotonic Reasoning (LPNMR 2013)*, volume 8148 of *Lecture Notes in Computer Science*, pages 530–542. Springer, 2013. doi:10.1007/978-3-642-40564-8_53. (Cited on p. 126.)

[DG19] K. M. Dannert and E. Grädel. Provenance analysis: A perspective for description logics? In *Description Logic, Theory Combination, and All That*, volume 11560 of *Lecture Notes in Computer Science*, pages 266–285. Springer, 2019. doi:10.1007/978-3-030-22102-7_12. (Cited on p. 128.)

[DG21] K. M. Dannert and E. Grädel. Semiring provenance for guarded logics. In *Hajnal Andréka and István Németi on Unity of Science: From Computing to Relativity Theory through Algebraic Logic*, Outstanding Contributions to Logic, pages 53–79. Springer, 2021. doi:10.1007/978-3-030-64187-0_3. (Cited on pp. 120 and 128.)

[DGMN24] A. Dawar, E. Grädel, L. Mrkonjić, and M. Naaf. Notes on compactness in semiring semantics, 2024. Unpublished notes. (Cited on pp. 17 and 183.)

[DGNT19] K. M. Dannert, E. Grädel, M. Naaf, and V. Tannen. Generalized absorptive polynomials and provenance semantics for fixed-point logic, 2019. Full version of [DGNT21]. `arXiv:1910.07910`. (Cited on pp. 34, 132, and 145.)

[DGNT21] K. M. Dannert, E. Grädel, M. Naaf, and V. Tannen. Semiring provenance for fixed-point logic. In *29th EACSL Annual Conference on Computer Science Logic (CSL 2021)*, volume 183 of *Leibniz International Proceedings in Informatics*, pages 17:1–17:22. Dagstuhl, 2021. `doi:10.4230/LIPIcs.CSL.2021.17`. (Cited on pp. 6, 8, 9, 10, 14, 29, 51, 52, 53, 73, 74, 82, 113, 115, 117, 119, 129, 132, and 134.)

[DJW97] S. Dziembowski, M. Jurdzinski, and I. Walukiewicz. How much memory is needed to win infinite games? In *Proceedings, 12th Annual IEEE Symposium on Logic in Computer Science (LICS 1997)*, pages 99–110. IEEE Computer Society, 1997. `doi:10.1109/LICS.1997.614939`. (Cited on p. 172.)

[DK09] M. Droste and W. Kuich. Semirings and formal power series. In *Handbook of Weighted Automata*, pages 3–28. Springer, 2009. `doi:10.1007/978-3-642-01492-5_1`. (Cited on p. 32.)

[DMRT14] D. Deutch, T. Milo, S. Roy, and V. Tannen. Circuits for datalog provenance. In *Proc. 17th International Conference on Database Theory (ICDT 2014)*, pages 201–212. OpenProceedings.org, 2014. `doi:10.5441/002/icdt.2014.22`. (Cited on pp. 14, 29, 51, 126, 127, and 128.)

[DP02] B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, 2002. `doi:10.1017/CBO9780511809088`. (Cited on p. 21.)

[Dup03] J. Duparc. Positive games and persistent strategies. In *Computer Science Logic (CSL 2003)*, volume 2803 of *Lecture Notes in Computer Science*, pages 183–196. Springer, 2003. `doi:10.1007/978-3-540-45220-1_17`. (Cited on p. 149.)

[EF95] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory: Second Edition*. Perspectives in Mathematical Logic. Springer, 1995. `doi:10.1007/3-540-28788-4`. (Cited on pp. 6, 8, 22, 108, 113, 121, 122, 127, 181, and 182.)

[EKL10] J. Esparza, S. Kiefer, and M. Luttenberger. Newtonian program analysis. *Journal of the ACM*, 57(6):33:1–33:47, 2010. `doi:10.1145/1857914.1857917`. (Cited on pp. 7, 12, 13, 81, 82, 84, 86, 98, 99, and 106.)

[EKL11] J. Esparza, S. Kiefer, and M. Luttenberger. Derivation tree analysis for accelerated fixed-point computation. *Theoretical Computer Science*, 412(28):3226–3241, 2011. `doi:10.1016/J.TCS.2011.03.020`. (Cited on pp. 14, 82, 98, and 126.)

[ELS14] J. Esparza, M. Luttenberger, and M. Schlund. A brief history of strahler numbers. In *Language and Automata Theory and Applications (LATA 2014)*, volume 8370 of *Lecture Notes in Computer Science*, pages 1–13. Springer, 2014. `doi:10.1007/978-3-319-04921-2_1`. (Cited on p. 98.)

[Fag76] R. Fagin. Probabilities on finite models. *Journal of Symbolic Logic*, 41(1):50–58, 1976. `doi:10.1017/S0022481200051756`. (Cited on pp. 16 and 176.)

[FGT08] J. N. Foster, T. J. Green, and V. Tannen. Annotated XML: queries and provenance. In *Proceedings of the Twenty-Seventh ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS 2008)*, pages 271–280. ACM, 2008. `doi:10.1145/1376916.1376954`. (Cited on p. 26.)

[Gai82] H. Gaifman. On local and non-local properties. In *Proceedings of the Herbrand Symposium*, volume 107 of *Studies in Logic and the Foundations of Mathematics*, pages 105–135. Elsevier, 1982. `doi:10.1016/S0049-237X(08)71879-2`. (Cited on p. 182.)

[GHNW22a] E. Grädel, H. Helal, M. Naaf, and R. Wilke. Zero-one laws and almost sure valuations of first-order logic in semiring semantics. In *LICS'22: 37th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 41:1–41:12. ACM, 2022. `doi:10.1145/3531130.3533358`. (Cited on pp. 17, 179, 191, 208, and 217.)

[GHNW22b] E. Grädel, H. Helal, M. Naaf, and R. Wilke. Zero-one laws and almost sure valuations of first-order logic in semiring semantics, 2022. Full version of [GHNW22a]. `arXiv:2203.03425`. (Cited on pp. 17 and 179.)

[GKL+07] E. Grädel, P. G. Kolaitis, L. Libkin, M. Marx, J. Spencer, M. Y. Vardi, Y. Venema, and S. Weinstein. *Finite Model Theory and Its Applications*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2007. `doi:10.1007/3-540-68804-8`. (Cited on pp. 6, 22, 113, 120, and 121.)

[GKLT69] Y. V. Glebskii, D. I. Kogan, M. I. Liogon'kii, and V. A. Talanov. Range and degree of realizability of formulas in the restricted predicate calculus. *Cybernetics*, 5:142–154, 1969. Translated from Kibernetika, Vol. 5, No. 2, pp. 17–27, 1969. `doi:10.1007/BF01071084`. (Cited on pp. 16 and 176.)

[GKT07] T. J. Green, G. Karvounarakis, and V. Tannen. Provenance semirings. In *Proceedings of the Twenty-Sixth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS 2007)*, pages 31–40. ACM, 2007. `doi:10.1145/1265530.1265535`. (Cited on pp. 3, 4, 7, 20, 26, 29, 32, 107, 113, and 125.)

[Gla21] B. Glavic. Data provenance. *Foundations and Trends in Databases*, 9(3-4):209–441, 2021. `doi:10.1561/1900000068`. (Cited on pp. 3 and 125.)

[GLN21] E. Grädel, N. Lücking, and M. Naaf. Semiring provenance for Büchi games: Strategy analysis with absorptive polynomials. In *Proceedings 12th International Symposium on Games, Automata, Logics, and Formal Verification (GandALF 2021)*, volume 346 of *Electronic Proceedings in Theoretical Computer Science*, pages 67–82, 2021. `doi:10.4204/EPTCS.346.5`. (Cited on pp. 16 and 147.)

[GLN24] E. Grädel, N. Lücking, and M. Naaf. Semiring provenance for Büchi games: Strategy analysis with absorptive polynomials. *Logical Methods In Computer*

Science, 20(1), 2024. doi:10.46298/LMCS-20(1:21)2024. (Cited on pp. 16 and 147.)

[GM08] M. Gondran and M. Minoux. *Graphs, dioids and semirings: new models and algorithms*, volume 41 of *Operations Research/Computer Science Interfaces Series*. Springer, 2008. doi:10.1007/978-0-387-75450-5. (Cited on pp. 12, 32, and 81.)

[GM21] E. Grädel and L. Mrkonjić. Elementary equivalence versus isomorphism in semiring semantics. In *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*, volume 198 of *Leibniz International Proceedings in Informatics*, pages 133:1–133:20. Dagstuhl, 2021. doi:10.4230/LIPICS.ICALP.2021.133. (Cited on pp. 9, 122, and 180.)

[GP10] F. Geerts and A. Poggi. On database query languages for k-relations. *Journal of Applied Logic*, 8(2):173–185, 2010. doi:10.1016/J.JAL.2009.09.001. (Cited on p. 126.)

[Gra83] E. Grandjean. Complexity of the first-order theory of almost all finite structures. *Information and Control*, 57(2):180–204, 1983. doi:10.1016/S0019-9958(83)80043-6. (Cited on p. 197.)

[Grä11] E. Grädel. Back and forth between logic and games. In *Lectures in Game Theory for Computer Scientists*, page 99–145. Cambridge University Press, 2011. doi:10.1017/CBO9780511973468.005. (Cited on pp. 8 and 132.)

[Grä16] E. Grädel. Logic and games. Lecture notes, RWTH Aachen University, 2016. Last accessed on November 11, 2024. URL: https://logic.rwth-aachen.de/files/LS-WS15/script.pdf. (Cited on p. 172.)

[Gre11] T. J. Green. Containment of conjunctive queries on annotated relations. *Theory of Computing Systems*, 49(2):429–459, 2011. doi:10.1007/s00224-011-9327-6. (Cited on p. 28.)

[GT] E. Grädel and V. Tannen. Provenance analysis and semiring semantics for first-order logic. In *Festschrift for Johann A. Makowsky*, Trends in Mathematics. Springer. Extended version of [GT17a], submitted for publication. (Cited on pp. 4, 8, 107, 108, and 111.)

[GT17a] E. Grädel and V. Tannen. Semiring provenance for first-order model checking, 2017. arXiv:1712.01980. (Cited on pp. 4, 8, 107, 108, 109, 111, 112, 129, 131, 162, 165, and 190.)

[GT17b] T. J. Green and V. Tannen. The semiring framework for database provenance. In *Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS 2017)*, pages 93–99. ACM, 2017. doi:10.1145/3034786.3056125. (Cited on pp. 3, 27, 29, and 125.)

[GT20] E. Grädel and V. Tannen. Provenance analysis for logic and games. *Moscow Journal of Combinatorics and Number Theory*, 9(3):203–228, 2020. doi:10.

2140/moscow.2020.9.203. (Cited on pp. 6, 8, 9, 51, 129, 131, 147, 148, 149, 165, and 166.)

[GTW02] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *Lecture Notes in Compter Science*. Springer, 2002. doi:10.1007/3-540-36387-4. (Cited on p. 173.)

[HG09] P. Halmos and S. Givant. *Introduction to Boolean algebras*. Undergraduate Texts in Mathematics. Springer, 2009. doi:10.1007/978-0-387-68436-9. (Cited on p. 43.)

[HK99] M. W. Hopkins and D. C. Kozen. Parikh's theorem in commutative Kleene algebra. In *Proceedings. 14th Symposium on Logic in Computer Science (LICS 1999)*, pages 394–401. IEEE, 1999. doi:10.1109/LICS.1999.782634. (Cited on pp. 12, 81, 82, 99, 100, and 103.)

[HS21] D. Hausmann and L. Schröder. Quasipolynomial computation of nested fixpoints. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2021)*, volume 12651 of *Lecture Notes in Computer Science*, pages 38–56. Springer, 2021. doi:10.1007/978-3-030-72016-2_3. (Cited on pp. 14, 105, 106, and 167.)

[KLZ13] S. Köhler, B. Ludäscher, and D. Zinn. First-order provenance games. In *In Search of Elegance in the Theory and Practice of Computation - Essays Dedicated to Peter Buneman*, volume 8000 of *Lecture Notes in Computer Science*, pages 382–399. Springer, 2013. doi:10.1007/978-3-642-41660-6_20. (Cited on pp. 126 and 129.)

[KNP+24] M. A. Khamis, H. Q. Ngo, R. Pichler, D. Suciu, and Y. R. Wang. Convergence of datalog over (pre-) semirings. *Journal of the ACM*, 71(2):8:1–8:55, 2024. doi:10.1145/3643027. (Cited on p. 126.)

[Koz94] D. Kozen. A completeness theorem for Kleene algebras and the algebra of regular events. *Information and Computation*, 110(2):366–390, 1994. doi:10.1006/inco.1994.1037. (Cited on pp. 35 and 36.)

[Kre04] S. Kreutzer. Expressive equivalence of least and inflationary fixed-point logic. *Annals of Pure and Applied Logic*, 130(1):61–78, 2004. doi:10.1016/J.APAL.2004.02.001. (Cited on p. 187.)

[KS85] M. Kaufmann and S. Shelah. On random models of finite power and monadic logic. *Discrete Mathematics*, 54(3):285–293, 1985. doi:10.1016/0012-365X(85)90112-8. (Cited on p. 177.)

[Kui97] W. Kuich. Semirings and formal power series: Their relevance to formal languages and automata. In *Handbook of Formal Languages, Volume 1: Word, Language, Grammar*, pages 609–677. Springer, 1997. doi:10.1007/978-3-642-59136-5_9. (Cited on p. 7.)

[KV92] P. G. Kolaitis and M. Vardi. Infinitary logics and 0-1 laws. *Information and Computation*, 98:258–294, 1992. doi:10.1016/0890-5401(92)90021-7. (Cited on p. 177.)

[KV00] P. G. Kolaitis and M. Y. Vardi. 0-1 laws for fragments of existential second-order logic: A survey. In *Mathematical Foundations of Computer Science (MFCS 2000)*, volume 1893 of *Lecture Notes in Computer Science*, pages 84–98. Springer, 2000. doi:10.1007/3-540-44612-5_6. (Cited on pp. 177 and 217.)

[Lac96] T. Lacoste. Finitistic proofs of 0-1 laws for fragments of second-order logic. *Information Processing Letters*, 58(1):1–4, 1996. doi:10.1016/0020-0190(96)00035-X. (Cited on p. 217.)

[Lac97] T. Lacoste. A simplified proof of the 0-1 law for existential second-order ackermann sentences. *Mathematical Logic Quarterly*, 43:413–418, 1997. doi:10.1002/malq.19970430314. (Cited on p. 217.)

[LBKS20] E. Livshits, L. E. Bertossi, B. Kimelfeld, and M. Sebag. The shapley value of tuples in query answering. In *23rd International Conference on Database Theory (ICDT 2020)*, volume 155 of *Leibniz International Proceedings in Informatics*, pages 20:1–20:19. Dagstuhl, 2020. doi:10.4230/LIPICS.ICDT.2020.20. (Cited on p. 126.)

[Lib18] L. Libkin. Certain answers meet zero-one laws. In *Proceedings of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS 2018)*, pages 195–207. ACM, 2018. doi:10.1145/3196959.3196983. (Cited on pp. 178 and 217.)

[LLG19] S. Lee, B. Ludäscher, and B. Glavic. PUG: a framework and practical implementation for why and why-not provenance. *The VLDB Journal*, 28(1):47–71, 2019. doi:10.1007/S00778-018-0518-5. (Cited on p. 126.)

[LV17] L. Libkin and M. Vardi. 2017 ACM PODS Alberto O. Mendelzon test-of-time award. In *PODS'17: Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, page 91. ACM, 2017. doi:10.1145/3034786.3056122. (Cited on p. 3.)

[Mar76] G. Markowsky. Chain-complete posets and directed sets with applications. *Algebra universalis*, 6:53–68, 1976. doi:10.1007/BF02485815. (Cited on pp. 31 and 32.)

[Mar79] G. Markowsky. Free completely distributive lattices. *Proceedings of the American Mathematical Society*, 74(2):227–228, 1979. doi:10.1090/S0002-9939-1979-0524290-9. (Cited on p. 68.)

[Moh02] M. Mohri. Semiring frameworks and algorithms for shortest-distance problems. *Journal of Automata, Languages and Combinatorics*, 7(3):321–350, 2002. doi:10.25596/jalc-2002-321. (Cited on p. 25.)

[Mrk20] L. Mrkonjić. The model theory of semiring semantics. Master's thesis, RWTH Aachen University, 2020. URL: `https://logic.rwth-aachen.de/pub/mrkonjic/masterarbeit.pdf`. (Cited on p. 52.)

[Mrk24] L. Mrkonjić. *Semiring Semantics: Algebraic Foundations, Model Theory, and Strategy Analysis.* PhD thesis, RWTH Aachen University, 2024. Forthcoming. (Cited on pp. 40, 183, 184, and 186.)

[MT02] J. Marcinkowski and T. Truderung. Optimal complexity bounds for positive LTL games. In *Computer Science Logic (CSL 2002)*, volume 2471 of *Lecture Notes in Computer Science*, pages 262–275. Springer, 2002. `doi:10.1007/3-540-45793-3_18`. (Cited on p. 149.)

[Naa19] M. Naaf. Semirings for provenance analysis of fixed-point logics and games. Master's thesis, RWTH Aachen University, 2019. URL: `http://logic.rwth-aachen.de/pub/naaf/masterarbeit.pdf`. (Cited on pp. 6, 8, 9, 10, 14, 29, 34, 51, 73, 113, 129, and 132.)

[Naa21a] M. Naaf. Computing least and greatest fixed points in absorptive semirings. In *Relational and Algebraic Methods in Computer Science - 19th International Conference (RAMiCS 2021)*, volume 13027 of *Lecture Notes in Computer Science*, pages 344–361. Springer, 2021. `doi:10.1007/978-3-030-88701-8_21`. (Cited on pp. 14 and 83.)

[Naa21b] M. Naaf. Computing least and greatest fixed points in absorptive semirings, 2021. Full version of [Naa21a]. `arXiv:2106.00399`. (Cited on pp. 14 and 83.)

[Sen17] P. Senellart. Provenance and probabilities in relational databases: From theory to practice. *SIGMOD Record*, 46(4):5–15, 2017. `doi:10.1145/3186549.3186551`. (Cited on p. 125.)

[Spe93] J. Spencer. Zero-one laws with variable probability. *Journal of Symbolic Logic*, 58(1):1–14, 1993. `doi:10.2307/2275320`. (Cited on pp. 177 and 216.)

[Wal02] I. Walukiewicz. Monadic second-order logic on tree-like structures. *Theoretical Computer Science*, 275(1):311–346, 2002. `doi:10.1016/S0304-3975(01)00185-2`. (Cited on pp. 2, 154, and 168.)

[WHT03] N. Wallmeier, P. Hütten, and W. Thomas. Symbolic synthesis of finite-state controllers for request-response specifications. In *Implementation and Application of Automata, 8th International Conference (CIAA 2003)*, volume 2759 of *Lecture Notes in Computer Science*, pages 11–22. Springer, 2003. `doi:10.1007/3-540-45089-0_3`. (Cited on p. 173.)

[XZAT18] J. Xu, W. Zhang, A. Alawini, and V. Tannen. Provenance analysis for missing answers and integrity repairs. *IEEE Data Engineering Bulletin*, 41(1):39–50, 2018. URL: `http://sites.computer.org/debull/A18mar/p39.pdf`. (Cited on pp. 111 and 162.)