Provenance Analysis and Semiring Semantics for Logic and Games
SS 2022

(Notes by Matthias Naaf)

Prof. Dr. Erich Grädel

Mathematische Grundlagen der Informatik
RWTH Aachen

Contents

1 Introduction 1
1.1 Motivation 1
1.2 Examples 2
1.3 Towards Semirings 7

2 Commutative Semirings 9
2.1 Absorptive Semirings 11
2.2 Positive Semirings 12
2.3 Naturally-ordered versus +-positive Semirings 13

3 Semiring Valuations 15
3.1 Excursion: Relational Algebra 15
3.2 Valuations and Homomorphisms 18

4 Provenance Semirings 19
4.1 One Semiring to Rule Them All 19
4.2 Other Provenance Semirings 20
4.3 Provenance Semirings as Quotients 23

5 Semiring Valuations of Games 25
5.1 Well-founded Games 26
5.2 Game Valuations in Provenance Semirings 30
5.3 Absorption on Strategies 32
5.4 Game Valuations in Application Semirings 33
5.5 Separating Valuations 36

6 Semiring Valuations for FO (with negation) 39
6.1 Game-theoretic view 40
6.2 Dual-indeterminate Polynomials 42

6.3 Example: Model-defining tracking 45
6.4 Example: Model-compatible tracking 48

7 Elementary Equivalence versus Isomorphism 51
7.1 A Counterexample 52
7.2 Separating Homomorphisms for PosBool$[X]$ 54
7.3 A different Story: The Viterbi Semiring 56

8 Least Fixed Points and ω-Continuous Semirings 61
8.1 Scenarios involving Least Fixed Points 61
8.2 Towards Semirings for Least Fixed Points 65
8.3 ω-continuous Semirings 67

9 Greatest Fixed Points 75
9.1 Semirings for LFP 75
9.2 Generalised Absorptive Polynomials 78
9.3 Case Study: Büchi Games 79

1 Introduction

1.1 Motivation

Idea: Evaluate logical statements not just by true/false, but annotate them by values from some algebraic structure $(S, +, \cdot, 0, 1)$.

Motivation: Get additional information beyond the truth/falsity of a statement

General question: Which combinations of atomic facts are responsible for the truth of a given statement? (This is not limited to logic: consider a computational process applied to a complex input consisting of multiple items.)

Example 1.1. Possible scenarios:

- evaluate a logical sentence on a large (finite) structure,
- compute the result of a database query,
- verify a specification on a transition system,
- determine the winner of a game (with a finite but large game graph). ■

Provenance analysis aims to explain how a particular result depends on the specific input items. Consider a model checking problem $\mathfrak{A} \models^? \psi$ or the evaluation of a database query on a large database. The input items are the atomic facts (of the structure or the database).

- Which facts are actually used in the evaluation?
- Can we derive the result also from different combinations of input items?
- In how many different ways can this output be computed?

Provenance analysis is also interesting for answering refined questions, beyond the truth or falsity of a statement:

Cost: How to minimize the cost of computing the output, based on prizes attached to the atomic facts?

Access restrictions: Suppose that atomic facts come with access restrictions (secret, top secret, confidential, ...). What clearance level is needed to determine that the statement is true?

Confidence: If atomic facts are labelled by some degree of trust (a value in the real interval $[0, 1]$), how to compute a degree of trust to the statement?

An important distinction is between the *joint* use of information, modelled by $\cdot$, and the *alternative* use of information, modelled by $+$. Which assumptions on these operations do we need to answer the questions posed above? In other words, what are appropriate algebraic structures $\mathcal{S} = (S, +, \cdot, 0, 1)$ for provenance analysis?

1.2 Examples

1.2.1 Access Restrictions

Let $\psi(x) = \exists y \exists z (Rxy \land Syz \land Rzx)$ and consider the structure $\mathfrak{A} = (A, R, S)$ with universe $A = \{a, b, c, d\}$ and relations $R = \{(a, b), (a, c), (c, a), (d, a)\}$ and $S = \{(a, a), (b, c), (b, d)\}$. Then $\mathfrak{A} \models \psi(a)$ and $\mathfrak{A} \models \psi(c)$. As a diagram (red for R, blue for S):

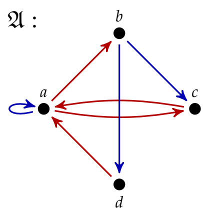

Now consider *access restrictions*: $0 < \mathsf{T} < \mathsf{S} < \mathsf{C} < \mathsf{P} = 1$, where 0 stands for *false* or *inaccessible*, T for *top secret*, S for *secret*, C for *confidential*, and P for *public*. We annotate the relations as follows:

$$R: \left|\begin{array}{cc|c} a & b & \mathsf{P} \\ a & c & \mathsf{T} \\ c & a & \mathsf{S} \\ d & a & \mathsf{P} \end{array}\right| \qquad S: \left|\begin{array}{cc|c} a & a & \mathsf{P} \\ b & c & \mathsf{C} \\ b & d & \mathsf{P} \end{array}\right| \qquad \blacksquare$$

Question: How do we propagate the access restrictions from the atomic facts to the full sentence?

To prove that $\mathfrak{A} \models \psi(a)$ we use

$$\left.\begin{array}{rl} \text{either} & Rab \wedge Sbd \wedge Rda, \\ \text{or} & \underbrace{Rab \wedge Sbc \wedge Rca.}_{\text{joint use of information}} \end{array}\right\} \text{alternative use of information}$$

Joint use: the access level of $\varphi_1 \wedge \varphi_2$ is the *minimum* of the access levels of φ_1 and φ_2. That is, $\pi[\![\varphi_1 \wedge \varphi_2]\!] = \min(\pi[\![\varphi_1]\!], \pi[\![\varphi_2]\!])$. Thus

$$\pi[\![Rab \wedge Sbd \wedge Rda]\!] = \min(\mathsf{P}, \mathsf{P}, \mathsf{P}) = \mathsf{P} \quad \text{(public)},$$
$$\pi[\![Rab \wedge Sbc \wedge Rca]\!] = \min(\mathsf{P}, \mathsf{C}, \mathsf{S}) = \mathsf{S} \quad \text{(secret)}.$$

Similarly, for $\mathfrak{A} \models \psi(c)$ we use $Rca \wedge Saa \wedge Rac$ and obtain

$$\pi[\![\psi(c)]\!] = \min(\mathsf{S}, \mathsf{P}, \mathsf{T}) = \mathsf{T} \quad \text{(top secret)}.$$

Alternative use: the access level of $\varphi_1 \vee \varphi_2$ is the *maximum* of the access levels of φ_1 and φ_2. Hence

$$\pi[\![\varphi_1 \vee \varphi_2]\!] = \max(\pi[\![\varphi_1]\!], \pi[\![\varphi_2]\!]),$$
$$\pi[\![\exists x\ \varphi(x)]\!] = \max_{a \in A} \pi[\![\varphi(a)]\!].$$

Thus

$$\begin{aligned} \pi[\![\psi(a)]\!] &= \max(\pi[\![Rab \wedge Sbd \wedge Rda]\!], \pi[\![Rab \wedge Sbc \wedge Rca]\!]) \\ &= \max(\mathsf{P}, \mathsf{S}) = \mathsf{P}, \end{aligned}$$

$$\begin{aligned} \pi[\![\exists x\ \psi(x)]\!] &= \max(\pi[\![\psi(a)]\!], \pi[\![\psi(b)]\!], \pi[\![\psi(c)]\!], \pi[\![\psi(d)]\!]) \\ &= \max(\quad \mathsf{P} \quad, \quad 0 \quad, \quad \mathsf{T} \quad, \quad 0 \quad) = \mathsf{P}. \end{aligned}$$

Thus, an appropriate structure for reasoning about access levels is

$$\mathbb{A} := (\{0 < \mathsf{T} < \mathsf{S} < \mathsf{C} < \mathsf{P}\}, \underbrace{\max}_{+}, \underbrace{\min}_{\cdot}, \underbrace{0}_{0}, \underbrace{\mathsf{P}}_{1}).$$

1.2.2 Number of Evaluation Strategies

We consider the number of "evaluation strategies" or "proofs" showing that a sentence is true. An appropriate structure here is

$$\mathbb{N} = (\mathbb{N}, +, \cdot, 0, 1).$$

and we label the true atomic facts by 1 (and non-facts by 0). Let $\pi[\![\varphi]\!] := \#_{\mathsf{St}}(\varphi)$ denote the number of strategies for establishing the truth of φ. Then

- $\#_{\mathsf{St}}(\varphi) = 0 \quad$ if φ is false,
- $\#_{\mathsf{St}}(\varphi) = 1 \quad$ if φ is a true atomic fact.

For the alternative use of information,

- $\#_{\mathsf{St}}(\varphi_1 \vee \varphi_2) = \#_{\mathsf{St}}(\varphi_1) + \#_{\mathsf{St}}(\varphi_2)$,
- $\#_{\mathsf{St}}(\exists x\ \varphi(x)) = \sum_{a \in A} \#_{\mathsf{St}}(\varphi(a))$,

since every strategy establishing the truth of *one* alternative can be used to establish the truth of the compound statement.

For the joint use of information,

- $\#_{\mathsf{St}}(\varphi_1 \wedge \varphi_2) = \#_{\mathsf{St}}(\varphi_1) \cdot \#_{\mathsf{St}}(\varphi_2)$,
- $\#_{\mathsf{St}}(\forall x\ \varphi(x)) = \prod_{a \in A} \#_{\mathsf{St}}(\varphi(a))$.

A strategy establishing the truth of a compound statement is *any combination* of strategies establishing the individual pieces of information. In the example:

$$\begin{aligned}\#_{\mathsf{St}}(\psi(a)) &= \#_{\mathsf{St}}(Rab \wedge Sbd \wedge Rda) + \#_{\mathsf{St}}(Rab \wedge Sbc \wedge Rca) \\ &= 1 + 1 = 2;\end{aligned}$$

$$\#_{\mathsf{St}}(\psi(c)) = 1, \quad \text{hence} \quad \#_{\mathsf{St}}(\psi(a) \vee \psi(c)) = 3.$$

A somewhat unintuitive consequence of this approach is that we *lose the idempotence* of $\vee$ and $\wedge$:

$$\#_{\mathrm{St}}(\varphi \vee \varphi) = \#_{\mathrm{St}}(\varphi) + \#_{\mathrm{St}}(\varphi) = 2 \cdot \#_{\mathrm{St}}(\varphi),$$
$$\#_{\mathrm{St}}(\varphi \wedge \varphi) = \#_{\mathrm{St}}(\varphi) \cdot \#_{\mathrm{St}}(\varphi) = \big(\#_{\mathrm{St}}(\varphi)\big)^2.$$

1.2.3 Bag Semantics

A further twist comes from *bag semantics* (or multiset semantics) as used in databases. Here, already the atomic facts may come with a multiplicity, i.e., occur several times in a database.

Consider the following bag relations:

$$R\colon \left|\begin{array}{cc|c} a & b & 1 \\ a & c & 2 \\ c & a & 1 \\ d & a & 3 \end{array}\right| \qquad\qquad S\colon \left|\begin{array}{cc|c} a & a & 2 \\ b & c & 3 \\ b & d & 1 \end{array}\right|$$

Then,

$$\begin{aligned} \pi[\![\psi(a)]\!] &= \pi[\![Rab \wedge Sbd \wedge Rda]\!] + \pi[\![Rab \wedge Sbc \wedge Rca]\!] \\ &= (1 \cdot 1 \cdot 3) + (1 \cdot 3 \cdot 1) = 6, \\ \pi[\![\psi(c)]\!] &= \pi[\![Rca \wedge Saa \wedge Rac]\!] \\ &= 1 \cdot 2 \cdot 2 = 4, \end{aligned}$$

and the query $\psi(x)$ thus results in the multiset $\{\!\{a, a, a, a, a, a, c, c, c, c\}\!\}$.

1.2.4 Cost Analysis

Assume that (true) atomic facts come with a cost value in $\mathbb{R}_+$ (or $\mathbb{R}_+^\infty = \{r \in \mathbb{R} \cup \{\infty\} \mid r \geq 0\}$). Facts that are free have cost 0 (untracked). Non-facts (or inaccessible ones) have cost ∞.

In our running example, we can annotate the relations with cost values:

$$R: \left|\begin{array}{cc|c} a & b & 0 \\ a & c & 0 \\ c & a & 1 \\ d & a & 5 \end{array}\right| \qquad S: \left|\begin{array}{cc|c} a & a & \infty \\ b & c & 2 \\ b & d & 1 \end{array}\right|$$

Then $\pi[\![\varphi]\!] = \text{cost}(\varphi) \in \mathbb{R}_+^\infty$ and we distinguish:

- alternative use of information:

$$\begin{aligned} \text{cost}(\varphi_1 \vee \varphi_2) &= \min(\text{cost}(\varphi_1), \text{cost}(\varphi_2)), \\ \text{cost}(\exists x\ \varphi(x)) &= \min_{a \in A} \text{cost}(\varphi(a)); \end{aligned}$$

- joint use of information:

$$\begin{aligned} \text{cost}(\varphi_1 \wedge \varphi_2) &= \text{cost}(\varphi_1) + \text{cost}(\varphi_2), \\ \text{cost}(\forall x\ \varphi(x)) &= \sum_{a \in A} \text{cost}(\varphi(a)). \end{aligned}$$

Following the example in 1.2.1, we obtain

$$\begin{aligned} \text{cost}(\psi(a)) &= \min(\text{cost}(Rab \wedge Sbd \wedge Rda), \text{cost}(Rab \wedge Sbc \wedge Rca)) \\ &= \min(0+1+5, 0+2+1) = 3 \end{aligned}$$

$$\begin{aligned} \text{cost}(\psi(c)) &= \text{cost}(Rca \wedge Saa \wedge Rac) \\ &= 1 + \infty + 0 = \infty \quad (\psi(c) \text{ cannot be established}) \quad \blacksquare \end{aligned}$$

An appropriate structure for cost analysis is

$$\mathbb{T} = (\mathbb{R}_+^\infty, \underbrace{\min}_{+}, \underbrace{+}_{\cdot}, \underbrace{\infty}_{0}, \underbrace{0}_{1}).$$

This structure is called the *tropical semiring*[1]

[1] In the literature, also other min-plus semirings, for instance over $\mathbb{N}^\infty$ or $\mathbb{R}^\infty$, are called "tropical", and there are fields such as tropical (algebraic) geometry or tropical analysis. The terminology was coined by French mathematicians, such as Jean-Eric Pin, Dominique Perrin, and Christian Choffrut, in honor of their Brazilian colleague Imre Simon, who was one of the pioneers in this area.

1.2.5 Trust

Attach to each atomic fact a level of trust in $[0,1] \subseteq \mathbb{R}$. What confidence do we have that a statement φ is true?

For the alternative use of information:

$$\begin{aligned}\text{trust}(\varphi_1 \vee \varphi_2) &= \max(\text{trust}(\varphi_1), \text{trust}(\varphi_2)),\\ \text{trust}(\exists x\ \varphi(x)) &= \max_{a\in A} \text{trust}(\varphi(a));\end{aligned}$$

For the joint use of information:

$$\begin{aligned}\text{trust}(\varphi_1 \wedge \varphi_2) &= \text{trust}(\varphi_1) \cdot \text{trust}(\varphi_2),\\ \text{trust}(\forall x\ \varphi(x)) &= \prod_{a\in A} \text{trust}(\varphi(a)).\end{aligned}$$

An appropriate structure for reasoning about confidence is thus

$$\mathbb{V} = ([0,1], \max, \cdot, 0, 1), \quad \text{which is called the } \textit{Viterbi semiring}.$$

The Viterbi semiring is isomorphic to the tropical semiring via $\mathbb{T} \to \mathbb{V}$, $x \mapsto e^{-x}$. Indeed, we then have

$$\begin{aligned}\min(x,y) &\mapsto e^{-\min(x,y)} = \max(e^{-x}, e^{-y}),\\ x+y &\mapsto e^{-x-y} = e^{-x} \cdot e^{-y},\\ \infty &\mapsto 0,\\ 0 &\mapsto 1.\end{aligned}$$

1.3 Towards Semirings

What are the general properties that a structure $\mathcal{S} = (S, +, \cdot, 0, 1)$ should have? Recall:

- $+$ models the alternative use of information,
- $\cdot$ models the joint use of information,
- 0 models false,
- all other elements of S represent some degree/shade of true,
- 1 models untracked facts.

We "derive" the following properties of $\mathcal{S}$:

(1) $+$ and $\cdot$ are associative and commutative,
(2) $0 \neq 1$,
(3) 0 is neutral for $+$, and 1 is neutral for $\cdot$,
(4) $0 \cdot a = 0$,
(5) distributive law: $a(b+c) = ab + ac$

(but while + distributes over $\cdot$, the dual property does not always hold: $a + bc \stackrel{?}{=} (a+b)(a+c)$; true for $\mathbb{A}$, but not for $\mathbb{N}, \mathbb{T}, \mathbb{V}$).

Definition 1.2. A structure $\mathcal{S} = (S, +, \cdot, 0, 1)$ satisfying (1)-(5) is called a *commutative semiring*.

Beyond the examples $\mathbb{A}, \mathbb{N}, \mathbb{T}, \mathbb{V}$ there are many others. In particular

$$\mathbb{B} = (\{\bot, \top\}, \vee, \wedge, \bot, \top),$$

the Boolean semiring, which is the natural habitat of mathematical logic (so we *always* did semiring semantics...).

2 Commutative Semirings

Recall that we have defined a *commutative semiring* as a structure $\mathcal{S} = (S, +, \cdot, 0, 1)$ such that

- $(S, +, 0)$ and $(S, \cdot, 1)$ are commutative monoids with $0 \neq 1$,
- $a(b + c) = ab + ac$ for all $a, b, c \in S$,
- $0 \cdot a = 0$ for all $a \in S$.

In the following, "semirings" always are commutative semirings.

Definition 2.1. A semiring $\mathcal{S}$ is called

- $+$-idempotent, or simply idempotent, if $a + a = a$ for all $a \in S$,
- $\cdot$-idempotent, or multiplicatively idempotent, if $a \cdot a = a$ for all $a \in S$,
- fully idempotent if it is both $+$- and $\cdot$-idempotent.

Examples of semirings:

- $\mathbb{B} = (\{\bot, \top, \vee, \wedge, \bot, \top)$ is the *Boolean semiring*.
- Min-max semirings. Let $(A, <)$ be a linear order with minimal element 0 and maximal element 1. Then $(A, \max, \min, 0, 1)$ is a fully idempotent semiring. Notice that it can be finite (as for the access control semiring $\mathbb{A}$ with $0 < \mathsf{T} < \mathsf{S} < \mathsf{C} < \mathsf{P}$) or infinite (as for the *fuzzy semiring* $\mathbb{F} = ([0, 1], \max, \min, 0, 1)$).
- $\mathbb{N} = (\mathbb{N}, +, \cdot, 0, 1)$, the *natural semiring*, used for counting evaluation strategies, or bag semantics in databases.
- $\mathbb{T} = (\mathbb{R}_+^\infty, \min, +, \infty, 0)$ used for min-cost computations. It is called the *tropical semiring* and it is isomorphic to the *Viterbi semiring* $\mathbb{V} = ([0, 1], \max, \cdot, 0, 1)$ used for reasoning about confidence.
- Instead of min-max semirings, induced by a total order $(A, \leq)$, we often consider *lattice semirings* $(A, \sqcup, \sqcap, 0, 1)$ induced by a bounded lattice, i.e., a partial order $(A, \leq)$ where $a \sqcup b$ and $a \sqcap b$ are supremum and infimum of a, b, and $0, 1$ are bottom and top elements.

- The Łukasiewicz semiring is $\mathbb{L} = ([0,1], \max, \odot, 0, 1)$ with $a \odot b = \max(a+b-1, 0)$. An isomorphic variant is $\mathbb{D} = ([0,1], \min, \oplus, 1, 0)$ with $a \oplus b = \min(a+b, 1)$. (It may be an alternative to the Viterbi semiring for reasoning about confidence/doubt).
- A particular example is the semiring $(\mathcal{P}(X), \cup, \cap, \emptyset, X)$ for an arbitrary set X.
- A further interesting semiring over a set X is $\mathrm{Why}(X) := (\mathcal{P}(\mathcal{P}(X)), \cup, \Cup, \emptyset, \{\emptyset\})$ where, for $M, M' \subseteq \mathcal{P}(X)$,

  $$M \Cup M' := \{m \cup m' \mid m \in M, m' \in M'\}.$$

  For instance, take $M = \{m\}$ and $M' = \{m'\}$ for distinct $m, m' \in \mathcal{P}(X)$. Then $M \cup M' = \{m, m'\}$ and $M \Cup M' = \{m \cup m'\}$. Moreover, $(M \cup M') \Cup (M \cup M') = \{m, m \cup m', m'\} \neq M \cup M'$, so $\mathrm{Why}(X)$ is *not* fully idempotent.

You noticed that we did not mention any rings or fields as examples. Indeed, in provenance analysis we are mainly interested in semirings that are *naturally ordered* (by addition): For any semiring $(S, +, \cdot, 0, 1)$, set $a \leq b$ if $a + c = b$ for some $c \in S$. Clearly, $\leq$ is reflexive and transitive:

- $a \leq a$
  (since $a + 0 = a$),
- $a \leq b \ \wedge\ b \leq c \implies a \leq c$
  (since $a + x = b$ and $b + y = c$ imply $a + (x + y) = c$).

However, $\leq$ need in general not be antisymmetric, i.e.

- $a \leq b \ \wedge\ b \leq a \implies a = b$ may fail.

Definition 2.2. A semiring $\mathcal{S}$ is *naturally ordered* if

$$a \leq b \quad \overset{\text{def}}{\iff} \quad \exists c(a + c = b)$$

is a partial order. (This excludes rings.)

A function $f\colon S^k \to S$ on a naturally ordered semiring S is *monotone* if for all $\mathbf{a}, \mathbf{b} \in S^k$ such that $a_i \leq b_i$ (for $i = 1, \ldots, k$) also $f(\mathbf{a}) \leq f(\mathbf{b})$.

Lemma 2.3. $+$ and $\cdot$ are monotone on naturally ordered semirings.

Proof. Let $a \leq b$. We have to prove for all c that $a + c \leq b + c$ and $a \cdot c \leq b \cdot c$. We have $a + x = b$ for some $x \in S$. Hence

$$a + c \leq (a + c) + x = (a + x) + c = b + c$$
$$a \cdot c \leq ac + xc = (a + x)c = b \cdot c$$

Q.E.D.

2.1 Absorptive Semirings

Definition 2.4. A semiring $\mathcal{S} = (S, +, \cdot, 0, 1)$ is *absorptive* if $a + ab = a$ for all $a, b \in S$ (or equivalently $1 + b = 1$ for all $b \in S$).

Notice that an absorptive semiring is also idempotent ($a + a = a + a \cdot 1 = a$) and both properties correspond to classical logical equivalences: $\varphi \vee \varphi \equiv \varphi$ and $\varphi \vee (\varphi \wedge \vartheta) \equiv \varphi$.

Lemma 2.5. Every idempotent (and hence in particular every absorptive) semiring is naturally ordered.

Proof. Assume $a \leq b$ by $a + c = b$, and $b \leq a$ by $b + d = a$. Hence

$$a + b = a + (a + c) = (a + a) + c = a + c = b$$
$$a + b = (b + d) + b = (b + b) + d = b + d = a$$

and thus $a = b$. Hence $\leq$ is antisymmetric and thus a partial order. Q.E.D.

Lemma 2.6. Let $(S, +, \cdot, 0, 1)$ be naturally ordered. The following are equivalent:

(1) S is absorptive
(2) $\cdot$ is decreasing: $ab \leq a$ for all a, b,
(3) 1 is maximal w.r.t. $\leq$.

Proof. For all $a, b \in S$:

- (1) $\Rightarrow$ (3): S absorptive $\implies a + 1 = 1 \implies a \leq 1$.
- (3) $\Rightarrow$ (2): $b \leq 1 \implies ab \leq a$.
- (2) $\Rightarrow$ (1): $a \leq a + ab = a(1 + b) \overset{(2)}{\leq} a$. Hence $a = a + ab$. Q.E.D.

2.2 Positive Semirings

Definition 2.7 (positive)**.** A semiring $(S, +, \cdot, 0, 1)$

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

Definition 2.8 (homomorphism)**.** A *semiring homomorphism* is a function $h \colon S \to T$ on semirings S, T such that

(1) $h(0) = 0$ and $h(1) = 1$,
(2) $h(a +_S b) = h(a) +_T h(b)$,
(3) $h(a \cdot b) = h(a) \cdot h(b)$.

Recall the Boolean semiring $\mathbb{B} = (\{\bot, \top\}, \vee, \wedge, \bot, \top)$. For a semiring $(S, +, \cdot, 0, 1)$, let the *truth-projection* be the function $t_S \colon S \to \mathbb{B}$ with

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

Then t_S is a semiring homomorphism if, and only if, S is positive.

Proof. Clearly $t_S(0) = \bot$ and $t_S(1) = \top$.

- Addition:

$$t_S(a + b) = \bot \iff a + b = 0$$
$$t_S(a) \vee t_S(b) = \bot \iff t_S(a) = t_S(b) = \bot \iff a = b = 0$$

  Hence t_S preserves $+$ (for all a, b) $\iff$ S is $+$-positive.
- Multiplication:

$$t_S(a \cdot b) = \bot \iff a \cdot b = 0$$
$$\begin{aligned} t_S(a) \wedge t_S(b) = \bot &\iff t_S(a) = \bot \text{ or } t_S(b) = \bot \\ &\iff a = 0 \text{ or } b = 0 \end{aligned}$$

  Hence t_S preserves $\cdot$ (for all a, b) $\iff$ S has no divisors of 0. Q.E.D.

2.3 Naturally-ordered versus +-positive Semirings

Lemma 2.9. Every naturally ordered semiring is +-positive.

Proof. Let S be naturally ordered, and assume that $a + b = 0$. Then $0 \leq a$ (by $0 + a = a$) and $a \leq 0$ (by $a + b = 0$), hence $a = 0$. Q.E.D.

The converse is not true: Let $(S, +, \cdot, 0, 1)$ with $S = \{0, 1, 2\}$ be defined as follows:

$+$	0	1	2
0	0	1	2
1	1	2	1
2	2	1	2

$\cdot$	0	1	2
0	0	0	0
1	0	1	2
2	0	2	2

This is a +-positive semiring that is *not* naturally ordered, since $1 \leq 2$ (by $1 + 1 = 2$) and $2 \leq 1$ (by $2 + 1 = 1$). Q.E.D.

3 Semiring Valuations

Let $\tau = \{R_1, \ldots, R_m\}$ be a finite relational vocabulary, A a finite universe. We denote by $\mathrm{Atoms}_A(\tau)$ the set of fully instantiated atoms $R\mathbf{a}$ with $R \in \tau$, $\mathbf{a} \in A^{\mathrm{arity}(R)}$.

Definition 3.1. An *S-valuation* (for τ and A) is a function of the form $\pi\colon \mathrm{Atoms}_A(\tau) \to S$ into a semiring $(S, +, \cdot, 0, 1)$.

It extends to valuations for more powerful expressions, for instance for positive first-order logic FO^+.

- $\mathrm{FO}^+(\tau)$: closure of τ-atoms $R\mathbf{x}$ and equalities under $\wedge, \vee, \exists, \forall$.
- $\mathrm{FO}^+_A(\tau) := \{(\varphi, \beta) \mid \varphi \in \mathrm{FO}^+(\tau),\ \beta\colon \mathrm{free}(\varphi) \to A\}$
  $= \{\varphi(\mathbf{a}) \mid \varphi(\mathbf{x}) \in \mathrm{FO}^+(\tau),\ \mathbf{a} \in A^n\}$.

Definition 3.2. An S-valuation $\pi\colon \mathrm{Atoms}_A(\tau) \to S$ extends to a function $\pi\colon \mathrm{FO}^+_A(\tau) \to S$ as follows:

$$\begin{aligned}
\pi[\![a = a']\!] &= \begin{cases} 1, & a = a' \\ 0, & \text{otherwise,} \end{cases} \\
\pi[\![\psi \vee \varphi]\!] &= \pi[\![\psi]\!] + \pi[\![\varphi]\!], \\
\pi[\![\psi \wedge \varphi]\!] &= \pi[\![\psi]\!] \cdot \pi[\![\varphi]\!], \\
\pi[\![\exists x\ \varphi]\!] &= \sum_{a \in A} \pi[\![\varphi(a)]\!], \\
\pi[\![\forall x\ \varphi]\!] &= \prod_{a \in A} \pi[\![\varphi(a)]\!].
\end{aligned}$$

3.1 Excursion: Relational Algebra

An alternative positive logical formalism, popular for instance in databases, is RA^+, positive relational algebra, a fragment of full relational algebra RA.

Syntax: Let $\tau = \{R_1, \dots, R_m\}$ be a finite relational vocabulary. $\mathrm{RA}(\tau)$ is a calculus of terms, each of which comes with an arity.

- each $R_i \in \tau$ is in $\mathrm{RA}(\tau)$,
- let $R, R' \in \mathrm{RA}(\tau)$ have arity r, let $T \in \mathrm{RA}(\tau)$ have arity t, and let $i, j, i_1, \dots, i_s \leq r$. Then $\mathrm{RA}(\tau)$ contains the terms

$$\begin{array}{lcccccc} & R - R', & R \cup R', & R \times T, & \pi_{i_1,\dots,i_s} R, & \sigma_{i=j} R, \\ \text{arities:} & r\ , & \multicolumn{4}{c}{\underbrace{r\ ,\quad r+t\ ,\quad s\ ,\quad r\ .}_{\mathrm{RA}^+(\tau)}} \end{array}$$

In addition, we admit the term $\emptyset$ (for any arity).

Semantics: Fix a (usually infinite) domain D. Let $(S, +, \cdot, 0, 1)$ be a semiring. The idea is to interpret each term $R \in \mathrm{RA}^+(\tau)$ by a function $R\colon D^r \to S$ with *finite support* $\mathrm{supp}(R) = \{\mathbf{d} \in D^r \mid R(\mathbf{d}) \neq 0\}$.

Starting from such valuations $R_i\colon D^{r_i} \to S$ for the basic terms $R_i \in \tau$ we define

$$\begin{aligned} (R \cup R')(\mathbf{d}) &:= R(\mathbf{d}) + R'(\mathbf{d}), \\ (R \times T)(\mathbf{d}, \mathbf{d}') &:= R(\mathbf{d}) \cdot T(\mathbf{d}'), \\ (\pi_{i_1,\dots,i_s} R)(\mathbf{d}') &:= \sum_{\mathbf{d}|_{i_1 \cdots i_s} = \mathbf{d}'} R(\mathbf{d}), \end{aligned}$$

where we write $\mathbf{d}|_{i_1,\dots,i_s} := (d_{i_1} \cdots d_{i_s}) \in D^s$, for $\mathbf{d} = (d_1, \dots, d_r) \in D^r$,

$$\begin{aligned} (\sigma_{i=j} R)(\mathbf{d}) &:= \begin{cases} R(\mathbf{d}), & \text{if } \mathbf{d}|_i = \mathbf{d}|_j, \\ 0, & \text{otherwise,} \end{cases} \\ \emptyset(\mathbf{d}) &:= 0. \end{aligned}$$

How to interpret $R - R'$? This is indeed a problem since in general, we do not have an appropriate semiring operation for that. That's why we only consider RA^+ for now. But we will get back to this, and also to negation in FO.

In the Boolean sense, RA is (more or less) equivalent to FO: Every term $R \in \mathrm{RA}(\tau)$ can be translated into a first-order formula $\varphi_R(\mathbf{x})$ such that, for any interpretation of the basic terms $R_1, \ldots, R_m \in \tau$ as *finite* relations $R_i \subseteq D^{r_i}$, we have that

$$R = \{\mathbf{d} \in D^r \mid (D, R_1, \ldots, R_m) \models \varphi(\mathbf{d})\}.$$

The translation is straightforward from the definition of the semantics of RA.

The converse is a little bit more subtle, since RA usually (at least as used in databases) assumes an infinite domain, but defines only finite relations. In particular, RA can only define the difference of relations, but not their (infinite) complement. We can get around this, by restricting $(D, R_1, \ldots, R_m)$ to the *active domain*, the set aD of elements that occur in some relation. It is RA^+-definable by

$$aD := \bigcup_{R \in \tau} \bigcup_{i \leq \mathrm{arity}(R)} \pi_i R.$$

One can then translate every formula $\varphi(x_1, \ldots, x_r) \in \mathrm{FO}$ into an r-ary term $R_\varphi \in \mathrm{RA}$ such that, for all $R_1, \ldots, R_m$ over D,

$$(aD, R_1, \ldots, R_m) \models \varphi(\mathbf{d}) \quad \Longleftrightarrow \quad \mathbf{d} \in R_\varphi$$

for all tuples $\mathbf{d} \in (aD)^r$.

Can we extend this equivalence to semiring semantics? So far we do not have a semiring semantics for full RA and FO, but what about RA^+ and FO^+? Embedding RA^+ in FO^+ poses no problem:

$$\begin{aligned}
\varphi_{R \cup R'}(\mathbf{x}) &:= \varphi_R(\mathbf{x}) \vee \varphi_{R'}(\mathbf{x}) \\
\varphi_{R \times T}(\mathbf{x}, \mathbf{y}) &:= \varphi_R(\mathbf{x}) \wedge \varphi_T(\mathbf{y}) \\
\varphi_{\pi_{i_1, \ldots, i_s}(R)}(\mathbf{x}) &:= \exists \mathbf{y}(\varphi_R(\mathbf{y}) \wedge y_{i_1} = x_1 \wedge \cdots \wedge y_{i_s} = x_s) \\
\varphi_{\sigma_{i=j} R}(\mathbf{x}) &:= \varphi_R(\mathbf{x}) \wedge x_i = x_j \\
\varphi_\emptyset(\mathbf{x}) &= \bot
\end{aligned}$$

For any valuation π that interprets the symbols $R_i \in \tau$ as functions

$R_i \colon D^{r_i} \to S$ (into a semiring S) we have that, for any $R \in \mathrm{RA}^+$ and $\mathbf{d} \in D^r$, $R(\mathbf{d}) = \pi[\![\varphi_R(\mathbf{d})]\!]$. Hence $\mathrm{RA}^+ \leq \mathrm{FO}^+$ in semiring semantics (for any semiring).

The converse is not true! The translation uses only $\vee, \wedge, \exists$ but not $\forall$, so RA^+ is actually embedded into the *existential positive* fragment of FO, denoted $\mathrm{FO}(\exists, \vee, \wedge)$. RA^+ is strictly weaker than FO^+, even in the Boolean sense. Consider $\tau = \{R_1, R_2\}$, both unary, and an instance with $R_1 = \{c\}$, $R_2 = \{c, d\}$ for distinct $c, d \in D$. By induction, one easily proves that for every term $R \in \mathrm{RA}^+(\tau)$ (that is not equivalent to $\emptyset$ on all instances) we have that $(c, \dots, c) \in R$. But let $\varphi(x) := R_2 x \wedge \forall y R_1 y \in \mathrm{FO}^+$. Clearly, $(aD, R_1, R_2) \not\models \varphi(c)$. Q.E.D.

3.2 Valuations and Homomorphisms

Theorem 3.3 (Fundamental Property)**.** Let $h \colon S \to T$ be a semiring homomorphism and $\pi \colon \mathrm{Atoms}_A(\tau) \to S$. Then we get a semiring valuation $h \circ \pi \colon \mathrm{Atoms}_A(\tau) \to T$ such that $h(\pi[\![\varphi]\!]) = (h \circ \pi)[\![\varphi]\!]$ for all $\varphi \in \mathrm{FO}^+(\tau)$.

$$\begin{tikzcd} & \mathrm{Atoms}_A(\tau) \arrow[dl, "\pi"'] \arrow[dr, "h \circ \pi"] & \\ S \arrow[rr, "h"] & & T \end{tikzcd} \quad \Longrightarrow \quad \begin{tikzcd} & \mathrm{FO}^+(\tau) \arrow[dl, "\pi"'] \arrow[dr, "h \circ \pi"] & \\ S \arrow[rr, "h"] & & T \end{tikzcd}$$

Proof. Simple induction. For instance, if $\psi = \exists x\ \varphi$, then

$$\begin{aligned} h(\pi[\![\psi]\!]) &= h\Big(\sum_{a \in A} \pi[\![\varphi(a)]\!]\Big) \\ &\overset{\text{hom}}{=} \sum_{a \in A} h(\pi[\![\varphi(a)]\!]) \\ &\overset{\text{I.H.}}{=} \sum_{a \in A} (h \circ \pi)[\![\varphi(a)]\!] = (h \circ \pi)[\![\psi]\!]. \end{aligned}$$

Q.E.D.

4 Provenance Semirings

4.1 One Semiring to Rule Them All

Let X be a finite set of indeterminates (or "provenance tokens"). The semiring $\mathbb{N}[X]$ of multivariate polynomials (with indeterminates from X and coefficients from $\mathbb{N}$) is the semiring that is freely generated by X. It has the following universal property.

Theorem 4.1 (Universal property)**.** For every semiring S and every map $h\colon X \to S$ there is a unique extension of h to a semiring homomorphism $\hat{h}\colon \mathbb{N}[X] \to S$.

Proof. Starting from $\hat{h}(x) = h(x)$ for $x \in X$, $\hat{h}(0) = 0$ and $\hat{h}(1) = 1$, one constructs by $\hat{h}(f + g) := \hat{h}(f) + \hat{h}(g)$ and $\hat{h}(f \cdot g) := \hat{h}(f) \cdot \hat{h}(g)$ a homomorphism $\hat{h}\colon \mathbb{N}[X] \to S$.

For two homomorphisms $h_1, h_2\colon \mathbb{N}[X] \to S$ that extend h, it follows by induction that $h_1 = h_2$: if $h_1(f) = h_2(f)$ and $h_1(g) = h_2(g)$, then also $h_1(f + g) = h_1(f) + h_1(g) = h_2(f) + h_2(g) = h_2(f + g)$, and analogously for $f \cdot g$. Q.E.D.

We can think of $\hat{h}(3x^2 + 2xy + y^3)$ as evaluating $3x^2 + 2xy + y^3$ in S, based on the given map $h\colon \{x, y\} \to S$.

Fix a set $X = \{X_\alpha \mid \alpha \in \mathrm{Atoms}_A(\tau)\}$ of indeterminates to label the atoms. The map $\ell\colon \mathrm{Atoms}_A(\tau) \to X$, $\alpha \mapsto X_\alpha$ extends to a valuation $\ell\colon \mathrm{FO}^+(\tau) \to \mathbb{N}[X]$ that maps every sentence $\psi(\mathbf{a})$ to a polynomial $f_{\psi(\mathbf{a})} \in \mathbb{N}[X]$. Let now $h\colon X \to S$ induce a valuation of the atoms in a semiring S. By the fundamental property, $\pi[\![\psi]\!] = \hat{h}(f_\psi)$.

$$
\begin{tikzcd}
 & X \arrow[dl, hook'] \arrow[dr, "h"] & \\
\mathbb{N}[X] \arrow[rr, "\hat{h}"] & & S
\end{tikzcd}
\quad \Longrightarrow \quad
\begin{tikzcd}
 & \mathrm{FO}^+(\tau) \arrow[dl, "\ell"'] \arrow[dr, "\pi = \hat{h} \circ \ell"] & \\
\mathbb{N}[X] \arrow[rr, "\hat{h}"] & & S
\end{tikzcd}
$$

To compute valuations of $\psi \in \mathrm{FO}^+(\tau)$ in various semirings, we can thus compute the valuation $f_\psi \in \mathbb{N}[X]$ of ψ in the most general semiring and specialise to valuations in application semirings S via homomorphisms $\hat{h}\colon \mathbb{N}[X] \to S$ induced by $h\colon X \to S$.

4.2 Other Provenance Semirings

From $\mathbb{N}[X]$, the most general semiring over X, we can get "simpler" and "less informative" semirings which have specific algebraic properties:

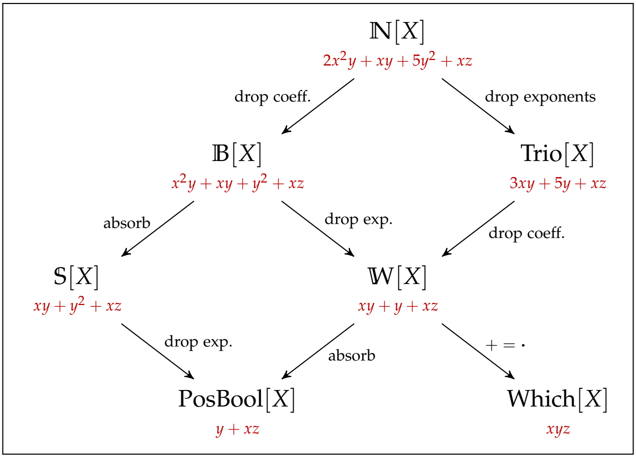

Figure 4.1. Overview on different provenance semirings

$\mathbb{N}[X]$:

Elements $f \in \mathbb{N}[X]$ can be written as sum of monomials $c \cdot x_1^{e_1} \cdots x_r^{e_r}$ with $x_1, \ldots, x_r \in X$ and $c, e_1, \ldots, e_r \in \mathbb{N}$.

$\mathbb{B}[X]$:

Sums of distinct monomials $x_1^{e_1} \cdots x_r^{e_r}$. The semiring $\mathbb{B}[X]$ is $+$-idempotent.

Trio$[X]$:

Sums of $c \cdot \underbrace{x_1 \cdots x_r}_{m \subseteq X}$, where

$$\begin{aligned}(c \cdot m) \bullet (c' \cdot m') &= cc' \cdot (m \cup m'),\\ (\sum_m c_m \cdot m) \bullet (\sum_{m'} c_{m'} \cdot m') &= \sum_{m,m'} c_m c_{m'} \cdot (m \cup m').\end{aligned}$$

$\mathbb{W}[X]$:

Sums/sets of monomials $m \subseteq X$, i.e., sets of subsets of X, where

$$\begin{aligned} + &= \cup, \qquad \bullet = \Cup, \qquad M \bullet M' = \{m \cup m' \mid m \in M, m' \in M'\},\\ 0 &= \emptyset, \qquad 1 = \{\emptyset\}.\end{aligned}$$

Note that $\mathbb{W}[X] \cong \mathrm{Why}(X) = (\mathcal{P}(\mathcal{P}(X)), \cup, \Cup, \emptyset, \{\emptyset\})$.

$\mathbb{S}[X]$:

Set of absorptive polynomials. We consider monomials $m = x_1^{e_1} \cdots x_r^{e_r}$ as a map $m\colon X \to \mathbb{N}$ with $m(x_i) = e_i$ (and $m(x) = 0$ if x does not occur in m). Multiplication $m \bullet m'$ is defined (as usual) by $(m \bullet m')(x) = m(x) + m'(x)$. For $m, m'\colon X \to \mathbb{N}$ we say that m *absorbs* m' ($m \succcurlyeq m'$) if $m(x) \leq m'(x)$ for all $x \in X$ (notice the order inversion).

Example 4.2. The monomial xy^2 absorbs x^3y^2 and xy^5z, but not x^2y. We write 1_m for the monomial with $1_m(x) = 0$ for all $x \in X$ and observe that 1_m absorbs every monomial. ∎

An *absorptive polynomial* $p \in \mathbb{S}[X]$ is a set ($\widehat{=}$ sum) of monomials, none of which absorbs any other one, i.e., an *antichain of monomials* w.r.t. $\succcurlyeq$. Let M be any set of monomials over X; we write $\mathrm{maximals}(M)$ for the set of absorption-maximal $m \in M$ (which is always an antichain, i.e., an element of $\mathbb{S}[X]$). Addition and multiplication of absorptive polynomials is defined as usual, except that we afterwards apply absorption:

$$\begin{aligned} p + q &:= \mathrm{maximals}(p \cup q),\\ p \bullet q &:= \mathrm{maximals}(\{m \bullet m' := m \in p, m' \in q\}),\\ 0 &:= \emptyset,\end{aligned}$$

$$1 := \{1_m\}.$$

Example 4.3. In $\mathbb{S}[X]$ with $X = \{x, y, z\}$:

$$\underbrace{(x^2y + xy^2)}_{p} \cdot \underbrace{(x + yz)}_{q} + z = \underbrace{(x^3y + \cancel{x^2y^2z} + x^2y^2 + xy^3z)}_{p \cdot q} + z$$
$$= x^3y + x^2y^2 + z \qquad \blacksquare$$

<u>$(\mathrm{PosBool}[X], \vee, \wedge, \bot, \top)$:</u>
Positive Boolean formulae with variables in X where we *identify equivalent formulae*. We can also view $\mathrm{PosBool}[X]$ as the set of absorptive polynomials in $\mathbb{S}[X]$ without exponents (i.e., with monomials $m\colon X \to \mathbb{B}$). This corresponds to the representation of positive Boolean formulae in irredundant DNF. For instance,

$$xy + xz \quad \rightsquigarrow \quad (x \wedge y) \vee (x \wedge z).$$

Absorption corresponds to the equivalence $x \vee (x \wedge y) \equiv x$.

Since $\mathrm{PosBool}[X]$ identifies equivalent Boolean expressions, semiring semantics in $\mathrm{PosBool}[X]$ is less dependent on the *syntax* of a formula: $\varphi \vee \varphi$, $\varphi \wedge \varphi$ and φ get the same value in $\mathrm{PosBool}[X]$ although their valuations in $\mathbb{N}[X]$, $\mathbb{W}[X]$ and $\mathbb{S}[X]$ may differ, since $\mathrm{PosBool}[X]$ is fully idempotent:

$$\underset{\vee}{(f + g)} \underset{\wedge}{\cdot} \underset{\vee}{(f + g)} = \underbrace{f \cdot f}_{f} + \underbrace{f \cdot g + f \cdot g}_{fg} + \underbrace{g \cdot g}_{g} \overset{\text{absorb}}{=} f + g$$

and also a *lattice semiring*, with partial order $f \leq g \iff f \models g$.

<u>$\mathrm{Which}(X) = \mathrm{Lin}(X) := (\mathcal{P}(X) \cup \{\bot\}, +, \cdot, \bot, \emptyset)$:</u>
The operations are defined as follows:

$$\begin{aligned} &S + \bot = S, \qquad S \cdot \bot = \bot, && \text{for all } S, \\ &S + S' = S \cdot S' = S \cup S', && \text{for all } S, S' \in \mathcal{P}(X). \end{aligned}$$

4.3 Provenance Semirings as Quotients

Definition 4.4 (congruence)**.** A *congruence* on S is an equivalence relation $\sim$ on S such that

$$a \sim a' \text{ and } b \sim b' \quad \implies \quad a + b \sim a' + b' \text{ and } a \cdot b \sim a' \cdot b'. \ (*)$$

If $\sim$ is a congruence on a semiring S, with $0 \not\sim 1$, then the quotient $S/_\sim$ is also a semiring with elements $[a] := \{b \in S \mid a \sim b\}$ and operations $[a] + [b] := [a + b]$ and $[a] \cdot [b] := [ab]$ (note that the operations are well-defined, i.e., independent of the choices of representatives in $[a]$ and $[b]$).

Notice that the intersection of congruences is again a congruence. In $\mathbb{N}[X]$, $x + x \sim x$ (for $x \in X$) generates a congruence (the smallest equivalence relation satisfying $(*)$ and containing $x + x \sim x$ for all $x \in X$). Then $\mathbb{N}[X]/_\sim \cong \mathbb{B}[X]$ by dropping coefficients.

$\mathbb{B}[X]$ has the universal property for idempotent semirings. Consider $\pi\colon X \to S$. It can be (uniquely) extended to $\hat{\pi}\colon \mathbb{B}[X] \to S$ if, and only if, S is idempotent:

$$\hat{\pi}(f) = \hat{\pi}(\underbrace{f + f}_{f}) = \hat{\pi}(f) + \hat{\pi}(f).$$

Also $\mathbb{S}[X]$ can be defined as a quotient semiring of $\mathbb{N}[X]$ or $\mathbb{B}[X]$. Let $\sim$ be the smallest congruence on $\mathbb{N}[X]$ (or $\mathbb{B}[X]$) such that $f \sim f + fg$ for all $f, g \in \mathbb{N}[X]$ (or $f, g \in \mathbb{B}[X]$). We claim that $\mathbb{S}[X] \cong \mathbb{N}[X]/_\sim$ (analogously, $\mathbb{S}[X] \cong \mathbb{B}[X]/_\sim$).

Proof. For $f \in \mathbb{N}[X]$, let p_f be the set of monomials that occur in f (with any coefficient $c > 0$), i.e., $f = \sum_{m \in p_f} c_m \cdot m$. Then $\text{maximals}(p_f) \in \mathbb{S}[X]$ which we identify with the polynomial $f_0 := \sum\{m \mid m \in \text{maximals}(p_f)\}$. Clearly,

$$\pi\colon \mathbb{N}[X] \to \mathbb{S}[X], \quad f \mapsto f_0 \mathrel{\hat{=}} \text{maximals}(p_f)$$

is surjective.

Notice that $f \approx g \overset{\text{def}}{\iff} f_0 = g_0$ defines a congruence on $\mathbb{N}[X]$ with $f \approx f + f \cdot g$ for all f, g (since $(f + fg)_0 = f_0$). But $\sim$ is the smallest congruence with this property, so $f \sim g \implies f \approx g \implies f_0 = g_0$.

We claim that, viewing f_0 as a polynomial in $\mathbb{N}[X]$, we have $f \sim f_0$. Indeed, for each $m \in p_f$ there is at least one $m_0 \in \text{maximals}(p_f)$ with $m = m_0 \cdot m'$ for some m', so $m_0 \sim m_0 + m$. But this implies that

$$\sum_{m \in p_f} c_m \cdot m = f \ \sim \ f_0 = \sum \{ m_0 \mid m_0 \in \text{maximals}(p_f) \}.$$

It follows that $f_0 = g_0 \implies f \sim g$ (by $f \sim f_0 \sim g_0 \sim g$). Thus $f \sim g \iff f_0 = g_0$, and we obtain:

$$\begin{tikzcd}
\mathbb{N}[X] \arrow[d, two heads] \arrow[dr, "\pi"] & \\
\mathbb{N}[X]/_{\sim} \arrow[r, "\sim"] & \mathbb{S}[X]
\end{tikzcd}$$

Q.E.D.

5 Semiring Valuations of Games

To understand what valuations in provenance semirings such as $\mathbb{N}[X]$, $\mathbb{B}[X]$, $\mathbb{S}[X]$, ... tell us about, say, first-order sentences $\psi(\mathbf{a}) \in \mathrm{FO}^+$ (and later $\psi(\mathbf{a}) \in \mathrm{FO}$ and other logics), it is instructive to consider the associated model-checking games. Given $\psi \in \mathrm{FO}^+(\tau)$ (or $\mathrm{FO}(\tau)$) and a finite structure $\mathfrak{A}$, we know the model-checking game $\mathcal{G}(\mathfrak{A}, \psi)$: positions are instantiated subformulae $\varphi(\mathbf{a})$, Player 0 (Verifier) moves at positions $\varphi_1 \vee \varphi_2$ and $\exists x\ \varphi$ (alternative use of information), Player 1 (Falsifier) moves at positions $\varphi_1 \wedge \varphi_2$ and $\forall x\ \varphi$ (joint use of information), and plays end at atomic formulae (or literals).

Notice that the *game graph* of $\mathcal{G}(\mathfrak{A}, \psi)$ only depends on ψ and the *universe* A of $\mathfrak{A}$. Only the labelling of the terminal positions, as either winning for Player 0 or for Player 1, that depends on which atoms/literals are true in $\mathfrak{A}$. Hence the definition of model-checking games readily generalises to semiring semantics. Given $\psi \in \mathrm{FO}^+(\tau)$ (or $\mathrm{FO}(\tau)$) and a finite universe A, we obtain a game graph $\mathcal{G}(A, \psi)$. The terminal positions of $\mathcal{G}(A, \psi)$ are the atoms (literals) in $\mathrm{Atoms}_A(\tau)$ (and their negations). A valuation $\pi\colon \mathrm{Atoms}_A(\tau) \to S$ provides a valuation of the terminal positions in games $\mathcal{G}(A, \psi)$ for $\psi \in \mathrm{FO}^+(\tau)$. From there we can compute valuations in S of other positions in $\mathcal{G}(A, \psi)$, i.e., of subformulae $\varphi(\mathbf{a})$ of ψ.

Towards semiring valuations for games, we consider general games, not restricted to model-checking games.

Definition 5.1. A *game graph* is a structure $\mathcal{G} = (V, V_0, V_1, T, E)$, where $V = V_0 \dot\cup V_1 \dot\cup T$ is the set of positions, partitioned into the sets V_0, V_1 of the two players and the set T of terminal positions, and where $E \subseteq V \times V$ is the set of moves. We denote the set of immediate successors of a position v by $vE := \{w : (v, w) \in E\}$ and assume that $vE = \emptyset$ if, and only if, $v \in T$.

A *play* is a finite or infinite path $v_0 v_1 v_2 \dots$ through $\mathcal{G}$ where the successor $v_{i+1} \in v_i E$ is chosen by Player 0 if $v_i \in V_0$ and by Player 1 if $v_1 \in V_1$. A play ends when it reaches a terminal node $v_m \in T$.

Definition 5.2. For every game graph $\mathcal{G} = (V, V_0, V_1, T, E)$, and every position $v_0 \in V$, the *tree unraveling* of $\mathcal{G}$ from v_0 is the game tree $\mathcal{T}(\mathcal{G}, v_0)$ consisting of all finite paths from v_0. More precisely, $\mathcal{T}(\mathcal{G}, v) = (V^\#, V_0^\#, V_1^\#, T^\#, E^\#)$, where $V^\#$ is the set of all finite paths $\rho = v_0 v_1 \dots v_m$ through $\mathcal{G}$, with $V_\sigma^\# = \{\rho v \in V^\# : v \in V_\sigma\}$, $T^\# = \{\rho t \in V^\# : t \in T\}$, and $E^\# = \{(\rho v, \rho v v') : (v, v') \in E\}$.

For most game-theoretic considerations, the games played on $\mathcal{G}$ and its unravelings are equivalent, via the canonical projection $\pi \colon \mathcal{T}(\mathcal{G}, v_0) \to \mathcal{G}$, $\rho v \mapsto v$ mapping every path to its end point.

There are several possibilities to define the notion of a strategy formally. Here, we identify a strategy with the set of plays that it admits and view it as a subtree of $\mathcal{T}(\mathcal{G}, v_0)$.

Definition 5.3. A *strategy* of Player σ (for $\sigma \in \{0, 1\}$) from v_0 in a game $\mathcal{G}$ is a subtree $\mathcal{S} = (W, F)$ of $\mathcal{T}(\mathcal{G}, v_0)$ with $W \subseteq V^\#$ and $F \subseteq (W \times W) \cap E^\#$ such that

- W is closed under predecessors: if $\rho v \in W$ then also $\rho \in W$;
- if $\rho v \in W \cap V_\sigma^\#$, then $|(\rho v)F| = 1$;
- if $\rho v \in W \cap V_{1-\sigma}^\#$, then $(\rho v)F = (\rho v)E^\#$.

We write $\mathrm{Strat}_\sigma(v_0)$ for the set of all strategies of Player σ from v_0, and $\mathrm{Plays}(\mathcal{S})$ for the set of plays admitted by $\mathcal{S}$.

5.1 Well-founded Games

We now consider games on finite, acyclic game graphs $\mathcal{G}$, which do not admit infinite plays. We are interested in valuations $f_0, f_1 \colon V \to S$ of the positions of $\mathcal{G}$ in some semiring S, describing a value (or cost) of a position, from the point of view of Players 0 and 1. These are induced by valuations

$$f_\sigma \colon T \to S \qquad \text{(of the terminal positions),}$$

$h_\sigma \colon E \to S \setminus \{0\}$ (of the moves).

The simplest example is a reachability game $\mathcal{G}$ with objective $T_\sigma \subseteq T$ and valuations

$$f_\sigma(t) := \begin{cases} 1, & t \in T_\sigma \\ 0, & t \in T \setminus T_\sigma \end{cases} \quad \text{and} \quad h_\sigma(e) := 1 \text{ for all } e \in E.$$

More complicated examples may associate with terminal positions and moves non-trivial costs, access restrictions or confidences.

Definition 5.4. Given a well-founded game $\mathcal{G} = (V, V_0, V_1, T, E)$ with $f_\sigma \colon T \to S$ and $h_\sigma \colon E \to S \setminus \{0\}$, we define a valuation $f_\sigma \colon V \to S$ for Player σ by backwards induction:

$$f_\sigma(v) := \begin{cases} \sum_{w \in vE} h_\sigma(vw) \cdot f_\sigma(w), & \text{if } v \in V_\sigma, \\ \prod_{w \in vE} h_\sigma(vw) \cdot f_\sigma(w), & \text{if } v \in V_{1-\sigma}. \end{cases}$$

That is, a move from v to w contributes to $f_\sigma(v)$ the value $h_\sigma(vw) \cdot f_\sigma(w)$. We again use summation for alternative use of moves, and product for joint use of moves.

The valuations f_σ and h_σ further induce valuations of plays and strategies. Consider a finite play $x = v_0 v_1 \dots v_m$, with $v_m \in T$. We put $f_\sigma(x) := h_\sigma(v_0 v_1) \cdots h_\sigma(v_{m-1} v_m) \cdot f_\sigma(v_m)$. In other words, $f_\sigma(x)$ consists of the product over all edges $\prod_{e \in x} h_\sigma(e)$ and the valuation $f_\sigma(\text{outcome}(x))$, where the *outcome* of the play is the terminal position v_m the play ends in. For strategies, we take multiplicities of moves and positions into account.

Definition 5.5. Let $\mathcal{S} = (W, F) \in \text{Strat}_\sigma(v_0)$ be a strategy. For $v \in V$ and $e = (v, w) \in E$, set

$$\#_v(\mathcal{S}) := |\{\rho v : \rho v \in \mathcal{S}\}| = |\pi_{\mathcal{S}}^{-1}(v)|,$$
$$\#_e(\mathcal{S}) := |\{\rho v \in \mathcal{S} : \rho v \to \rho v w \text{ is a move in } \mathcal{S}\}| = |\pi_{\mathcal{S}}^{-1}(e)|,$$

where $\pi_{\mathcal{S}} \colon (W, F) \to (V, E)$ is the restriction of the canonical projection

$\pi\colon \mathcal{T}(\mathcal{G}, v_0) \to \mathcal{G}$ to $\mathcal{S}$. Then

$$F(\mathcal{S}) := \prod_{e \in E} h_\sigma(e)^{\#_e(\mathcal{S})} \cdot \prod_{t \in T} f_\sigma(t)^{\#_t(\mathcal{S})}$$

is the valuation of $\mathcal{S}$ in the semiring S.

Remark 5.6. In some important special cases the valuation of strategies coincides with the product of the valuations of plays they admit.

Lemma 5.7. If $h_\sigma(e) = 1$ for all $e \in E$, or if the semiring S is multiplicatively idempotent, then $F(\mathcal{S}) = \prod_{x \in \mathrm{Plays}(\mathcal{S})} f_\sigma(x)$.

Proof. For each terminal position $t \in T$, let

$$\#_t(\mathcal{S}) = |\{x \in \mathrm{Plays}(\mathcal{S}) : \mathrm{outcome}(x) = t\}|.$$

For $e \in E$, we have

$$\#_e(\mathcal{S}) > 0 \quad \Longleftrightarrow \quad e \in x \text{ for some play } x \in \mathrm{Plays}(\mathcal{S}).$$

Since $h_\sigma(e)^n = h_\sigma(e)$ if $n > 0$ (and $h_\sigma(e)^n = 1$ if $n = 0$), we have

$$h_\sigma(e)^{\#_e(\mathcal{S})} = \begin{cases} h_\sigma(e), & e \in x \text{ for some } x \in \mathrm{Plays}(\mathcal{S}), \\ 1, & \text{otherwise.} \end{cases}$$

Thus,

$$\begin{aligned} &\prod_{x \in \mathrm{Plays}(\mathcal{S})} f_\sigma(x) \\ &= \prod_{x \in \mathrm{Plays}(\mathcal{S})} \left(\prod_{e \in x} h_\sigma(e) \right) \cdot f_\sigma(\mathrm{outcome}(x)) \\ &= \prod_{\substack{e \in x, \\ x \in \mathrm{Plays}(\mathcal{S})}} h_\sigma(e) \cdot \prod_{t \in T} f_\sigma(t)^{|\{x \in \mathrm{Plays}(\mathcal{S}) \,:\, \mathrm{outcome}(x) = t\}|} \\ &= \prod_{e \in E} h_\sigma(e)^{\#_e(\mathcal{S})} \cdot \prod_{t \in T} f_\sigma(t)^{\#_t(\mathcal{S})} = F(\mathcal{S}). \end{aligned}$$

Q.E.D.

However, there are simple examples where $F(\mathcal{S}) \neq \prod_{x \in \mathrm{Plays}(\mathcal{S})} f_\sigma(x)$ if the semiring is not multiplicatively idempotent.

Example 5.8. Consider the following game in a semiring with an element a such that $a^2 \neq a$.

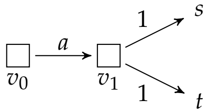

$$f_0(s) = f_0(t) = 1$$
$$h_0(v_0 v_1) = a$$

For the unique strategy $\mathcal{S}$ of Player 0 (do nothing), we have $F(\mathcal{S}) = a$, but there are two plays in $\text{Plays}(\mathcal{S})$ each with value a, so $\prod_{x \in \text{Plays}(\mathcal{S})} f_\sigma(x) = a^2$. ∎

Theorem 5.9 (Sum-of-Strategies)**.** Let $\mathcal{G}$ be any finite acyclic game with valuations $f_\sigma \colon T \to S$ and $h_\sigma \colon E \to S \setminus \{0\}$, inducing the valuation $f_\sigma \colon V \to S$. Then for all $v \in V$,

$$f_\sigma(v) = \sum_{\mathcal{S} \in \text{Strat}_\sigma(v)} F(\mathcal{S}).$$

Proof. For terminal positions v the claim is trivially true. So suppose that $v \in V_\sigma$. Then any strategy $\mathcal{S} \in \text{Strat}_\sigma(v)$ can be written in the form $\mathcal{S} = v \cdot \mathcal{S}'$ for some successor $w \in vE$ and some strategy $\mathcal{S}' \in \text{Strat}_\sigma(w)$.

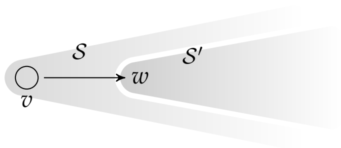

Clearly, $\#_t(\mathcal{S}) = \#_t(\mathcal{S}')$ for every terminal position $t \in T$. For the moves we have that $\#_e(\mathcal{S}) = \#_e(\mathcal{S}')$ for all $e \neq (v, w)$ but $\#_e(\mathcal{S}) = 1$ and $\#_e(\mathcal{S}') = 0$ for $e = (v, w)$. This implies that $F(S) = h(vw) \cdot F(\mathcal{S}')$. By induction hypothesis $f_\sigma(w) = \sum_{\mathcal{S}' \in \text{Strat}_\sigma(w)} F(\mathcal{S}')$. Hence

$$\begin{aligned} f_\sigma(v) &\overset{\text{def}}{=} \sum_{w \in vE} h_\sigma(vw) \cdot f_\sigma(w) \\ &\overset{\text{IH}}{=} \sum_{w \in vE} h_\sigma(vw) \cdot \sum_{\mathcal{S}' \in \text{Strat}_\sigma(w)} F(\mathcal{S}') \\ &= \sum_{w \in vE} \sum_{\mathcal{S}' \in \text{Strat}_\sigma(w)} h_\sigma(vw) \cdot F(\mathcal{S}') = \sum_{\mathcal{S} \in \text{Strat}_\sigma(v)} F(\mathcal{S}). \end{aligned}$$

Finally, let $v \in V_{1-\sigma}$ with $vE = \{w_1, \dots, w_n\}$. Every strategy $\mathcal{S} \in \mathrm{Strat}_\sigma(v)$ has the form $\mathcal{S} = v(\mathcal{S}_1 \cup \dots \cup \mathcal{S}_n)$ with $\mathcal{S}_i \in \mathrm{Strat}_\sigma(w_i)$. For the terminal nodes $t \in T$ we have that $\#_t(\mathcal{S}) = \sum_{i \le n} \#_t(\mathcal{S}_i)$; similarly, for all moves e from a different position than v, we have $\#_e(\mathcal{S}) = \sum_{i \le n} \#_e(\mathcal{S}_i)$, but for the moves $e = (v, w_i)$ we have $\#_e(\mathcal{S}) = 1$ and $\#_e(\mathcal{S}_i) = 0$ for all i. Thus

$$F(\mathcal{S}) = \prod_{w_i \in vE} h_\sigma(vw_i) \cdot F(\mathcal{S}_i).$$

It follows that

$$\begin{aligned}
f_\sigma(v) &= \prod_{w_i \in vE} h_\sigma(vw_i) \cdot f_\sigma(w_i) \\
&= \prod_{w_i \in vE} \left(h_\sigma(vw_i) \cdot \sum_{\mathcal{S}_i \in \mathrm{Strat}_\sigma(w_i)} F(\mathcal{S}_i) \right) \\
&= \prod_{w_i \in vE} \left(\sum_{\mathcal{S}_i \in \mathrm{Strat}_\sigma(w_i)} h_\sigma(vw_i) \cdot F(\mathcal{S}_i) \right) \\
&\overset{(*)}{=} \sum_{v \cdot (\mathcal{S}_1 \cup \dots \mathcal{S}_n) \in \mathrm{Strat}_\sigma(v)} \; \prod_{w_i \in vE} h_\sigma(vw_i) \cdot F(\mathcal{S}_i) \\
&= \sum_{\mathcal{S} \in \mathrm{Strat}_\sigma(v)} F(\mathcal{S}),
\end{aligned}$$

where $(*)$ uses the distributive law

$$\prod_{i=1}^{n} \left(\sum_{j \in A_i} a_{ij} \right) = \sum_{\substack{(j_1, \dots, j_n) \\ \in A_1 \times \dots \times A_n}} a_{1j_1} \cdot a_{2j_2} \cdots a_{nj_n} = \sum_{\substack{(j_1, \dots, j_n) \\ \in A_1 \times \dots \times A_n}} \left(\prod_{i=1}^{n} a_{ij_i} \right).$$

Q.E.D.

5.2 Game Valuations in Provenance Semirings

Let X be a set of variables X_e for moves $e \in E$ and X_t for terminal positions T, and consider $f_\sigma \colon T \to X \cup \{1\}$ and $h_\sigma \colon E \to X \cup \{1\}$ with

$$\begin{aligned} f_\sigma(t) &= X_t \qquad \text{or} \qquad && f_\sigma(t) = 1, \\ h_\sigma(e) &= X_e \qquad \text{or} \qquad && \underbrace{h_\sigma(e) = 1.}_{\text{for untracked } t/e} \end{aligned}$$

We get a valuation $f_\sigma\colon V \to \mathbb{N}[X]$ by the rules

$$f_\sigma(v) = \begin{cases} \sum\limits_{w \in vE} h_\sigma(vw) \cdot f_\sigma(w), & \text{for } v \in V_\sigma, \\ \prod\limits_{w \in vE} h_\sigma(vw) \cdot f_\sigma(w), & \text{for } v \in V_{1-\sigma}. \end{cases}$$

Similarly, we get valuations $f_\sigma\colon V \to K(X)$ for other provenance semirings $K(X) = \mathbb{B}[X], \mathrm{Trio}[X], \mathbb{W}[X], \dots$ (by the same rules), and of course:

$$\begin{tikzcd} V \arrow[r, "f_\sigma"] \arrow[dr, "f_\sigma"'] & \mathbb{N}[X] \arrow[d, two heads] \\ & K(X) \end{tikzcd}$$

What do these valuations tell us? To simplify notation, we write e for X_e and t for X_t.

$\underline{\mathbb{N}[X]:}$

Each $f_\sigma(v) \in \mathbb{N}[X]$ is a sum of monomials

$$c_m \cdot m = c_m \cdot e_1^{\alpha_1} \cdots e_r^{\alpha_r} \cdot t_1^{\beta_1} \cdots t_s^{\beta_s}$$

which, by the Sum-of-Strategies Theorem, tells us that Player σ has exactly c_m strategies $\mathcal{S} \in \mathrm{Strat}_\sigma(v)$ that

- use the moves $e_1, \dots, e_r$ (and possibly further untracked ones),
- have the outcomes $t_1, \dots, t_s$ (and possibly other untracked ones),
- use the move e_i precisely α_i times, and have precisely β_j plays with outcome t_j:

$$\#_{e_i}(\mathcal{S}) = \alpha_i \quad \text{and} \quad \#_{t_j}(\mathcal{S}) = \beta_j. \tag{$*$}$$

Example 5.10. Consider a game $\mathcal{G}$ where Player 0 has a reachability objective $W \subseteq T$. We just want to track the terminal positions, i.e. we put $f_0 \colon t \mapsto X_t \hat{=} t$, and $h_0 \colon e \mapsto 1$ (for all $e \in E$), so that we get valuations $f_0(v) \in \mathbb{N}[T]$.

Specialize these valuations by setting $t := 0$ for $t \in T \setminus W$. More precisely, consider $g \colon T \to \mathbb{N}[T]$ with $g(t) = t$ for $t \in W$ and $g(t) = 0$ otherwise. By the universal property, this extends to $\hat{g} \colon \mathbb{N}[T] \to \mathbb{N}[T]$ such that $\hat{g} \colon f_0(v) \mapsto f_0^W(v) := f_0(v)[t \mapsto 0]_{t \in T \setminus W}$.

Observe that $f_0^W(v)$ is a sum of monomials $c_m \cdot t_1^{\beta_1} \cdots t_s^{\beta_s}$ for $t_1, \dots, t_s \in W$; each such monomial tells us that Player 0 has c_m *winning* strategies, admitting precisely β_j plays ending in $t_j \in W$ (for each j). ■

<u>$\mathbb{B}[X]$:</u>
$f_\sigma(v) \in \mathbb{B}[X]$ is a sum of monomials $m = e_1^{\alpha_1} \cdots e_r^{\alpha_r} \cdot t_1^{\beta_1} \cdots t_s^{\beta_s}$ telling us that Player σ has *at least one* strategy $\mathcal{S} \in \mathrm{Strat}_\sigma(v)$ with property $(*)$.

<u>$\mathrm{Trio}[X]$:</u>
$f_\sigma(v) \in \mathrm{Trio}[X]$ is a sum of monomials $c_m \cdot e_1 \cdots e_r \cdot t_1 \cdots t_s$ telling us that Player σ has precisely c_m strategies making use of precisely the moves in $\{e_1, \dots, e_r\}$ and with outcomes $\{t_1, \dots, t_s\}$.

<u>$\mathbb{W}[X]$:</u>
Analogously (without coefficients).

To understand the valuations in $\mathbb{S}[X]$ and $\mathrm{PosBool}[X]$, we first have to discuss the notion of absorption among strategies.

5.3 Absorption on Strategies

Definition 5.11. Let $Y \subseteq E \cup T$ and $\mathcal{S}, \mathcal{S}' \in \mathrm{Strat}_\sigma(v)$. We say that $\mathcal{S}$ *absorbs* $\mathcal{S}'$ (w.r.t. Y), denoted $\mathcal{S} \succcurlyeq_Y \mathcal{S}'$, if $\#_e(\mathcal{S}) \leq \#_e(\mathcal{S}')$ and $\#_t(\mathcal{S}) \leq \#_t(\mathcal{S}')$ for all $e, t \in Y$.

An *absorption-dominant* strategy $\mathcal{S} \in \mathrm{Strat}_\sigma(v)$ (w.r.t. Y) is one that is *not* absorbed by any other strategy $\mathcal{S}'$ with

$$\#_y(\mathcal{S}') \leq \#_y(\mathcal{S}) \qquad \text{for all } y \in Y,$$

$$\#_y(\mathcal{S}') < \#_y(\mathcal{S}) \qquad \text{for at least one } y \in Y.$$

For acyclic games (and especially model-checking games for FO) we are mainly interested in absorption(-dominance) w.r.t. $Y := T$ (this will be different for games with infinite plays). Absorption-dominant strategies are in a sense strategies with minimal effort: to establish $\varphi \vee (\varphi \wedge \psi)$ we can either choose a strategy that establishes φ, or one that establishes $\varphi \wedge \psi$. The second one will be absorbed by the first one.

$\mathbb{S}[X]$:

Assuming that we only track the terminal positions, $f_\sigma(v) \in \mathbb{S}[T]$ is a sum of monomials $m = t_1^{\beta_1} \cdots t_r^{\beta_r}$ each of which stands for an *absorption-dominant* strategy $\mathcal{S} \in \mathrm{Strat}_\sigma(v)$ with precisely β_i plays with outcome t_i (for $i = 1, \dots, r$). This means that every other strategy $\mathcal{S}' \in \mathrm{Strat}_\sigma(v)$ either admits a play with outcome $t \notin \{t_1, \dots, t_r\}$ or has *more* than β_i plays with outcome t_i, for *some* $i \le r$, or also has precisely β_i plays with outcome t_i, for *all* $i \le r$.

$\mathrm{PosBool}[X]$:

Valuations $f_\sigma(v) \in \mathrm{PosBool}[T]$ lose the information about the number of plays with a particular outcome. Monomials are of the form $t_1 \dots t_r$ and say that $\{t_1, \dots, t_r\}$ is a *minimal set of outcomes* for strategies in $\mathrm{Strat}_\sigma(v)$.

5.4 Game Valuations in Application Semirings

5.4.1 Cost of Strategies

Given a game $\mathcal{G} = (V, V_0, V_1, T, E)$, associate with Player 0 cost functions $f_0 \colon T \to \mathbb{R}_+^\infty$ and $h_0 \colon E \to \mathbb{R}_+^\infty$ for terminal positions and moves. Define the cost of a strategy $\mathcal{S} \in \mathrm{Strat}_0(v)$ as the sum of all moves and outcomes it admits, weighted by the number of their occurrences:

$$\mathrm{cost}(\mathcal{S}) := \sum_{e \in E} \#_e(\mathcal{S}) \cdot h_0(e) + \sum_{t \in T} \#_t(\mathcal{S}) \cdot f_0(t).$$

Proposition 5.12. The cost of an optimal strategy from v for Player 0 is given by the valuation $f_0(v)$ in the tropical semiring $\mathbb{T} = (\mathbb{R}_+^\infty, \min, +, \infty, 0)$.

Proof. We defined the valuation of a strategy $\mathcal{S} \in \mathsf{Strat}_0(v)$ in a semiring $(S, +, \cdot, 0, 1)$ as

$$F(\mathcal{S}) := \prod_{e \in E} h_\sigma(e)^{\#_e(\mathcal{S})} \cdot \prod_{t \in T} f_\sigma(t)^{\#_t(\mathcal{S})}.$$

Product in $\mathbb{T}$ is addition in $\mathbb{R}_+^\infty$, so $\mathrm{cost}(\mathcal{S}) = F(\mathcal{S})$ (in $\mathbb{T}$). Addition in $\mathbb{T}$ is minimization, so by the Sum-of-Strategies Theorem,

$$f_0(v) = \min_{\mathcal{S} \in \mathsf{Strat}_0(v)} \mathrm{cost}(\mathcal{S}),$$

which describes the minimal cost of a strategy for Player 0 from v. Q.E.D.

5.4.2 Clearance Levels

Recall the access control semiring $\mathbb{A} = (\{\mathsf{0} < \mathsf{T} < \mathsf{S} < \mathsf{C} < \mathsf{P} = 1\}, \max, \min, \mathsf{0}, \mathsf{P})$ and consider functions $f_\sigma : T \to \mathbb{A}$ and $h_\sigma : E \to \mathbb{A} \setminus \{0\}$ to define access levels for the moves and terminal positions (Player σ can make a move e if, and only if, her personal clearance level is at least $h(e)$, similarly for $t \in T$). We say that Player σ wins if she can reach an accessible terminal position.

Proposition 5.13. The valuation $f_\sigma(v) \in \mathbb{A}$ describes the *minimal clearance level* that Player σ needs to win from position v.

Proof. Since $\mathbb{A}$ is multiplicatively idempotent, we have

$$h_\sigma(e)^{\#_e(\mathcal{S})} = \begin{cases} h_\sigma(e), & e \text{ occurs in } \mathcal{S}, \\ 0, & \text{otherwise} \end{cases}$$

$$f_\sigma(t)^{\#_t(\mathcal{S})} = \begin{cases} f_\sigma(t), & t \text{ occurs in } \mathcal{S}, \\ 0, & \text{otherwise} \end{cases}$$

Hence, for every strategy $\mathcal{S} \in \mathrm{Strat}_\sigma(v)$,

$$F(\mathcal{S}) = \min(\{h_\sigma(e) : e \in \mathcal{S}\} \cup \{f_\sigma(t) : t \in \mathcal{S}\})$$

and, by the Sum-of-Strategies Theorem,

$$f_\sigma(v) = \max_{\mathcal{S} \in \mathrm{Strat}_\sigma(v)} F(\mathcal{S}).$$

If the clearance level of Player σ is at least $F_\sigma(v)$ then there exists a strategy $\mathcal{S}$ from v such that Player σ can use all moves in $\mathcal{S}$ and access the outcome of each compatible play $x \in \mathrm{Plays}(\mathcal{S})$. Q.E.D.

5.4.3 Counting Winning Strategies

Consider a reachability game with objective $W \subseteq T$ for Player σ. Put

$$f_\sigma(t) = \begin{cases} 1, & t \in W \\ 0, & t \in T \setminus W \end{cases} \quad \text{and} \quad h_\sigma(e) = 1, \text{ for all } e \in E.$$

Then the valuation $f_\sigma(v)$ in the natural semiring $(\mathbb{N}, +, \cdot, 0, 1)$ describes the *number of winning strategies* of Player σ from v. Indeed, $F(\mathcal{S}) = 1$ if $\mathcal{S}$ is winning from v, and $F(\mathcal{S}) = 0$ otherwise. The claim follows by the Sum-of-Strategies Theorem.

Note that we count *all* winning strategies, not just the positional ones.

Example 5.14. Consider the following game:

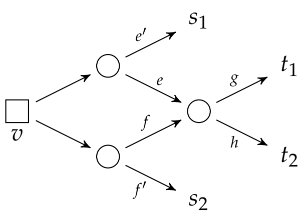

$$T = \{s_1, s_2, t_1, t_2\}$$
$$W = \{t_1, t_2\}$$

Player 0 has 9 strategies, 4 of which are winning (and absorption-dominant w.r.t. T), but only 2 winning strategies are positional:

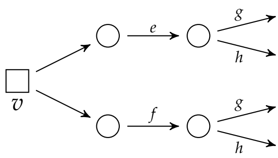

eg/fg (pos.)
eg/fh (not pos.)
eh/fg (not pos.)
eh/fh (pos.)

5.5 Separating Valuations

The valuations f_0, f_1 for the two players in a game $\mathcal{G}$, as defined by Definition 5.4, are *a priori* completely independent of each other. This admits the treatment of a wide variety of games, without any restrictions on how the objectives of the two players relate to each other.

In *antagonistic games*, valuations f_0, f_1 shall reflect conflicting objectives.

Definition 5.15. Let f_0, f_1 be the valuations of the two players in a game $\mathcal{G}$, and let $U \subseteq V$ be a set of positions. We say that

(1) f_0, f_1 are *weakly antagonistic* on U if

$$\forall u \in U\big(f_0(u) \cdot f_1(u) = 0\big),$$

(2) f_0, f_1 are *antagonistic* on U if

$$\forall u \in U\big(f_0(u) = 0 \text{ or } f_1(u) = 0\big),$$

(3) f_0 and f_1 are *additively positive* on U if

$$\forall u \in U\big(f_0(u) + f_1(u) \neq 0\big),$$

(4) f_0 and f_1 are *strongly antagonistic* on U if they are both antagonistic and additively positive U. In other words,

$$\forall u \in U\big(f_0(u) = 0 \text{ if, and only if, } f_1(u) \neq 0\big).$$

Notice that if the underlying semiring $\mathcal{S}$ has no divisors of 0, then weakly antagonistic valuations are in fact antagonistic.

Proposition 5.16. (1) If f_0, f_1 are weakly antagonistic on T, then they are so on all positions.

(2) If f_0, f_1 are antagonistic on T, then they are so on all positions.

(3) Assume that $\mathcal{S}$ is positive. If f_0, f_1 are additively positive on T, then they are so on all positions. It follows that if f_0, f_1 are strongly antagonistic on T, then they are so on all positions.

Proof. Since $h_\sigma \colon E \to S \setminus \{0\}$ and we are only interested in whether valuations are 0 or not, we can assume that $h_\sigma(e) = 1$ for all $e \in E$. Recall that if $v \in V_\sigma$, then

$$f_\sigma(v) = \sum_{w \in vE} f_\sigma(w) \quad \text{and} \quad f_{1-\sigma}(v) = \prod_{w \in vE} f_{1-\sigma}(w).$$

For Claim (1), assume that f_0, f_1 are weakly antagonostic on vE. Then it follows that f_0, f_1 are also weakly antagonistic on v:

$$\begin{aligned} f_\sigma(v) \cdot f_{1-\sigma}(v) &= \Big(\sum_{w \in vE} f_\sigma(w)\Big) \cdot \Big(\prod_{w \in vE} f_{1-\sigma}(w)\Big) \\ &= \sum_{w \in vE} \Big(f_\sigma(w) \cdot \prod_{w' \in vE} f_{1-\sigma}(w')\Big) \\ &= \sum_{w \in vE} \big(\underbrace{f_\sigma(w) \cdot f_{1-\sigma}(w)}_{=0} \cdot \prod_{w' \in vE \setminus \{w\}} f_{1-\sigma}(w')\big) = 0. \end{aligned}$$

For Claim (2), assume that f_0 and f_1 are antagonistic on vE. Then

$$\begin{aligned} f_\sigma(v) \neq 0 \quad &\Longrightarrow \quad (\exists w \in vE)\ f_\sigma(w) \neq 0 \\ &\Longrightarrow \quad (\exists w \in vE)\ f_{1-\sigma}(w) = 0 \quad \Longrightarrow \quad f_{1-\sigma}(v) = 0. \end{aligned}$$

The corresponding implication for strongly antagonistic valuations does not hold for all semirings, but it holds for positive ones. So assume that $\mathcal{S}$ is positive, $v \in V_\sigma$, and that f_0, f_1 are additively positive on vE.

Then $f_\sigma(v) + f_{1-\sigma}(v) = 0$ holds if, and only if, $f_\sigma(w) = 0$ for all $w \in vE$ (+-positivity) and $f_{1-\sigma}(w) = 0$ for at least one $w \in vE$ (no divisors of 0). Hence $f_\sigma(v) + f_{1-\sigma}(v) = 0$ would imply that $f_0(w) + f_1(w) = 0$ for some $w \in vE$, which contradicts our assumption. Q.E.D.

We remark that if the underlying semirings is not positive, then

Claim (3) does not hold for all games. A simple counterexample is given by game consisting of just one position $v \in V_1$, with two possible moves to the terminal positions t and t', with $f_0(t) = a$, $f_0(t') = a'$ and $f_1(t) = b$, $f_1(t') = b'$. Clearly $f_0(v) = a \cdot a'$ and $f_1(v) = b + b'$. If the semiring is not $+$-positive, then we may have that $a = a' = 0$, but b, b' are non-zero elements with $b + b' = 0$, so $f_0(v) = f_1(v) = 0$ although f_0, f_1 are additively positive on T. If the semiring has divisors of 0, then we may have $b = b' = 0$ and a, a' are non-trivial divisors of 0, so again, f_0, f_1 are additively positive on T but not on v.

Consider the case that $S = \mathbb{B}$. Strongly antagonistic valuations $f_0, f_1 \colon T \to \mathbb{B}$ define a strictly antagonistic reachability game. A play is won by Player σ if, and only if, it is lost by Player $1 - \sigma$. We just proved that then, every position v has strongly antagonistic valuations f_0, f_1, i.e. $f_\sigma(v) \neq 0 \iff f_{1-\sigma} = 0$. This is Zermelo's Theorem: in well-founded, strictly antagonistic games, one of the two players has a winning strategy, i.e. the game is determined.

6 Semiring Valuations for FO (with negation)

Negation poses a problem for semiring semantics, because in the semiring setting, negation is not a *compositional* logical operation. What does this mean? For $\psi \wedge \varphi$ and $\psi \vee \varphi$ the valuations $\pi[\![\psi \wedge \varphi]\!]$ and $\pi[\![\psi \vee \varphi]\!]$ are $\pi[\![\psi]\!] \cdot \pi[\![\varphi]\!]$ and $\pi[\![\psi]\!] + \pi[\![\varphi]\!]$ and are thus completely determined by $\pi[\![\psi]\!]$ and $\pi[\![\varphi]\!]$ (independent of the syntax of ψ and φ; we call this a compositional definition).

By the idea that in any semiring, the element 0 stands for *false* and all other elements for some shade of *true*, we would have that for any ψ with $\pi[\![\psi]\!] \neq 0$, necessarily $\pi[\![\neg\psi]\!] = 0$. But if $\pi[\![\psi]\!] = 0$ we cannot infer the value $\pi[\![\neg\psi]\!]$, without examining ψ itself (unless we are in the Boolean semiring where we have a unique value $\neq 0$; or if we give up the "axiom" that $\neg\neg\psi \equiv \psi$, as we could then set $\pi[\![\neg\psi]\!] = 1$ if $\pi[\![\psi]\!] = 0$, and $\pi[\![\neg\psi]\!] = 0$ otherwise, so that $\pi[\![\neg\neg\psi]\!] \in \{0, 1\}$ for all ψ).

We deal with this problem via transformation to negation normal form (nnf), i.e. we put $\pi[\![\neg\psi]\!] := \pi[\![\mathrm{nnf}(\neg\psi)]\!]$, or, equivalently, only consider formulae in FO that *are* in negation normal form.

Let τ be a relational vocabulary, A a (finite) universe, S a semiring. Let $\mathrm{Lit}_A(\tau) = \mathrm{Atoms}_A(\tau) \cup \mathrm{NegAtoms}_A(\tau)$, where $\mathrm{NegAtoms}_A(\tau) = \{\neg\alpha \mid \alpha \in \mathrm{Atoms}_A(\tau)\}$.

Definition 6.1. An S-interpretation (for τ and A) is a function $\pi\colon \mathrm{Lit}_A(\tau) \to S$. It extends to a valuation for all first-order sentences $\psi(\mathbf{a}) \in \mathrm{FO}(\tau \cup A)$ by interpreting equalities and inequalities by their truth values,

$$\pi[\![a_i = a_j]\!] = \begin{cases} 1, & \text{if } a_i = a_j, \\ 0, & \text{if } a_i \neq a_j, \end{cases} \qquad \pi[\![a_i \neq a_j]\!] = \begin{cases} 1, & \text{if } a_i \neq a_j, \\ 0, & \text{if } a_i = a_j, \end{cases}$$

logical connectives and quantifiers by the semiring operations (e.g. $\pi[\![\psi \vee \varphi]\!] = \pi[\![\psi]\!] + \pi[\![\varphi]\!]$) and by interpreting negation by transformation to negation normal form:

$$\pi[\![\neg\psi]\!] = \pi[\![\mathrm{nnf}(\neg\psi)]\!].$$

6.1 Game-theoretic view

Let $\mathcal{G}(A, \psi)$ be the game graph induced by $\psi \in \mathrm{FO}(\tau)$ and A. Its terminal positions are literals in $\mathrm{Lit}_A(\tau)$, or equalities and inequalities. An S-interpretation $\pi\colon \mathrm{Lit}_A(\tau) \to S$ thus provides a valuation $f_0\colon T \to S$ of the terminal positions of $\mathcal{G}(A, \psi)$. We put $f_1(\alpha) := f_0(\neg\alpha) = \pi(\neg\alpha)$. Further we put $h_\sigma(e) = 1$ for all moves e of $\mathcal{G}(A, \psi)$. We thus obtain valuations $f_0, f_1\colon V \to S$ for all positions v of $\mathcal{G}(A, \psi)$. Notice that these positions are instantiated subformulae $\varphi(\mathbf{a})$ of ψ.

Theorem 6.2. For all first-order sentences $\psi \in \mathrm{FO}(\tau)$, all semiring interpretations $\pi\colon \mathrm{Lit}_A(\tau) \to S$ and all positions $\varphi(\mathbf{a})$ of $\mathcal{G}(A, \psi)$ we have

$$f_0(\varphi(\mathbf{a})) = \pi[\![\varphi(\mathbf{a})]\!],$$
$$f_1(\varphi(\mathbf{a})) = \pi[\![\neg\varphi(\mathbf{a})]\!].$$

Proof. Obvious induction. Q.E.D.

Although this theorem holds without any restriction on the semiring interpretation $\pi\colon \mathrm{Lit}_A(\tau) \to S$, not all such interpretations are meaningful for logic. Normally we require that the value $\pi(R\mathbf{a})$ and $\pi(\neg R\mathbf{a})$ are related in a reasonable way.

Definition 6.3. An S-interpretation $\pi\colon \mathrm{Lit}_A(\tau) \to S$ is *model-defining*, if for each pair $\alpha, \neg\alpha$ of complementary literals, $\pi(\alpha) = 0 \iff \pi(\neg\alpha) \neq 0$. In that case π defines a unique τ-structure $\mathfrak{A}_\pi$ with universe A, with $R^{\mathfrak{A}_\pi} = \{\mathbf{a} \mid \pi(R\mathbf{a}) \neq 0\}$ for every $R \in \tau$.

Clearly, if $S \neq \mathbb{B}$, several model-defining S-interpretations may define the same structure.

Proposition 6.4. If S is positive and $\pi\colon \mathrm{Lit}_A(\tau) \to S$ is model-defining, then for all $\varphi \in \mathrm{FO}(\tau)$,

$$\mathfrak{A}_\pi \models \varphi \quad \Longleftrightarrow \quad \pi[\![\varphi]\!] \neq 0$$

Proof. In fact for any semiring S (positive or not) and any model-defining interpretation $\pi\colon \mathrm{Lit}_A(\tau) \to S$ we have that

$$\pi[\![\varphi]\!] \neq 0 \quad \Longrightarrow \quad \mathfrak{A}_\pi \models \varphi \tag{1}$$

(by trivial induction). For the converse, we observe that on each game graph $\mathcal{G}(A, \varphi)$, the valuations

$$f_0\colon \alpha \mapsto \pi(\alpha),$$
$$f_1\colon \alpha \mapsto \pi(\neg\alpha),$$

are strongly antagonistic on the terminal positions. Since S is positive, they are strongly antagonistic on all positions, in particular on φ. Hence

$$\pi[\![\varphi]\!] = f_0(\varphi) = 0 \quad \Longleftrightarrow \quad \pi[\![\neg\varphi]\!] = f_1(\varphi) \neq 0.$$

It follows that

$$\pi[\![\varphi]\!] = 0 \Longrightarrow \pi[\![\neg\varphi]\!] \neq 0 \overset{(1)}{\Longrightarrow} \mathfrak{A}_\pi \models \neg\varphi \Longrightarrow \mathfrak{A}_\pi \not\models \varphi. \tag{2}$$

The claim follows from (1) and (2). Q.E.D.

In semiring semantics, also S-interpretations that do not define a single structure are interesting. Additional issues arise for such interpretations $\pi\colon \mathrm{Lit}_A(\tau) \to S$:

- If for every atom $\alpha \in \mathrm{Atoms}_A(\tau)$ either $\pi(\alpha) = 0$ or $\pi(\neg\alpha) = 0$, then there is no $\varphi \in \mathrm{FO}(\tau)$ with $\pi[\![\varphi]\!] \neq 0$ and $\pi[\![\neg\varphi]\!] \neq 0$.
  (Valuations that are antagonistic on the terminal positions are so on all positions.)

- If for every atom $\alpha \in \mathrm{Atoms}_A(\tau)$ it holds $\pi(\alpha) \cdot \pi(\neg\alpha) = 0$, then also $\pi[\![\varphi]\!] \cdot \pi[\![\neg\varphi]\!] = 0$ for all $\varphi \in \mathrm{FO}(\tau)$.
  (Valuations that are weakly antagonistic on the terminal positions are so on all positions.)
- If S is positive, and for all atoms $\alpha \in \mathrm{Atoms}_A(\tau)$ we have that $\pi(\alpha) + \pi(\neg\alpha) \neq 0$, then also $\pi[\![\varphi]\!] + \pi[\![\neg\varphi]\!] \neq 0$ for all $\varphi \in \mathrm{FO}(\tau)$.
  (Additively positive valuations on the terminal positions are additively positive on all positions.)

6.2 Dual-indeterminate Polynomials

In the presence of negation, the semirings $\mathbb{N}[X]$, $\mathbb{B}[X]$, ... are not really appropriate anymore since they do not adequately represent the relationship between positive and negative literals.

Let $X, \overline{X}$ be two disjoint sets of indeterminates with a bijection $X \leftrightarrow \overline{X}$, $x \mapsto \overline{x}$. We use these indeterminates to annotate literals and we usually write x_α for the variable that annotates α.

Convention. If $x \in X$ is used to annotate α, we can use $\overline{x}$ only annotate the literal $\neg\alpha$. In other words, we forbid annotations of the form

$$\begin{array}{lll} \alpha \mapsto x & \beta \mapsto 0 & \\ \neg\alpha \mapsto 0 & \neg\beta \mapsto \overline{x} & \text{for } \alpha \neq \beta. \end{array}$$

On $K(X \cup \overline{X}) = \mathbb{N}[X \cup \overline{X}], \mathbb{B}[X \cup \overline{X}], \mathbb{S}[X \cup \overline{X}], \ldots$ let $\sim$ be the congruence generated by $x \cdot \overline{x} \sim 0$ for all $x \in X$. A monomial m in any $K(X \cup \overline{X})$ is *conflicting* if it contains both x and $\overline{x}$, for some $x \in X$ (with positive exponents, i.e., $m(x), m(\overline{x}) > 0$). For $f \in K(X \cup \overline{X})$, let $\hat{f}$ be obtained by deleting all conflicting monomials from f. Clearly $f \sim \hat{f}$ (since $m \sim 0$ for every conflicting monomial m).

Further observe that $f \sim g \Longleftrightarrow \hat{f} = \hat{g}$. (Indeed: $f \approx g \stackrel{\text{def}}{\Longleftrightarrow} \hat{f} = \hat{g}$ is a congruence with $x \cdot \overline{x} \approx 0$. Compatibility with $+$ is clear since for $f + g = h$ also $\hat{f} + \hat{g} = \hat{h}$. For $f \cdot g = h$ observe that $\hat{h}$ contains all non-conflicting monomials mm' for $m \in f$ and $m' \in g$ which is the same as the non-conflicting monomial mm' for $m \in \hat{f}$ and $m' \in \hat{g}$, thus $\hat{h} = \hat{f} \cdot \hat{g}$. Conversely, if $\hat{f} = \hat{g}$, then $f \sim \hat{f} = \hat{g} \sim g$.)

Thus $\mathbb{N}[X, \overline{X}] := \mathbb{N}[X \cup \overline{X}]/_\sim$ is in one-to-one correspondence with those polynomials in $\mathbb{N}[X \cup \overline{X}]$ that do not contain any conflicting monomials, and similarly for $\mathbb{B}[X, \overline{X}]$, $\mathbb{S}[X, \overline{X}]$, $\dots$. These are called semirings of *dual-indeterminate polynomials*.

Theorem 6.5 (Universal property)**.** Any function $h\colon X \cup \overline{X} \to S$ with $h(x) \cdot h(\overline{x}) = 0$ for all $x \in X$ extends to a unique homomorphism $\hat{h}\colon \mathbb{N}[X, \overline{X}] \to S$ (with $\hat{h}|_{X \cup \overline{X}} = h$).

Definition 6.6. A *provenance-tracking* labelling of $\mathrm{Lit}_A(\tau)$ is a map $\pi\colon \mathrm{Lit}_A(\tau) \to X \cup \overline{X} \cup \{0, 1\}$ mapping atoms α to $\pi(\alpha) \in X \cup \{0, 1\}$ and negated atoms to $\pi(\neg\alpha) \in \overline{X} \cup \{0, 1\}$.

If π maps $\alpha / \neg\alpha$ to an indeterminate x or $\overline{x}$, then this literal is *tracked* through the model-checking computation. If it is mapped to $0/1$, then we do not track it, but we still need to take into account whether or not it holds in a given structure. A provenance-tracking labelling gives, for each semiring $K(X, \overline{X})$ of dual-indeterminate polynomials, a $K(X, \overline{X})$-interpretation $\pi\colon \mathrm{Lit}_A(\tau) \to K(X, \overline{X})$.

Assume that $\pi\colon \mathrm{Lit}_A(\tau) \to X \cup \overline{X} \cup \{0, 1\}$ is both provenance-tracking and model-defining. For every first-order sentence $\psi \in \mathrm{FO}(\tau)$ and every semiring $K(X, \overline{X})$ we obtain a valuation $\pi[\![\psi]\!] \in K(X, \overline{X})$ which gives us information about the winning strategies of Player 0 in the model-checking game for $\mathfrak{A}_\pi \models \psi$.

Consider $\pi[\![\psi]\!] \in \mathbb{N}[X, \overline{X}]$; it is a sum of monomials of the form $c \cdot x_1^{j_1} \cdots x_r^{j_r}$, each of which informs us that Player 0 has precisely c *winning strategies* $\mathcal{S} \in \mathrm{Strat}_0(\psi)$ with the property that

- all plays in $\mathrm{Plays}(\mathcal{S})$ end at a literal α with $\pi(\alpha) \in \{x_1, \dots, x_r\} \cup \{1\}$,
- there are precisely j_i plays in $\mathrm{Plays}(\mathcal{S})$ that end in a literal α with $\pi(\alpha) = x_i$.

Similarly for other dual-indeterminate semirings. For instance $\pi[\![\psi]\!] \in \mathbb{S}[X, \overline{X}]$ is a sum of monomials of the form $x_1^{j_1} \cdots x_r^{j_r}$ saying that

Player 0 has an absorption-dominant winning strategy for $\mathfrak{A}_\pi \models \psi$ with that property. Further, although the semirings $K(X, \overline{X})$ are not positive (since $x \cdot \overline{x} = 0$), we still have $\pi[\![\psi]\!] \neq 0 \Longleftrightarrow \mathfrak{A}_\pi \models \psi$.

Composing $\pi\colon \mathrm{Lit}_A(\tau) \to X \cup \overline{X} \cup \{0,1\}$ with a map $h\colon X \cup \overline{X} \cup \{0,1\} \to S$ (into an arbitrary semiring S) with $h(x) \cdot h(\overline{x}) = 0$, $h(0) = 0$ and $h(1) = 1$ results in S-valuations for $\mathrm{FO}(\tau)$:

$$
\begin{tikzcd}
\mathrm{Lit}_A(\tau) \arrow[d, "\pi"] \arrow[dr, "h \circ \pi"] & \\
X \cup \overline{X} \cup \{0,1\} \arrow[r, "h"] & S
\end{tikzcd}
\qquad
\begin{tikzcd}
\mathrm{FO}(\tau) \arrow[d, "\pi"] \arrow[dr, "\hat{h} \circ \pi"] & \\
\mathbb{N}[X, \overline{X}] \arrow[r, "\hat{h}"] & S
\end{tikzcd}
$$

$S = \mathbb{T}$: cost computation for $\mathrm{FO}(\tau)$,
$S = \mathbb{A}$: clearance levels for $\mathrm{FO}(\tau)$,
$S = \mathbb{V}$: confidence scores,
$S = \mathbb{N}$: counting evaluation strategies.

What about $K(X, \overline{X})$-interpretations that are not model-defining?

Definition 6.7. A provenance-tracking labelling $\pi\colon \mathrm{Lit}_A(\tau) \to X \cup \overline{X} \cup \{0,1\}$ is *model-compatible* if for each atom α, either

- $\pi(\alpha) = x$ and $\pi(\neg\alpha) = \overline{x}$ for some $x \in X$, or
- $\pi(\alpha) = 0$ and $\pi(\neg\alpha) = 1$, or
- $\pi(\alpha) = 1$ and $\pi(\neg\alpha) = 0$.

A structure $\mathfrak{A}$ (with universe A) is *compatible* with π if $\mathfrak{A} \models \alpha$ for every literal $\alpha \in \mathrm{Lit}_A(\tau)$ with $\pi(\alpha) = 1$, and we set $\mathrm{Mod}_\pi := \{\mathfrak{A} \mid \mathfrak{A} \text{ is compatible with } \pi\}$. We say that $\psi \in \mathrm{FO}(\psi)$ is

- *Mod_π-satisfiable* if $\mathfrak{A} \models \psi$ for *some* $\mathfrak{A} \in \mathrm{Mod}_\pi$,
- *Mod_π-valid* if $\mathfrak{A} \models \psi$ for *all* $\mathfrak{A} \in \mathrm{Mod}_\pi$.

Let $\pi\colon \mathrm{Lit}_A(\tau) \to X \cup \overline{X} \cup \{0,1\}$ be model-compatible, and $\psi \in \mathrm{FO}(\tau)$. For each provenance semiring $K(X, \overline{X})$, the valuation $\pi[\![\psi]\!] \in K(X, \overline{X})$ gives information about evaluation strategies and their use of literals mapped to $X \cup \overline{X}$. For $\mathbb{N}[X, \overline{X}]$, each monomial $c \cdot x_1^{j_1} \cdots x_r^{j_r} \in$

$\pi[\![\psi]\!]$ indicates that there are c strategies for the game $\mathcal{G}(A, \psi)$ from ψ all whose outcomes are literals α with $\pi(\alpha) \in \{x_1, \ldots, x_r\} \cup \{1\}$ (and there are precisely j_i plays with outcome labelled by x_i).

Corollary 6.8. Let π be model-compatible and $\psi \in \mathrm{FO}(\tau)$. Then,

- ψ is Mod_π-satisfiable if $\pi[\![\psi]\!] \neq 0$,
- ψ is Mod_π-valid if $\pi[\![\neg\psi]\!] = 0$.

6.3 Example: Model-defining tracking

Consider the sentence $\psi \in \mathrm{FO}(\{E\})$ defined as follows:

$$\begin{aligned} \mathrm{dom}(x) &:= \forall y (x = y \vee (Exy \wedge \neg Eyx)) \\ \psi &:= \forall x\, \neg \mathrm{dom}(x) && \text{“no dominant vertex”} \\ \mathrm{nnf}(\psi) &= \forall x \exists y (x \neq y \wedge (\neg Exy \vee Eyx)) \end{aligned}$$

We evaluate ψ in the following provenance-tracking labelling:

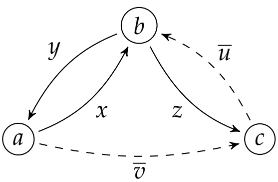

$A = \{a, b, c\}$

$X = \{x, y, z, u, v\}$

$$\pi \colon \mathrm{Lit}_A(\{E\}) \longrightarrow X \cup \overline{X} \cup \{0, 1\}$$

$$\left.\begin{aligned} Eab &\longmapsto x \\ Eba &\longmapsto y \\ Ebc &\longmapsto z \end{aligned}\right\} \text{tracked edges}$$

$$\left.\begin{aligned} \neg Ecb &\longmapsto \overline{u} \\ \neg Eac &\longmapsto \overline{v} \end{aligned}\right\} \text{tracked non-edges}$$

which becomes model-defining by mapping

$$\begin{aligned} \neg Eab, \neg Eba, \neg Ebc &\longmapsto 0 \\ Ecb, Eac &\longmapsto 0 \\[1ex] \text{all other atoms} &\longmapsto 0 \\ \text{all other neg. atoms} &\longmapsto 1 \end{aligned} \quad \left.\begin{matrix} \\ \\ \end{matrix}\right\} \text{untracked}$$

The model induced by π is:

$\mathfrak{A}_\pi$:

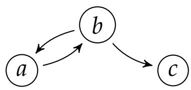

We can compute the value of ψ as follows:

$$\begin{aligned} \pi[\![\psi]\!] &= \pi[\![\neg\mathrm{dom}(a) \wedge \neg\mathrm{dom}(b) \wedge \neg\mathrm{dom}(c)]\!] \\[1ex] &= \pi[\![(\neg \underset{0}{Eab} \vee \underset{y}{Eba}) \vee (\neg \underset{\overline{v}}{Eac} \vee \underset{0}{Eca})]\!] \\ &\quad \cdot \pi[\![(\neg \underset{0}{Eba} \vee \underset{x}{Eab}) \vee (\neg \underset{0}{Ebc} \vee \underset{0}{Ecb})]\!] \\ &\quad \cdot \pi[\![(\neg \underset{1}{Eca} \vee \underset{0}{Eac}) \vee (\neg \underset{\overline{u}}{Ecb} \vee \underset{z}{Ebc})]\!] \\ &= (y + \overline{v}) \cdot x \cdot (1 + \overline{u} + z) \\ &= xy + x\overline{v} + xy\overline{u} + x\overline{v}\overline{u} + xyz + x\overline{v}z \end{aligned}$$

Each monomial corresponds to a winning strategy for proving that $\mathfrak{A}_\pi \models \psi$.

- For instance, the monomial $xy\overline{u}$ describes the following strategy:

$$\begin{array}{llcccc} \text{from} & \neg\mathrm{dom}(a) & \rightsquigarrow & b & \qquad & Eba \triangleq y \\ & \neg\mathrm{dom}(b) & \rightsquigarrow & a & & Eab \triangleq x \\ & \neg\mathrm{dom}(c) & \rightsquigarrow & b & & \neg Ecb \triangleq \overline{u}, \end{array}$$

- The monomial $x\overline{v}$ is associated with a different strategy:

$$\begin{array}{rcl@{\qquad\qquad}r}
\neg\mathrm{dom}(a) & \rightsquigarrow & c & \neg Eac \mathrel{\hat{=}} \overline{v} \\
\neg\mathrm{dom}(b) & \rightsquigarrow & a & Eab \mathrel{\hat{=}} x \\
\neg\mathrm{dom}(c) & \rightsquigarrow & a & \neg Eca \mathrel{\hat{=}} 1.
\end{array}$$

We also see that the fact Eab labelled by x is *necessary* for proving $\mathfrak{A}_\pi \models \psi$ (it occurs in all winning strategies). Indeed, removing Eab from $\mathfrak{A}_\pi$ makes b a dominant vertex.

Cost computation in $\mathbb{T}$. Let $h\colon X \cup \overline{X} \cup \{0,1\} \to \mathbb{R}_+^\infty$ be given by

$$\begin{array}{lll}
x \mapsto 0, & y \mapsto 1, & z \mapsto 2, \\
\overline{u} \mapsto \infty, & \overline{v} \mapsto 5, &
\end{array}$$

resulting in

$$\begin{aligned}
\mathrm{cost}(\psi) &= \hat{h} \circ \pi[\![\psi]\!] \\
&= \min(\underset{1}{\hat{h}(xy)}, \underset{5}{\hat{h}(x\overline{v})}, \underset{\infty}{\hat{h}(xy\overline{u})}, \underset{\infty}{\hat{h}(x\overline{vu})}, \underset{3}{\hat{h}(xyz)}, \underset{7}{\hat{h}(x\overline{v}z)}) \\
&= 1 \qquad\qquad\qquad\qquad (\text{recall: } \hat{h}(xy) = h(x) + h(y)).
\end{aligned}$$

Access restriction in $\mathbb{A}$. Now consider the mapping

$$\begin{array}{lll}
x \mapsto \mathsf{P}, & y \mapsto \mathsf{T}, & z \mapsto \mathsf{C}, \\
\overline{u} \mapsto \mathsf{P}, & \overline{v} \mapsto \mathsf{P}. &
\end{array}$$

The required clearance level for ψ is then

$$\text{cl.-level}(\psi) = \max(\mathsf{T}, \mathsf{P}, \mathsf{T}, \mathsf{P}, \mathsf{T}, \mathsf{C}) = \mathsf{P}.$$

That is, $\mathfrak{A}_\pi \models \psi$ can be checked with public information. However, if your clearance level is not for top secret information, only the strategies corresponding to $x\overline{v}$, $x\overline{vu}$, and $x\overline{v}z$ are available, whose costs are 5, ∞, and 7. So only agents cleared for top secret information have a strategy with cost 1.

6.4 Example: Model-compatible tracking

We next move to applications of model-compatible interpretations. Modify the map $\pi\colon \mathrm{Lit}_A(\{E\}) \to X \cup \overline{X} \cup \{0,1\}$ of the previous example on the tracked edges and tracked non-edges of the model $\mathfrak{A}_\pi$ above by mapping the complementary literals to the complementary tokens:

$$\pi\colon\quad \begin{aligned} Eab/\neg Eab &\longmapsto x/\overline{x}, \\ Eba/\neg Eba &\longmapsto y/\overline{y}, \\ Ebc/\neg Ebc &\longmapsto z/\overline{z}, \\ Ecb/\neg Ecb &\longmapsto u/\overline{u}, \\ Eac/\neg Eac &\longmapsto v/\overline{v}, \end{aligned}$$

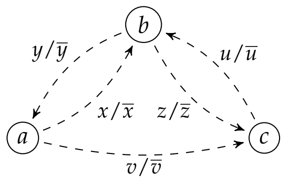

and the rest as above. Then

$$\mathrm{Mod}_\pi = \{\mathfrak{A} = (A, E) \mid \mathfrak{A} \models \neg Eaa \land \neg Ebb \land \neg Ecc \land \neg Eca\}.$$

We can again compute

$$\begin{aligned} \pi[\![\psi]\!] &= \pi[\![\neg\mathrm{dom}(a) \land \neg\mathrm{dom}(b) \land \neg\mathrm{dom}(c)]\!] \\ &= (\overline{x} + y + \overline{v} + 0)(\overline{y} + x + \overline{z} + u)(1 + v + \overline{u} + z) \\ &= \ldots \text{ 48 monomials, which reduces to 30 monomials} \\ &\qquad\ \text{by elimination of conflicting monomials.} \end{aligned}$$

On the other side,

$$\pi[\![\neg\psi]\!] = \pi[\![\mathrm{dom}(a) \lor \mathrm{dom}(b) \lor \mathrm{dom}(c)]\!] = x\overline{y}v + \overline{x}yz\overline{u} + 0.$$

Each of the two monomials in $\pi[\![\neg\psi]\!]$ gives us models for $\neg\psi$ in Mod_π:

- $x\overline{y}v$: models of the form

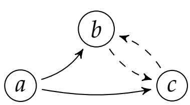

(*Ebc* and *Ecb* may or may not hold in these models),

- $\overline{x}yz\overline{u}$: models of the form

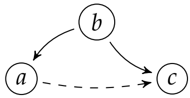

(Eac may or may not hold).

Consider the structure

$\mathfrak{B}$:

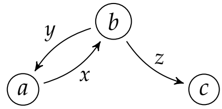

$\mathfrak{B}$ is *not* a model of $\neg\psi$. For $X = \{x, y, z, u, v\}$ let the edges/non-edges with labels in $X \cup \overline{X}$ be the one that can be added to / deleted from $\mathfrak{B}$. A *minimal repair* (for $\mathfrak{B}, X, \overline{X}$ and $\neg\psi$) is a minimal subset of $X \cup \overline{X}$ such that switching these edges/non-edges will update $\mathfrak{B}$ to a model of $\neg\psi$. We can read off the minimal repairs from $\pi[\![\neg\psi]\!]$:

- $x\overline{y}v$: delete Eba $(y \rightsquigarrow \overline{y})$,
  add Eac $(\overline{v} \rightsquigarrow v)$;

- $\overline{x}yz\overline{u}$: delete Eab $(x \rightsquigarrow \overline{x})$.

Given a cost labelling $h\colon X \cup \overline{X} \cup \{0, 1\} \to \mathbb{T}$ we may select:

- the repair with minimal cost, or
- the repaired model inducing the minimal cost for $\neg\psi$.

7 Elementary Equivalence versus Isomorphism

Definition 7.1. Let $\pi_A \colon \mathrm{Lit}_A(\tau) \to S$ and $\pi_B \colon \mathrm{Lit}_B(\tau) \to S$ be two S-interpretations.

- *Isomorphism:* $\pi_A \cong \pi_B$ if there is a bijection $h \colon A \xrightarrow{\sim} B$ such that for every τ-literal $\alpha(\mathbf{x})$ and all $\mathbf{a} \in A^k$: $\pi_A(\alpha(\mathbf{a})) = \pi_B(\alpha(h\mathbf{a}))$. We write $h \colon \pi_A \xrightarrow{\sim} \pi_B$.
- *Elementary equivalence:* $\pi_A, \mathbf{a} \equiv \pi_B, \mathbf{b}$ if for all $\varphi(\mathbf{x}) \in \mathrm{FO}(\tau)$: $\pi_A[\![\varphi(\mathbf{a})]\!] = \pi_B[\![\varphi(\mathbf{b})]\!]$.

Obviously, the Isomorphism Lemma also holds for semiring semantics.

Lemma 7.2 (Isomorphism Lemma)**.** If $h \colon \pi_A \xrightarrow{\sim} \pi_B$ is an isomorphism of S-interpretations, then for all tuples $\mathbf{a} \in A^k$: $\pi_A, \mathbf{a} \equiv \pi_B, h(\mathbf{a})$, and in particular $\pi_A \equiv \pi_B$.

In classical Boolean semantics, the converse holds for finite structures. For every *finite* τ-structure $\mathfrak{A}$ there exists a sentence $\psi_a \in \mathrm{FO}(\tau)$ such that $\mathfrak{B} \models \psi_a \iff \mathfrak{B} \cong \mathfrak{A}$. Hence $\mathfrak{B} \equiv \mathfrak{A} \iff \mathfrak{B} \cong \mathfrak{A}$.

Questions for any given semiring S:

(1) Are elementary equivalent finite S-interpretations always isomorphic?
(2) Is every finite S-interpretation π_A FO-axiomatisable? That is, is there a set $\Phi_A \subseteq \mathrm{FO}$ such that whenever $\pi_B[\![\varphi]\!] = \pi_A[\![\varphi]\!]$ for all $\varphi \in \Phi_A$, then $\pi_A \cong \pi_B$?
(3) Does every finite S-interpretation admit an axiomatisation by a *finite* set of axioms?
(4) ... by a *single* axiom?

Obviously, the following implications would hold for positive answers: (4) $\Longrightarrow$ (3) $\Longrightarrow$ (2) $\Longrightarrow$ (1). We shall prove that the answers depend on the semiring S.

- For min-max semirings with ≥ 3 elements, there exists $\pi_A \not\cong \pi_B$ with $\pi_A \equiv \pi_B$, so all answers are negative.
- For $\mathbb{V}$, $\mathbb{T}$, $\mathbb{N}$ and $\mathbb{N}[X]$, any finite interpretation is FO-axiomatisable, so (1) and (2) have positive answers.
- For $\mathbb{V}$ and $\mathbb{T}$, finite axiomatisations are possible, but not axiomatisations by a single axiom (so (3) and (4) are not always equivalent).

7.1 A Counterexample

Let $\tau = \{P, Q\}$ consist of two unary predicates, and let $S_4 := (\{0 < 1 < 2 < 3\}, \max, \min, 0, 3)$ (the min-max-semiring with four elements). For $A = \{a, b, c\}$, consider the following interpretations (notice that the interpretation of P and Q is switched):

$$\pi_{PQ} := \begin{array}{c|cccc} & P & Q & \neg P & \neg Q \\ \hline a & 1 & 3 & 0 & 0 \\ b & 2 & 1 & 0 & 0 \\ c & 3 & 2 & 0 & 0 \end{array} \qquad \pi_{QP} := \begin{array}{c|cccc} & P & Q & \neg P & \neg Q \\ \hline a & 3 & 1 & 0 & 0 \\ b & 1 & 2 & 0 & 0 \\ c & 2 & 3 & 0 & 0 \end{array}$$

Clearly $\pi_{PQ} \not\cong \pi_{QP}$. How can we prove that $\pi_{PQ} \equiv \pi_{QP}$? Idea: If $\pi_A \not\equiv \pi_B$, then this is witnessed by some pair $i \neq j$ in S and $\psi \in$ FO, with $\pi_A[\![\psi]\!] = i$ and $\pi_B[\![\psi]\!] = j$. To exclude (i, j) as such a witness, we look for homomorphisms $h_A, h_B \colon S \to S'$ with

- $h_A(i) \neq h_B(j)$,
- $h_A \circ \pi_A \equiv h_B \circ \pi_B$ (elementary equivalent S'-interpretations).

(This excludes (i, j), since $\pi_A[\![\psi]\!] = i$, $\pi_B[\![\psi]\!] = j$ imply (by the Fundamental Property) $(h_A \circ \pi_A)[\![\psi]\!] = h_A(i) \neq h_B(j) = (h_B \circ \pi_B)[\![\psi]\!]$, contradicting $h_A \circ \pi_A \equiv h_B \circ \pi_B$.)

We want to find enough homomorphisms to exclude *all* pairs $i \neq j$ as possible witnesses. For $S' = \mathbb{B}$, the equivalences $h_A \circ \pi_A \equiv h_B \circ \pi_B$ amount to isomorphism.

Definition 7.3. A set $H \subseteq \mathrm{Hom}^2(S, S')$ of homomorphism pairs $h, h' \colon S \to S'$ is *separating* if for all $i \neq j$ in S, there is a pair $(h, h') \in H$ with $h(i) \neq h'(j)$.

A *diagonal separating* set is a separating set H where $h = h'$ for all $(h, h') \in H$. We can present it as $D = \{h \in \mathrm{Hom}(S, S') \mid (h, h) \in H\}$.

We then obtain the following reduction: Let $\pi_A \colon \mathrm{Lit}_A(\tau) \to S$, $\pi_B \colon \mathrm{Lit}_B(\tau) \to S$, $\mathbf{a} \in A^k$, $\mathbf{b} \in B^k$, and $H \subseteq \mathrm{Hom}^2(S, S')$ be a separating set, $\varphi(\mathbf{x}) \in \mathrm{FO}(\tau)$. If $(h \circ \pi_A)[\![\varphi(\mathbf{a})]\!] = (h' \circ \pi_B)[\![\varphi(\mathbf{b})]\!]$ for all $(h, h') \in H$, then $\pi_A[\![\varphi(\mathbf{a})]\!] = \pi_B[\![\varphi(\mathbf{b})]\!]$ (if $\pi_A[\![\varphi(\mathbf{a})]\!] = i \neq j = \pi_B[\![\varphi(\mathbf{b})]\!]$, then $(h \circ \pi_A)[\![\varphi(\mathbf{a})]\!] = h(i) \neq h'(j) = (h' \circ \pi_B)[\![\varphi(\mathbf{b})]\!]$ for some $(h, h') \in H$, contradiction). Hence $h \circ \pi_A, \mathbf{a} \equiv h' \circ \pi_B, \mathbf{b}$ for all $(h, h') \in H$ implies that $\pi_A, \mathbf{a} \equiv \pi_B, \mathbf{b}$.

Let S_m be the min-max semiring over $\{0, \dots, m-1\}$. A diagonal separating set for S_m and $\mathbb{B}$ is $D = \{h_j \mid 1 \leq j \leq m-1\}$ with $h_j(i) = \bot$ if $i < j$, and $h_j(i) = \top$ if $i \geq j$, since every pair $i < j$ is separated by h_j. Recall the S_4-interpretations π_{PQ} and π_{QP} given above and consider their resulting $\mathbb{B}$-interpretations (i.e., finite structures) $h_j \circ \pi_{PQ}$ and $h_j \circ \pi_{QP}$, for $j \in \{1, 2, 3\}$. We have to show that $h_j \circ \pi_{PQ} \equiv h_j \circ \pi_{QP}$ for all $j \in \{1, 2, 3\}$. To this end, we show that they are isomorphic (which implies elementary equivalence):

$$h_1 \circ \pi_{PQ} : \begin{array}{c|cc} & P & Q \\ \hline a & \top & \top \\ b & \top & \top \\ c & \top & \top \end{array} \qquad = \qquad h_1 \circ \pi_{QP} : \begin{array}{c|cc} & P & Q \\ \hline a & \top & \top \\ b & \top & \top \\ c & \top & \top \end{array}$$

$$h_2 \circ \pi_{PQ} : \begin{array}{c|cc} & P & Q \\ \hline a & \bot & \top \\ b & \top & \bot \\ c & \top & \top \end{array} \qquad \underset{\text{(switch } a \text{ and } b\text{)}}{\cong} \qquad h_2 \circ \pi_{QP} : \begin{array}{c|cc} & P & Q \\ \hline a & \top & \bot \\ b & \bot & \top \\ c & \top & \top \end{array}$$

$$h_3 \circ \pi_{PQ} : \begin{array}{c|cc} & P & Q \\ \hline a & \bot & \top \\ b & \bot & \bot \\ c & \top & \bot \end{array} \quad \underset{\text{(switch } a \text{ and } c\text{)}}{\cong} \quad h_3 \circ \pi_{QP} : \begin{array}{c|cc} & P & Q \\ \hline a & \top & \bot \\ b & \bot & \bot \\ c & \bot & \top \end{array}$$

Corollary 7.4. $\pi_{PQ} \equiv \pi_{QP}$ (although $\pi_{PQ} \not\cong \pi_{QP}$).

7.2 Separating Homomorphisms for PosBool$[X]$

A similar approach works also for more general semirings, such as PosBool$[X]$ and $\mathbb{W}[X]$, and thus, by the universal property, also for all fully idempotent semirings S with $|S| \geq 3$. We illustrate this for PosBool$[x, y]$, $A = \{a, b, c, d\}$ and $\tau = \{P, Q\}$ (as above). Note that for PosBool$[X]$, any subset $Y \subseteq X$ induces a homomorphism

$$h_Y \colon \mathrm{PosBool}[X] \to \mathbb{B}, \qquad h_Y(x) = \begin{cases} \top, & \text{if } x \in Y, \\ \bot, & \text{if } x \in X \setminus Y. \end{cases}$$

Then $h_Y(f) = \top$ if, and only if, f contains a monomial all whose variables are in Y.

Lemma 7.5. $D_X := \{h_Y \mid Y \subseteq X\} \subseteq \mathrm{Hom}(\mathrm{PosBool}[X], \mathbb{B})$ is a diagonal separating set of homomorphisms.

Proof. Let $f, g \in \mathrm{PosBool}[X]$ with $f \neq g$. Let $m = \prod_{y \in Y} y$ be a monomial that occurs in one of f, g but not in the other, with minimal Y. We can assume that m occurs in f but not in g. Clearly $h_Y(f) = \top$, and we claim that $h_Y(g) = \bot$.

Otherwise g would contain a monomial m' with only variables from Y. Since m' has less variables than m, it must also occur in f. But m' absorbs m, so m does not occur in f, contradiction. Q.E.D.

Consider the following PosBool$[x, y]$-interpretations:

π_{xy} :

A	P	Q	$\neg P$	$\neg Q$
a	0	y	x	0
b	x	0	0	y
c	y	x	0	0
d	0	0	y	x

π_{yx} :

A	P	Q	$\neg P$	$\neg Q$
a	y	0	0	x
b	0	x	y	0
c	x	y	0	0
d	0	0	x	y

Obviously $\pi_{xy} \not\cong \pi_{yx}$. To prove that $\pi_{xy} \equiv \pi_{yx}$ it suffices to show that for all $h \in D_{xy} = \{h_\emptyset, h_x, h_y, h_{xy}\}$, we have $h \circ \pi_{xy} \cong h \circ \pi_{yx}$.

$h_\emptyset \circ \pi_{xy}$:

A	P	Q	$\neg P$	$\neg Q$
a	$\bot$	$\bot$	$\bot$	$\bot$
b	$\bot$	$\bot$	$\bot$	$\bot$
c	$\bot$	$\bot$	$\bot$	$\bot$
d	$\bot$	$\bot$	$\bot$	$\bot$

$h_\emptyset \circ \pi_{yx}$:

A	P	Q	$\neg P$	$\neg Q$
a	$\bot$	$\bot$	$\bot$	$\bot$
b	$\bot$	$\bot$	$\bot$	$\bot$
c	$\bot$	$\bot$	$\bot$	$\bot$
d	$\bot$	$\bot$	$\bot$	$\bot$

$h_x \circ \pi_{xy}$:

A	P	Q	$\neg P$	$\neg Q$
a	$\bot$	$\bot$	$\top$	$\bot$
b	$\top$	$\bot$	$\bot$	$\bot$
c	$\bot$	$\top$	$\bot$	$\bot$
d	$\bot$	$\bot$	$\bot$	$\top$

$h_x \circ \pi_{yx}$:

A	P	Q	$\neg P$	$\neg Q$
a	$\bot$	$\bot$	$\bot$	$\top$
b	$\bot$	$\top$	$\bot$	$\bot$
c	$\top$	$\bot$	$\bot$	$\bot$
d	$\bot$	$\bot$	$\top$	$\bot$

$h_y \circ \pi_{xy}$:

A	P	Q	$\neg P$	$\neg Q$
a	$\bot$	$\top$	$\bot$	$\bot$
b	$\bot$	$\bot$	$\bot$	$\top$
c	$\top$	$\bot$	$\bot$	$\bot$
d	$\bot$	$\bot$	$\top$	$\bot$

$h_y \circ \pi_{yx}$:

A	P	Q	$\neg P$	$\neg Q$
a	$\top$	$\bot$	$\bot$	$\bot$
b	$\bot$	$\bot$	$\top$	$\bot$
c	$\bot$	$\top$	$\bot$	$\bot$
d	$\bot$	$\bot$	$\bot$	$\top$

$h_{xy} \circ \pi_{xy}$:

A	P	Q	$\neg P$	$\neg Q$
a	$\bot$	$\top$	$\top$	$\bot$
b	$\top$	$\bot$	$\bot$	$\top$
c	$\top$	$\top$	$\bot$	$\bot$
d	$\bot$	$\bot$	$\top$	$\top$

$h_{xy} \circ \pi_{yx}$:

A	P	Q	$\neg P$	$\neg Q$
a	$\top$	$\bot$	$\bot$	$\top$
b	$\bot$	$\top$	$\top$	$\bot$
c	$\top$	$\top$	$\bot$	$\bot$
d	$\bot$	$\bot$	$\top$	$\top$

Since PosBool$[X]$ has the universal property for lattice semirings $(S, \sqcup, \sqcap, 0, 1)$, it follows that for each such semiring with at least three elements, there exist S-interpretations π_{rs} and π_{sr} which are elementary equivalent but not isomorphic: choose $r, s \neq 0$, $r \neq s$ and let π_{rs}, π_{sr} be obtained from π_{xy}, π_{yx} by $x \mapsto r$, $y \mapsto s$.

7.3 A different Story: The Viterbi Semiring

Recall the Viterbi semiring $\mathbb{V} = ([0,1], \max, \cdot, 0, 1)$. We want to show that every finite $\mathbb{V}$-interpretation is FO-axiomatisable up to isomorphism.

Let τ be a finite relational vocabulary and let $\mathrm{Lit}_n(\tau)$ be the set of literals $R\mathbf{x}/\neg R\mathbf{x}$ for $R \in \tau$ and $\mathbf{x}$ a tuple of variables from $x_1, \dots, x_n$. Recall that for every finite structure $\mathfrak{A}$ with universe $A = \{a_1, \dots, a_n\}$ there is a characteristic sentence $\chi_{\mathfrak{A}} \in \mathrm{FO}(\tau)$ such that $\mathfrak{B} \models \chi_{\mathfrak{A}} \iff \mathfrak{B} \cong \mathfrak{A}$. The characteristic sentence has the form

$$\chi_{\mathfrak{A}} \coloneqq \exists x_1 \dots \exists x_n (\varphi(\mathbf{x}) \wedge \psi(\mathbf{x}))$$

with

$$\begin{aligned}
\varphi(\mathbf{x}) &= \bigwedge_{1 \le i < j \le n} \big(x_i \neq x_j \wedge \forall y \bigvee_{i \le n} y = x_i\big), \\
\psi(\mathbf{x}) &= \bigwedge \{\alpha(\mathbf{x}) \in \mathrm{Lit}_n(\tau) \mid \mathfrak{A} \models \alpha(\mathbf{a})\}, \quad \mathbf{a} = \mathbf{x}[x_1/a_1, \dots, x_n/a_n].
\end{aligned}$$

Here, $\varphi(\mathbf{x})$ asserts that the universe has precisely n elements. This can be reused as-is for any semiring.

Lemma 7.6. For every S-interpretation $\pi_B \colon \mathrm{Lit}_B(\tau) \to S$ into an arbitrary semiring S and every tuple $\mathbf{b} = (b_1, \dots, b_n)$:

$$\pi_B[\![\varphi(\mathbf{b})]\!] = \begin{cases} 1, & \text{if } B = \{b_1, \dots, b_n\} \text{ and } b_i \neq b_j \text{ for } i \neq j, \\ 0, & \text{otherwise.} \end{cases}$$

Proof. The computation

$$\pi_B[\![\varphi(\mathbf{b})]\!] = \prod_{i<j} \pi_B[\![b_i \neq b_j]\!] \cdot \prod_{b \in B} \sum_{i \le n} \pi_B[\![b = b_i]\!]$$

evaluates to 1 if $b_1, \dots, b_n$ is a distinct enumeration of all elements of B, and to 0 otherwise. Q.E.D.

On the other side, the conjunction over the "true" literals $\psi(\mathbf{x})$ does not suffice to characterise a semiring interpretation up to isomorphism. There is a trivial example over universe $A = \{a\}$, $\tau = \{P, Q\}$ with model-defining $\mathbb{V}$-interpretations

$$\pi\colon Pa \mapsto 0.1 \qquad\qquad \pi\colon Qa \mapsto 0.9$$
$$\pi'\colon Pa \mapsto 0.9 \qquad\qquad \pi'\colon Qa \mapsto 0.1.$$

Clearly $\pi \not\cong \pi'$, but constructing $\chi_\pi, \chi_{\pi'}$ as above (using the induced model $\mathfrak{A}_\pi = \mathfrak{A}_{\pi'}$) would give

$$\chi_\pi = \exists x(\varphi(x) \land Px \land Qx) = \chi_{\pi'}.$$

Idea: repeat different "true" literals a different number of times, so that a different collection of values for the literals guarantees to give a different product. We associate with every finite $\mathbb{V}$-interpretation $\pi_A\colon \mathrm{Lit}_A(\tau) \to \mathbb{V}$ and every $\varepsilon > 0$ a characteristic sentence

$$\chi_{\pi_A,\varepsilon} := \exists x_1 \cdots \exists x_n(\varphi(\mathbf{x}) \land \psi_\varepsilon(\mathbf{x}))$$

with $n = |A|$ and $\varphi(\mathbf{x})$ as above, but with a more sophisticated construction of $\psi_\varepsilon(\mathbf{x})$.

Let $\mathbf{a} = (a_1, \dots, a_n)$ be a fixed enumeration of A and $\alpha_1(\mathbf{a}), \dots, \alpha_k(\mathbf{a})$ an enumeration of the "true" literals of π_A (with $\pi_A(\alpha(\mathbf{a})) \neq 0$). Further, fix a sequence $f(1), \dots, f(k) \in \mathbb{N}$ subject to the condition that

- $f(1) = 1$,
- $f(i+1)$ is large enough so that $(1-\varepsilon)^{f(i+1)} < \varepsilon^{f(1)+\cdots+f(i)}$. $(*)$

Then put

$$\psi_\varepsilon(\mathbf{x}) := \bigwedge_{i=1}^{k} \alpha_i(\mathbf{x})^{f(i)} \qquad \text{(where } \alpha^m := \underbrace{\alpha \land \alpha \land \cdots \land \alpha}_{m \text{ times}}\text{)}.$$

Proposition 7.7. Let $\pi_A \colon \mathrm{Lit}_A(\tau) \to \mathbb{V}$, $\pi_B \colon \mathrm{Lit}_B(\tau) \to \mathbb{V}$ be two finite model-defining $\mathbb{V}$-interpretations which induces the finite set of values

$$V := \{\pi_A(\alpha) \mid \alpha \in \mathrm{Lit}_A(\tau)\} \cup \{\pi_B(\alpha) \mid \alpha \in \mathrm{Lit}_B(\tau)\}.$$

Let $0 < \varepsilon \leq \min\{|r - s| \mid r, s \in V, r \neq s\}$. Then $\pi_A[\![\chi_{\pi_A,\varepsilon}]\!] = \pi_B[\![\chi_{\pi_B,\varepsilon}]\!]$ implies $\pi_A \cong \pi_B$.

Proof. Assume $\pi_A[\![\chi_{\pi_A,\varepsilon}]\!] = \pi_B[\![\chi_{\pi_B,\varepsilon}]\!]$. By construction, both values are > 0. By the lemma above, and since "$\exists$" is interpreted by max, we have $|A| = |B|$ and we further have enumerations $\mathbf{a} = (a_1, \ldots, a_n)$, $\mathbf{b} = (b_1, \ldots, b_n)$ of A and B such that

$$\pi_A[\![\psi_\varepsilon(\mathbf{a})]\!] = \pi_A[\![\chi_{\pi_A,\varepsilon}]\!] = \pi_B[\![\chi_{\pi_B,\varepsilon}]\!] = \pi_B[\![\psi_\varepsilon(\mathbf{b})]\!].$$

Observe that $\pi_A[\![\psi_\varepsilon(\mathbf{a})]\!] = \prod_{i=1}^k \pi_A(\alpha_i(\mathbf{a})^{f(i)}) > 0$ implies $\pi_A(\alpha_i(\mathbf{a})) > 0$ for all i, and analogously $\pi_B(\alpha_i(\mathbf{b})) > 0$ for all i. For $i = 1, \ldots, k$, let $r_i := \pi_A(\alpha_i(\mathbf{a})) > 0$ and $s_i := \pi_B(\alpha_i(\mathbf{b})) > 0$. It remains to show that $r_i = s_i$ for all $i \leq k$; then $\mathbf{a} \mapsto \mathbf{b}$ is indeed an isomorphism from π_A to π_B.

Assume that this is *not* the case. Let $j = \max\{i \leq k \mid r_i \neq s_i\}$. We can assume that $r_j < s_j$. Further $s_j - r_j \geq \varepsilon$ and $\varepsilon \leq r_i, s_i \leq 1$ for all i. Hence

$$r_j \leq s_j - \varepsilon \leq s_j - \varepsilon \cdot s_j = (1 - \varepsilon) s_j,$$

and

$$\begin{aligned} r_1^{f(1)} \cdots r_j^{f(j)} \leq r_j^{f(j)} &\leq (1-\varepsilon)^{f(j)} s_j^{f(j)} \\ &\overset{(*)}{<} \varepsilon^{f(1)+\cdots+f(j-1)} \cdot s_j^{f(j)} \leq s_1^{f(1)} \cdots s_j^{f(j)}. \end{aligned}$$

But $r_i = s_i$ for $i = j+1, \ldots, k$ so it follows that

$$\pi_A[\![\psi_\varepsilon(\mathbf{a})]\!] = \prod_{1 \leq i \leq k} r_i^{f(i)} \neq \prod_{1 \leq i \leq k} s_i^{f(i)} = \pi_B[\![\psi_\varepsilon(\mathbf{b})]\!],$$

and hence $\pi_A[\![\chi_{\pi_A,\varepsilon}]\!] \neq \pi_B[\![\chi_{\pi_B,\varepsilon}]\!]$, contradiction. Q.E.D.

While none of the sentences $\chi_{\pi_A,\varepsilon}$ alone can characterise π_A, the (of course countable) set $\Phi_{\pi_A} = \{\chi_{\pi_A,\varepsilon} \mid \varepsilon > 0\}$ provides an axiomatisation. Indeed: No infinite $\mathbb{V}$-interpretation π_B agrees with π_A on $\chi_{\pi_A,\varepsilon}$ (due to $\varphi(\mathbf{x})$), and for a finite $\mathbb{V}$-interpretation π_B we have

- if $\pi_A \cong \pi_B$ then $\pi_B[\![\chi_{\pi_A,\varepsilon}]\!] = \pi_A[\![\chi_{\pi_A,\varepsilon}]\!]$ for all ε,
- if $\pi_A \ncong \pi_B$ then we find an appropriate ε such that $\pi_A[\![\chi_{\pi_A,\varepsilon}]\!] \neq \pi_B[\![\chi_{\pi_A,\varepsilon}]\!]$ by Proposition 7.7.

One can show that for each finite $\mathbb{V}$-interpretation, also an axiomatisation by a *finite* set of sentences is possible, so also question (3) has a positive answer for $\mathbb{V}$. We finally show that, however, no axiomatisation by a *single* sentence is possible, in general.

Proposition 7.8. There exists a $\mathbb{V}$-interpretation $\pi\colon \mathrm{Lit}_A(\tau) \to \mathbb{V}$ such that for every sentence $\psi \in \mathrm{FO}(\tau)$, there exists a $\mathbb{V}$-interpretation $\pi'\colon \mathrm{Lit}_A(\tau) \to \mathbb{V}$ such that $\pi \ncong \pi'$, but $\pi[\![\psi]\!] = \pi'[\![\psi]\!]$.

Proof. Let $\pi\colon Pa \mapsto p$, $Qa \mapsto q$ with $0 < p, q < 1$ and p, q are multiplicatively independent: there are no $k, l \in \mathbb{Z} \setminus \{0\}$ with $p^k q^l = 1$.

Consider the corresponding $\mathbb{B}[x,y]$-interpretation $\pi_B\colon Pa \mapsto x$, $Qa \mapsto y$. For $\psi \in \mathrm{FO}(\tau)$, we have $\pi_B[\![\psi]\!] \in \mathbb{B}[x,y]$ and $h\colon x \mapsto p$, $y \mapsto q$ extends to $\hat{h}\colon \mathbb{B}[x,y] \to \mathbb{V}$ with $\hat{h}(\pi_B[\![\psi]\!]) = \pi[\![\psi]\!]$ (by the universal property of $\mathbb{B}[x,y]$ for idempotent semirings). $\pi_B[\![\psi]\!]$ is a sum of monomials of the form $m = x^i y^j$, and $\pi[\![\psi]\!] = p^i q^j$ is the *maximal* value $m(p,q)$ for the monomials m in $\pi_B[\![\psi]\!]$. Since p, q are multiplicatively independent, no other monomial can take the same value, i.e. $m'(p,q) < m(p,q)$ for all other monomials m' occurring in $\pi_B[\![\psi]\!]$.

We can certainly find another pair of values r, s with $r \neq p$ and $r^i s^j = p^i q^j$ such that r is sufficiently close to p that $m'(r,s) < m(r,s) = m(p,q)$ for all other monomials m' in $\pi_B[\![\psi]\!]$. For the $\mathbb{V}$-interpretation π' with $\pi'\colon Pa \mapsto r$, $Qa \mapsto s$ this implies $\pi'[\![\psi]\!] = r^i s^j = p^i q^j = \pi[\![\psi]\!]$, but $\pi \ncong \pi'$. Q.E.D.

8 Least Fixed Points and ω-Continuous Semirings

8.1 Scenarios involving Least Fixed Points

We discuss several scenarios that go beyond the semiring valuations considered so far.

8.1.1 *Games with Cycles*

Consider finite game graphs $\mathcal{G} = (V, V_0, V_1, T, E)$ as in Chapter 5, but drop the condition that $\mathcal{G}$ is acyclic. Given $f_\sigma\colon T \to S$, $h_\sigma\colon E \to S \setminus \{0\}$, we defined

$$f_\sigma(v) = \sum_{w \in vE} h_\sigma(vw) \cdot f_\sigma(w), \qquad \text{for } v \in V_\sigma,$$
$$f_\sigma(v) = \prod_{w \in vE} h_\sigma(vw) \cdot f_\sigma(w), \qquad \text{for } v \in V_{1-\sigma}.$$

Without the acyclicity condition this is no longer a simple backwards induction. Instead, putting $x_v := f_\sigma(v)$ we get, for $X = \{x_v \mid v \in V\}$, an equation system $X = G(X)$:

$$(*) \qquad \begin{aligned} x_t &= f_\sigma(t) && \text{for } t \in T,\\ x_v &= \sum_{w \in vE} h_\sigma(vw) \cdot x_w && \text{for } v \in V_\sigma,\\ x_v &= \prod_{w \in vE} h_\sigma(vw) \cdot x_w && \text{for } v \in V_{1-\sigma}. \end{aligned}$$

Under what condition can we compute a solution $G\colon X \to S$? The obvious idea is to use a least fixed-point induction. We define:

- $G^0(x_v) = 0$, for all v,
- $G^{i+1}(x_v) = \begin{cases} f_\sigma(x_v) & \text{for } v \in T, \\ \sum_{w \in vE} h_\sigma(vw) \cdot x_w & \text{for } v \in V_\sigma, \\ \prod_{w \in vE} h_\sigma(vw) \cdot x_w & \text{for } v \in V_{1-\sigma}. \end{cases}$

Question: Does the sequence $G^0, G^1, \ldots$ converge to a fixed-point solution $G^m = G^{m+1} =: G^\infty$?

8.1.2 Query languages with recursion: Datalog

Definition 8.1. A *Datalog program* Π is a finite set of *rules*

$$r: \quad H(\mathbf{x}) \leftarrow \alpha_1, \ldots, \alpha_m$$

where $H(\mathbf{x})$, the *head* of the rule, is an atomic formula, and $\alpha_1, \ldots, \alpha_m$, the *body* of the rule, is a collection of atomic formulae, containing all variables of the head predicate (and possibly more).

Example:

$$Txy \leftarrow Exy$$
$$Txy \leftarrow Txz, Tzy$$

We use the following notation:
σ: predicates that occur in the head of some rule (here: T),
τ: predicates that occur only in the bodies of the rules (here: E).

Given a database D of vocabulary τ, one recursively computes values for the head predicates and then produces a database $\Pi(D)$ of vocabulary $\tau \cup \sigma$: Write the body of a rule as a conjunction

$$r: \quad H(\mathbf{x}) \leftarrow \beta(\mathbf{x}, \mathbf{y}), \text{ where } \beta(\mathbf{x}, \mathbf{y}) = \alpha_1 \wedge \cdots \wedge \alpha_m.$$

Start with $\Pi(D) = (D, \emptyset)$, i.e. let the set of σ-predicates be empty. Whenever an instantiation $\beta(\mathbf{a}, \mathbf{b})$ of the body of a rule is true in $\Pi(D)$, add the fact $H(\mathbf{a})$ to $\Pi(D)$. This process terminates (after a polynomial number of steps w.r.t. $|D|$).

A (Boolean) Datalog query is given by a Datalog program Π and a head atom $H(\mathbf{a})$. It is true in a database D if the fact $H(\mathbf{a})$ holds in $\Pi(D)$; one often writes $D \cup \Pi \models H(\mathbf{a})$. In the example above, a database $D = (D, E)$ with finite $E \subseteq D \times D$ satisfies the query *Tab* if (a, b) is in the transitive closure of E. Thus Datalog can express queries that are not FO-definable.

On the other side, Datalog is limited by the absence of negation and universal quantification. In particular, Datalog queries are *monotone*: If $D \subseteq D'$ (viewed as sets of atomic facts), then every Datalog query true in D is also true in D'. There exist numerous extensions and variants of Datalog that add, for instance, negation in one form or another.

Question: How to define semiring valuations for Datalog?

8.1.3 Logic: Positive least-fixed point logic

Here we consider posLFP, a fragment of a much more general fixed-point logic (LFP), to be considered later. Given a formula $\psi(R, \mathbf{x})$ of vocabulary $\tau \cup \{R\}$ in which R occurs only positively and where $|\mathbf{x}| = k$ matches the arity of R, we build the formula

$$[\mathbf{lfp}\, R\mathbf{x}.\ \psi(R, \mathbf{x})](\mathbf{z}).$$

Given a τ-structure $\mathfrak{A}$, we associate with $\psi(R, \mathbf{x})$ an operator

$$\begin{aligned} F_\psi \colon \mathcal{P}(A^k) &\longrightarrow \mathcal{P}(A^k), \\ \overset{\Cup}{R} \quad &\longmapsto \{\mathbf{a} \in A^k \mid (\mathfrak{A}, R) \models \psi(R, \mathbf{a})\}. \end{aligned}$$

Due to the assumption that R only occurs positively in ψ, the operator F_ψ is monotone: if $R \subseteq R'$, then $F_\psi(R) \subseteq F_\psi(R')$. Therefore, starting with $R^0 = \emptyset$, the sequence $R^0, R^1, R^2, \ldots$ with $R^{i+1} := F_\psi(R^i)$ is increasing: $R^0 \subseteq R^1 \subseteq \ldots$ and therefore reaches a fixed-point R^α with $R^{\alpha+1} = F_\psi(R^\alpha) = R^\alpha =: R^\infty$. If $\mathfrak{A}$ is finite, this happens after a polynomial number of steps: $\alpha \leq |A|^k$. In fact, this inductively computed fixed point R^∞ is always the *least* fixed point of F_ψ, so

$$R^\infty = \mathbf{lfp}(F_\psi) = \bigcap \{R \subseteq A^k \mid F_\psi(R) = R\}$$

$$= \bigcap \{R \subseteq A^k \mid F_\psi(R) \subseteq R\}.$$

Semantics of lfp-formulae:

$$\mathfrak{A} \models [\mathbf{lfp}\, R\mathbf{x}.\ \psi(R, \mathbf{x})](\mathbf{a}) \quad \overset{\text{def}}{\iff} \quad \mathbf{a} \in \mathbf{lfp}(F_\psi).$$

Example 8.2. Let $\psi(T, xy) := Exy \lor \exists z(Txz \land Tzy)$. For a directed graph $G = (V, E)$, we have

$$\begin{aligned} &G \models [\mathbf{lfp}\, Txy.\ \psi(T, xy)]((a, b)) \\ \iff &(a, b) \in \mathrm{TC}(E) \\ \iff &\text{there is a path } a \xrightarrow{+} b \text{ in } G. \end{aligned}$$ ■

We define posLFP(τ) as the set of formulae built from τ-literals $Px, \neg Px$, $x = y$, $x \neq y$ and fixed-point atoms $R\mathbf{x}$, combined with $\land$, $\lor$, $\exists$, $\forall$, **lfp**. Some facts (more details in the upcoming lecture "Algorithmic Model Theory"):

- Let $\psi \in$ posLFP(τ). On a given finite τ-structure $\mathfrak{A}$ it can be decided in polynomial time (w.r.t. $|A|$) whether $\mathfrak{A} \models \psi$.
- Datalog $\leq$ posLFP: Every Datalog query can be translated into an equivalent posLFP-formula.

Example 8.3. As a further important example, consider reachability games $\mathcal{G} = (V, V_0, V_1, T, E)$ with a winning condition $W \subseteq T$ for Player 0. Player 0 has a winning strategy for $(\mathcal{G}, W)$ from $v \iff$ $(\mathcal{G}, W) \models [\mathbf{lfp}\, Rx.\ \varphi(R, x)](v)$, where

$$\begin{aligned} \varphi(R, x) := Wx &\lor (V_0 x \land \exists y(Exy \land Ry)) \\ &\lor (V_1 x \land \forall y(Exy \to Ry)). \end{aligned}$$

Indeed, for the fixed-point induction $R^0 \subseteq R^1 \subseteq \dots$ of F_φ on $(\mathcal{G}, W)$ we have that $R^i = \{v \mid \text{Pl. 0 can win from } v \text{ in } \leq i - 1 \text{ moves.}\}$. ■

Question: How to define semiring valuations for posLFP?

8.2 Towards Semirings for Least Fixed Points

Not all semirings are appropriate for these three scenarios. Consider the following game:

$$s \longleftarrow v \rightleftarrows \boxed{w} \longrightarrow t$$

8.2.1 Natural Numbers

Using the semiring $S = \mathbb{N}$, we set $h_0(e) = 1$ for all edges e, and $h_0(s) = a$, $h_0(t) = b$ for some $a, b \in \mathbb{N}$ (e.g. $a = 2$, $b = 5$). We get an equation system for x_v, x_w (and x_s, x_t):

$$\begin{aligned} x_s &= a, \\ x_t &= b, \\ x_v &= x_w + a, \\ x_w &= x_v \cdot b. \end{aligned}$$

This system has no solution in $\mathbb{N}$. The least fixed-point induction, starting with $G^0(x_v) = G^1(x_w) = 0$ leads to $G^i(x_s) = a$, $G^i(x_t) = b$ for all i, and further

	x_v	x_w
G^0	0	0
G^1	a	0
G^2	a	ab
G^3	$a + ab$	ab
G^4		$ab + ab^2$
G^5	$a + ab + ab^2$	
G^6		$ab + ab^2 + ab^3$
	$\vdots$	$\vdots$
G^{2n-1}	$a + ab + \cdots + ab^{n-1}$	
G^{2n}		$ab + ab^2 + \cdots + ab^n$
	$\vdots$	$\vdots$

But we do get a solution in the extended semiring $\mathbb{N}^\infty$ over $\mathbb{N} \cup \{\infty\}$ with operations extended as follows:

$$\infty + n = \infty \quad \text{and} \quad \infty \cdot n = \begin{cases} \infty, & n \neq 0, \\ 0, & n = 0. \end{cases}$$

The solution gives:

$$f_0(v) = G^\infty(v) = \begin{cases} 0, & a = 0, \\ a, & b = 0, \\ \infty, & a, b \neq 0, \end{cases}$$

$$f_0(w) = G^\infty(w) = \begin{cases} 0, & a = 0 \text{ or } b = 0, \\ \infty, & a, b \neq 0. \end{cases}$$

8.2.2 Viterbi Semiring

We also get solutions in many other semirings, for instance in the Viterbi semiring:

$$f_0(v) = G^\infty(v) = a, \qquad f_0(w) = G^\infty(w) = ab.$$

Indeed:

$$\begin{aligned} f_0(v) &= f_0(w) + a && \text{since } a = \max(ab, a), \\ f_0(w) &= f_0(v) \cdot b && \text{since } ab = a \cdot b. \end{aligned}$$

8.2.3 Formal Power Series

What about the most general semirings $\mathbb{N}[X]$, here $\mathbb{N}[s, t]$? The fixed-point induction gives polynomials

$$\begin{aligned} F^{2n+1}(v) &= s + st + \cdots + st^n, \\ F^{2n+1}(w) &= st + st^2 + \cdots + st^n. \end{aligned}$$

This does not converge to a solution in $\mathbb{N}[s,t]$, but it does converge in $\mathbb{N}^\infty[\![s,t]\!]$, the semiring of formal power series (i.e., possibly infinite sums of monomials) over s,t. The solution is given by the infinite sums

$$f_0(v) = s + st + st^2 + \dots,$$
$$f_0(w) = st + st^2 + \dots.$$

Indeed, $f_0(v) = s + f_0(w)$ and $f_0(w) = f_0(v) \cdot t$. And this valuation gives us the desired information about the strategies of Player 0 from v and w:

- From v, Player 0 has strategies $\mathcal{S}_n$ for each $n \in \mathbb{N}$: move n times to w, then to s. The monomial associated with $\mathcal{S}_n$ is $s \cdot t^n$ (one play to s, n plays to t). Thus $F(\mathcal{S}_n) = s \cdot t^n$ and $f_0(v) = \sum_{n\in\mathbb{N}} F(\mathcal{S}_n)$.
- From w this is analogous, but for $\mathcal{S}'_n$ (move n times from v to w, then to s) the valuation is $F(\mathcal{S}'_n) = s \cdot t^{n+1}$.

8.3 ω-continuous Semirings

Definition 8.4. A semiring S is *ω-continuous* if it is

- naturally ordered ($a \leq b \overset{\text{def}}{\iff} \exists c\ a + c = b$ is a partial order);
- ω-complete: every ascending ω-chain $C = (a_n)_{n<\omega}$ with $a_n \leq a_{n+1}$, for all n, has a supremum in S (denoted $\bigsqcup C$);
- and its operations $+$ and $\cdot$ are ω-continuous in both arguments (see below), i.e., $a + \bigsqcup C = \bigsqcup(a + C)$ and $a \cdot \bigsqcup C = \bigsqcup(a \cdot C)$ for all ascending ω-chains C.

Definition 8.5. A function $f\colon S \to S$ on ω-complete semirings is *ω-continuous*, if $f(\bigsqcup C) = \bigsqcup f(C)$ for each ascending ω-chain $C = (a_n)_{n\in\omega}$ with supremum $\bigsqcup C$.

Notice that an ω-continuous function f is in particular monotone ($a \leq b \implies f(a) \leq f(b)$). To see this, consider the chain $\{a,b\}$ with $a \leq b$. Then $\bigsqcup\{f(a), f(b)\} = f(\bigsqcup\{a,b\}) = f(b)$, hence $f(a) \leq f(b)$. (For $+$ and $\cdot$, we already know that they are monotone on naturally ordered semirings.)

Theorem 8.6 (Kleene)**.** If S is ω-complete and $f\colon S \to S$ is ω-continuous, then $\mathbf{lfp}(f) = \bigsqcup\{f^n(0) \mid n < \omega\}$.

Proof. As f is monotone, the iteration $0, f(0), f^2(0), \dots, f^n(0), \dots$ is an ascending ω-chain with a supremum $z \in S$. By the ω-continuity of f:

$$f(z) = f\left(\bigsqcup\{f^n(0) \mid n < \omega\}\right) = \bigsqcup\left\{f^{n+1}(0) \,\middle|\, n < \omega\right\} = z.$$

For each other fixed point $x = f(x)$, we have that $f^n(0) \le x$ for all n, since clearly $f^0(0) = 0 \le x$ and if $f^n(0) \le x$ then also $f^{n+1}(0) = f(f^n(0)) \le f(x) = x$. Thus also $z = \bigsqcup\{f^n(0) \mid n < \omega\} \le x$, so $z = \mathbf{lfp}(f)$. Q.E.D.

This readily generalizes to systems $(f_1, \dots, f_n)$ with $f_i\colon S^n \to S$ (see below). In an ω-continuous semiring, we also have a well-defined infinite summation operator

$$\sum_{i \le \omega} b_i := \bigsqcup\{\sum_{i \le n} b_i \mid n < \omega\}$$

and the Kleene star operation

$$a^* := \sum_{i < \omega} a^i = \bigsqcup\{(1 + a + a^2 + \dots + a^i \mid i < \omega\}.$$

Definition 8.7. Given a semiring S and a finite set X of indeterminates, we define the semiring $S[\![X]\!]$ of *formal power series* (possibly infinite sums of monomials) with coefficients in S and indeterminates in X. Addition and multiplication are defined in the obvious way.

If S is ω-continuous, then so is $S[\![X]\!]$. Further, for $n = |X|$, each $f \in S[\![X]\!]$ induces a function

$$f\colon S^n \to S, \quad (a_1, \dots, a_n) \mapsto f(a_1, \dots, a_n)$$

which is ω-continuous in each argument.

We can deal with negation in the same way as in polynomial semirings, using dual-indeterminate power series: $\mathbb{N}[\![X, \overline{X}]\!]$ is the quotient of $\mathbb{N}[\![X \cup \overline{X}]\!]$ by the congruence generated by $x \cdot \overline{x} = 0$ for $x \in X$. This semiring has the following universal property:

Proposition 8.8 (Universal Property)**.** Every function $f\colon X \cup \overline{X} \to S$ into an *ω-continuous* semiring S with $f(x) \cdot f(\overline{x}) = 0$ for all $x \in X$ extends uniquely to an *ω-continuous* semiring homomorphism $\hat{f}\colon \mathbb{N}[\![X, \overline{X}]\!] \to S$ that coincides with f on $X \cup \overline{X}$.

Let S be ω-continuous. A system of power series (or polynomials) with indeterminates $X = (X_1, \ldots, X_n)$ is a sequence $F = (f_1, \ldots, f_n)$ with $f_i \in S[\![X]\!]$. It induces a function $F\colon S^n \to S^n$ that is ω-continuous in each argument. By Kleene's Fixed-Point Theorem it has a least fixed point $\mathbf{lfp}(F) \in S^n$ which coincides with the supremum $\bigsqcup_{i<\omega} F^i$ of the Kleene approximants with $F^0 = 0$ and $F^{i+1} = F(F^i)$. We also refer to $\mathbf{lfp}(F)$ as the least fixed-point *solution* of $X = F(X)$.

8.3.1 Reachability Games

Consider now a reachability game $\mathcal{G}$ (potentially with cycles), valuations $f_\sigma\colon T \to S$ and $h_\sigma\colon E \to S \setminus \{0\}$ into an ω-continuous semiring S, and the associated equation system $G_\sigma(X) = X$ for $X = \{x_v \mid v \in V\}$. The least fixed-point solution $\mathbf{lfp}(G_\sigma)$ gives us the desired valuation $f_\sigma\colon V \to S$ with $f_\sigma(v) = (\mathbf{lfp}\, G_\sigma)(x_v)$. Does the Sum-of-Strategies Theorem hold also in this case? To generalise this result to games $\mathcal{G}$ with cycles, we need to extend valuations of plays and strategies to such games.

Plays:
For a finite play $x = v_0 \ldots v_m$ ending in a terminal node $v_m \in T$, we put $f_\sigma(x) = h_\sigma(v_0v_1) \cdots h_\sigma(v_{m-1}v_m) f_\sigma(v_m)$ as before. For an infinite play x we put $f_\sigma(x) = 0$.

Strategies:
For a strategy $\mathcal{S} \in \mathrm{Strat}_\sigma(v)$ we put $F(\mathcal{S}) = 0$ if $\mathcal{S}$ admits any infinite play. Hence a strategy $\mathcal{S}$ can have a non-zero valuation only if it admits only finite plays. By Kőnig's Lemma it then only admits a finite number

of plays, and putting

$$F(\mathcal{S}) = \prod_{e \in E} h_\sigma(e)^{\#_e(\mathcal{S})} \cdot \prod_{v \in T} f_\sigma(t)^{\#_t(\mathcal{S})}$$

is well-defined, as $\#_e(\mathcal{S})$ and $\#_t(\mathcal{S})$ are always finite.

Although the *number* of strategies in $\mathrm{Strat}_\sigma(v)$ may well be infinite, the Sum-of-Strategies Theorem generalizes to reachability games with cycles. The proof relies on Kleene's Fixed-Point Theorem and the unravelings of $\mathcal{G}$ to finite truncated acyclic games $\mathcal{G}^{(n)}$.

Definition 8.9. Given $\mathcal{G} = (V, V_0, V_1, T, E)$ with basic valuations $f_\sigma \colon T \to S$ and $h_\sigma \colon E \to S \setminus \{0\}$, the *truncation* $\mathcal{G}^{(n)} = (V^{(n)}, V_0^{(n)}, V_1^{(n)}, T^{(n)}, E^{(n)})$ for $n > 0$ is the restriction of the forest $\bigcup_{v \in V} \mathcal{T}(\mathcal{G}, v)$ to paths of less than n moves, and $\rho^{(n)} \colon \mathcal{G}^{(n)} \to \mathcal{G}$ is the restriction of the canonical homomorphisms ρ to $\mathcal{G}^{(n)}$.

The basic valuations for $\mathcal{G}^{(n)}$ are defined by

$$\begin{aligned} h^n_\sigma \colon E^{(n)} &\longrightarrow S \setminus \{0\}, & f^n_\sigma \colon T^{(n)} &\longrightarrow S, \\ e &\longmapsto h_\sigma(\rho^{(n)}(e)), & \pi v &\longmapsto \begin{cases} f_\sigma(v) & \text{if } v \in T \\ 0 & \text{if } v \notin T, \end{cases} \end{aligned}$$

i.e. if $\pi v \in T^{(n)}$ is an initial segment with $n-1$ moves of a play in $\mathcal{G}$ that has not reached a terminal position.

The games $\mathcal{G}^{(n)}$ are finite acyclic games and the basic valuations h^n_σ, f^n_σ extend to valuations $f^n_\sigma \colon V \to S$. Let $G^n \colon V \to S$ be the Kleene approximants of the equation system $G_\sigma(X) = X$ of the game $\mathcal{G}$. By an easy induction, we get the

Lemma 8.10. For all $n > 0$, we have $f^n_\sigma = G^n$.

Let $\mathrm{Strat}^{(n)}_\sigma(v)$ be the set of strategies of Player σ from v in $\mathcal{G}^{(n)}$. Since the games $\mathcal{G}^{(n)}$ are acyclic, the *Sum-of-Strategies Theorem* from Chapter 5 holds:

$$f^n_\sigma(v) = \sum_{\mathcal{T} \in \mathrm{Strat}^{(n)}_\sigma(v)} F(\mathcal{T})$$

Every strategy $\mathcal{S} \in \mathrm{Strat}_\sigma(v)$ of $\mathcal{G}$ induces strategies $\mathcal{S}^{(n)} \in \mathrm{Strat}_\sigma^{(n)}(v)$ for $\mathcal{G}^{(n)}$ (for all n).

Lemma 8.11. For every strategy $\mathcal{S} \in \mathrm{Strat}_\sigma(v)$ in $\mathcal{G}$ with $F(\mathcal{S}) \neq 0$ there exists some $n_\mathcal{S} < \omega$ with

- $\mathcal{S} = \mathcal{S}^{(n)}$ for all $n \geq n_\mathcal{S}$,
- $F(\mathcal{S}^{(m)}) = 0$ for all $m < n_\mathcal{S}$.

Proof. If $F(\mathcal{S}) \neq 0$ then $\mathrm{Plays}(\mathcal{S})$ is finite and has only finite plays. Let $n_\mathcal{S} = \max\{|x| : x \in \mathrm{Plays}(\mathcal{S})\}$. For $n \geq n_\mathcal{S}$, all $x \in \mathrm{Plays}(\mathcal{S})$ are also contained in $\mathrm{Plays}(\mathcal{S}^{(n)})$. For $m < n_\mathcal{S}$, $\mathrm{Plays}(\mathcal{S}^{(m)})$ contains an unfinished play, whose last position has value $f_\sigma^m(w) = 0$, hence $F(\mathcal{S}^{(m)}) = 0$. Q.E.D.

Every strategy $\mathcal{T} \in \mathrm{Strat}_\sigma^{(n)}(v)$ is induced by some $\mathcal{S} \in \mathrm{Strat}_\sigma(v)$, so that $\mathcal{T} = \mathcal{S}^{(n)}$. In general $\mathcal{S}$ is not uniquely determined by $\mathcal{T}$ and n. Nevertheless:

Lemma 8.12. For all positions v of $\mathcal{G}$ and all n,

$$\sum_{\mathcal{S} \in \mathrm{Strat}_\sigma(v)} F(\mathcal{S}^{(n)}) = \sum_{\mathcal{T} \in \mathrm{Strat}_\sigma^{(n)}(v)} F(\mathcal{T}).$$

Proof. If $\mathcal{S}_1 \neq \mathcal{S}_2 \in \mathrm{Strat}_\sigma(v)$ with $\mathcal{T} = \mathcal{S}_1^{(n)} = \mathcal{S}_2^{(n)}$, then $\mathrm{Plays}(\mathcal{T})$ contains an unfinished play (otherwise $\mathcal{T} = \mathcal{S}_1 = \mathcal{S}_2$) which implies $F(\mathcal{T}) = 0$. Thus, although the strategy spaces $\mathrm{Strat}_\sigma(v)$ may be infinite, whereas $\mathrm{Strat}_\sigma^{(n)}(v)$ is finite for each n, those strategies that provide non-zero valuations are in one-to-one correspondence, and the two sums have the same values. Q.E.D.

Theorem 8.13 (Sum of Strategies)**.** For every finite game graph $\mathcal{G}$ with basic valuations $f_\sigma \colon T \to S$ and $h_\sigma \colon E \to S \setminus \{0\}$ into an ω-continuous semiring $\mathcal{S}$,

$$f_\sigma(v) = (\mathbf{lfp}\, G_\sigma)(v) = \sum_{\mathcal{S} \in \mathrm{Strat}_\sigma(v)} F(\mathcal{S}).$$

Proof. For every $n < \omega$,

$$G^n(v) = f_\sigma^n(v) = \sum_{\mathcal{T} \in \mathrm{Strat}_\sigma^{(n)}(v)} F(\mathcal{T}) = \sum_{\mathcal{S} \in \mathrm{Strat}_\sigma(v)} F(\mathcal{S}^{(n)}).$$

Further, for every $\mathcal{S} \in \mathrm{Strat}_\sigma(v)$, we have $F(\mathcal{S}) = F(\mathcal{S}^{(n)})$ for sufficiently large n. Take suprema on both sides. Q.E.D.

For basic valuations $f_\sigma \colon V \to \mathbb{N}^\infty[\![T]\!]$ with $f_\sigma(t) = t$ and $h_\sigma(e) = 1$ for all $e \in E$, the value $f_\sigma(v) \in \mathbb{N}^\infty[\![T]\!]$ is an infinite sum of monomials $m \cdot t_1^{j_1} \cdots t_k^{j_k}$ with $m \in \mathbb{N}^\infty$ and $j_1, \ldots, j_k > 0$. Each such monomial indicates that Player σ has m strategies $\mathcal{S}$ from v with set of outcomes $\{t_1, \ldots, t_k\}$, and precisely j_i plays consistent with $\mathcal{S}$ have outcome t_i.

8.3.2 *Valuations for* posLFP

Let $\pi \colon \mathrm{Lit}_A(\tau) \to S$ be a semiring intepretation into an ω-continuous semiring S, which provides valuations $\pi[\![\psi]\!] \in S$ for $\psi \in \mathrm{FO}(\tau)$. We extend this to posLFP as follows:

Consider $\psi(\mathbf{a}) = [\mathbf{lfp}\, R\mathbf{x}.\ \varphi(R, \mathbf{x})](\mathbf{a})$. Assume that a valuation for φ are already defined; if R has arity m, this gives us a function $g \colon A^m \to S$. For two such functions g, g' we say that $g \le g'$ if $g(\mathbf{a}) \le g'(\mathbf{a})$ for all $\mathbf{a} \in A^m$. Given $\pi \colon \mathrm{Lit}_A(\tau) \to S$, and $g \colon A^m \to S$, let

$$\pi[R \mapsto g] \colon\ \mathrm{Lit}_A(\tau) \cup \mathrm{Atoms}_A(\{R\}) \longrightarrow S$$

be obtained from π by adding valuations $R\mathbf{c} \mapsto g(\mathbf{c})$. This provides, for each $\mathbf{a} \in A^m$ a valuation $\pi[R \mapsto g][\![\varphi(R, \mathbf{a})]\!] \in S$ and thus an update operator F_π^φ on functions $g \colon A^m \to S$:

$$F_\pi^\varphi \colon g \mapsto F_\pi^\varphi(g), \text{ where } F_\pi^\varphi(g) \colon \mathbf{a} \mapsto \pi[R \mapsto g][\![\varphi(R, \mathbf{a})]\!].$$

This operator is monotone; by Kleene's Fixed-Point Theorem, it has a least fixed point $\mathbf{lfp}(F_\pi^\varphi) \colon A^m \to S$ which coincides with the supremum of $(g^n)_{n<\omega}$ with $g^0 = 0$ and $g^{n+1} := F_\pi^\varphi(g^n)$, and which we define as the S-valuation of ψ, i.e.,

$$\pi[\![\psi(\mathbf{a})]\!] := \mathbf{lfp}(F_\pi^\varphi)(\mathbf{a}).$$

8.3.3 Definition via Model-Checking Game

Alternatively, we can define the S-valuation via model-checking games. Extend model-checking games $\mathcal{G}(A,\psi)$ from FO to posLFP by the following moves, for formulae $[\mathbf{lfp}\, R\mathbf{x}.\ \varphi(R,\mathbf{x})](\mathbf{a})$, assuming that each fixed-point variable R is used only once by an **lfp**-operator, so that φ is uniquely determined by R:

$$\begin{aligned}[][\mathbf{lfp}\, R\mathbf{x}.\ \varphi(R,\mathbf{x})](\mathbf{a}) &\longmapsto \varphi(R,\mathbf{a}),\\ R\mathbf{b} &\longmapsto \varphi(R,\mathbf{b}).\end{aligned}$$

Notice that $\mathcal{G}(A,\psi)$ may have cycles, but the terminal positions are just the τ-literals $P\mathbf{c}, \neg P\mathbf{c}$ and $c = d, c \neq d$. A valuation $\pi\colon \mathrm{Lit}_A(\tau) \to S$ provides a valuation f_0 of the terminal positions of $\mathcal{G}(A,\psi)$; as for FO-model-checking games, we evaluate edges trivially. We get a valuation $f_0\colon V \to S$ of all positions of $\mathcal{G}(A,\psi)$. But these positions are instantiated subformulae $\varphi(\mathbf{a})$ of ψ. In particular, $f_0(\psi(\mathbf{a})) \in S$ gives us a semiring valuation for $\psi(\mathbf{a}) \in$ posLFP. These two definitions coincide:

$$\pi[\![\psi(\mathbf{a})]\!] = f_0(\psi(\mathbf{a})).$$

8.3.4 Datalog

Valuations of Datalog queries are defined in a similar way. Consider a Datalog program Π with head vocabulary σ and body predicates from τ, and an ω-continuous semiring S. A τ-database over S with domain A is given by a function $D\colon \mathrm{Atoms}_A(\tau) \to S$. The Datalog program Π extends this to a function $(D,\Pi)\colon \mathrm{Atoms}_A(\tau\cup\sigma) \to S$ which is the least fixed point of an update operator F_π^D on functions $g\colon \mathrm{Atoms}_A(\tau\cup\sigma) \to S$.

Given a rule $r\colon H(\mathbf{x}) \leftarrow \gamma_1(E_1),\dots,\gamma_m(E_m)$ with $E_1\cup\dots\cup E_m = \mathbf{x}\cup\mathbf{y}$, an instantiation $r(\mathbf{a},\mathbf{b})$ of r is obtained by a map $x\mapsto \mathbf{a}, y \mapsto \mathbf{b}$. Put

$$\begin{aligned}\mathrm{head}(r(\mathbf{ab})) &= H\mathbf{a},\\ \mathrm{body}(r(\mathbf{ab})) &= \{\gamma_1(\mathbf{ab}),\dots,\gamma_m(\mathbf{ab})\}.\end{aligned}$$

The update function $F_\Pi^D \colon g \mapsto g'$ for $g \colon \mathrm{Atoms}_A(\tau \cup \sigma) \to S$ is given by the equation system

$$\begin{aligned} g(x_\alpha) &= D(\alpha), \quad \text{for } \alpha \in \mathrm{Atoms}_A(\tau), \\ g(x_\alpha) &= \sum_{\substack{r(\mathbf{ab}) \in \Pi \\ \text{with head}(r(\mathbf{ab}))=\alpha}} \prod_{\gamma \in \mathrm{body}(r(\mathbf{ab}))} g(x_\gamma). \end{aligned}$$

The update operator is monotone, and we put $(D, \Pi)(\alpha) = \big(\mathbf{lfp}\, F_\Pi^D\big)(x_\alpha)$ for every $\alpha \in \mathrm{Atoms}_A(\sigma \cup \tau)$.

We can also view this as a reachability game with positions

$$\begin{aligned} &\underbrace{\{x_\alpha \mid \alpha \in \mathrm{Atoms}_A(\sigma)\}}_{V_0} \quad \cup \quad \underbrace{\{x_\alpha \mid \alpha \in \mathrm{Atoms}_A(\tau)\}}_{T} \\ \cup \quad &\underbrace{\{r(\mathbf{a}, \mathbf{b}) \mid r \text{ a rule of } \Pi,\ \mathbf{a} \in A^i,\ \mathbf{b} \in A^j \text{ for appropriate } i, j\}}_{V_1}. \end{aligned}$$

Player 0 moves from $x_\alpha \in V_0$ to some $r(\mathbf{a}, \mathbf{b})$ with $\mathrm{head}(r(\mathbf{a}, \mathbf{b})) = \alpha$ and from there Player 1 moves to some x_γ with $\gamma \in \mathrm{body}(r(\mathbf{a}, \mathbf{b}))$. An S-database $D \colon \mathrm{Atoms}_A(\tau) \to S$ provides a valuations of Player 0 for the positions in T. The induced valuation $f_0 \colon V_0 \to S$ provides values $\big(\Pi, D\big)(\alpha) = f_0(x_\alpha)$. Again the valuations coincide.

9 Greatest Fixed Points, Fully-Continuous Semirings and Generalized Absorptive Polynomials

9.1 Semirings for LFP

The full LFP-logic is based on both least and greatest fixed points (or, equivalently, the interleaving of least fixed points and negation). This means that for formulae $\psi(R, \mathbf{x})$, where R occurs only positively in ψ, we can build not only $[\mathbf{lfp}\, R\mathbf{x}.\ \psi(R, \mathbf{x})](\mathbf{z})$, but also $[\mathbf{gfp}\, R\mathbf{x}.\ \psi(R, \mathbf{x})](\mathbf{z})$ with

$$\mathfrak{A} \models \psi(\mathbf{a}) \quad \stackrel{\text{def}}{\Longleftrightarrow} \quad \mathbf{a} \in \mathbf{gfp}(F_\psi),$$

where $\mathbf{gfp}(F)$ is the greatest fixed point of a monotone operator F.

Consider a powerset lattice $\mathcal{P}(B)$. There is a duality between least and greatest fixed points: For $X \in \mathcal{P}(B)$, let $\overline{X} := B \setminus X$, and let the dual operator of $F\colon \mathcal{P}(B) \to \mathcal{P}(B)$ be $F^d\colon \mathcal{P}(B) \to \mathcal{P}(B)$ with $F^d(X) := \overline{F(\overline{X})}$. If F is monotone, then so is F^d, and we have $\mathbf{lfp}(F^d) = \overline{\mathbf{gfp}(F)}$, and $\mathbf{gfp}(F^d) = \overline{\mathbf{lfp}(F)}$. In terms of LFP-formulae, this means that

$$\begin{aligned}
\neg\big[\mathbf{lfp} R\,\mathbf{x}.\ \psi(R,\mathbf{x})\big](\mathbf{z}) &\equiv \big[\mathbf{gfp} R\,\mathbf{x}.\ \neg\psi(R,\mathbf{x})[R/\neg R]\big](\mathbf{z}),\\
\neg\big[\mathbf{gfp} R\,\mathbf{x}.\ \psi(R,\mathbf{x})\big](\mathbf{z}) &\equiv \big[\mathbf{lfp} R\,\mathbf{x}.\ \neg\psi(R,\mathbf{x})\,\underbrace{[R/\neg R]}_{\text{replace atoms } R\mathbf{a} \text{ by } \neg R\mathbf{a}}\big](\mathbf{z}).
\end{aligned}$$

For a monotone operator $F\colon \mathcal{P}(B) \to \mathcal{P}(B)$, the greatest fixed point $\mathbf{gfp}(F)$ can be computed by dual induction (compared to $\mathbf{lfp}(F)$):

$$Y^0 := B, \quad Y^{\alpha+1} := F(Y^\alpha), \quad Y^\lambda := \bigcap_{\alpha<\lambda} Y^\alpha,$$

for ordinals α and limit ordinals λ. This produces a decreasing sequence

$$Y^0 \supseteq Y^1 \supseteq \dots Y^\alpha \supseteq Y^{\alpha+1} \supseteq \dots Y^\infty = F(Y^\infty) = \mathbf{gfp}(F).$$

Example 9.1. Given a game $\mathcal{G} = (V, V_0, V_1, T, E)$, a *safety* condition for Player 0 is given by a set $S \subseteq V$ of safe positions; Player 0 wins those plays that never leave S. The winning region of $(\mathcal{G}, S)$ for Player 0 is defined by

$$\begin{aligned}\psi(x) := \big[\mathbf{gfp}\, Wx.\ Sx \wedge (V_0 x \to \exists y(Exy \wedge Wy))\\ \wedge (V_1 x \to \forall y(Exy \to Wy))\big](x).\end{aligned}$$ ∎

To define semiring valuations for posLFP-formulae, we considered for $\varphi(R, \mathbf{x})$ and $\pi\colon \mathrm{Lit}_A(\tau) \to S$ the monotone update operator F_π^φ on functions $g\colon A^m \to S$ with $F_\pi^\varphi(g)\colon \mathbf{a} \mapsto \pi[R \mapsto g][\![\varphi(R, \mathbf{a})]\!]$ and then put $\pi[\![[\mathbf{lfp}\, R\mathbf{x}.\ \varphi(R, \mathbf{x})](\mathbf{a})]\!] := \mathbf{lfp}(F_\pi^\varphi)(\mathbf{a})$. Can we use the same idea for greatest fixed points? What properties of semirings are necessary to guarantee that that greatest fixed points $\mathbf{gfp}(F_\pi^\varphi)$ are *well-defined* and *informative* on S?

Given a naturally ordered semiring S, a *chain* $C \subseteq S$ is any totally ordered subset.

Definition 9.2. A naturally ordered semiring S is *fully chain-complete* if every chain $C \subseteq S$ has supremum $\bigsqcup C$ and an infimum $\bigsqcap C$. Moreover, S is *fully continuous* if in addition, its operations $\circ \in \{+, \cdot\}$ are fully continuous, i.e., $a \circ \bigsqcup C = \bigsqcup(a \circ C)$ and $a \circ \bigsqcap C = \bigsqcap(a \circ C)$ for all $a \in S$ and all non-empty chains $C \subseteq S$.

Example 9.3. The semirings $\mathbb{V}$, $\mathbb{N}^\infty$, $\mathbb{N}^\infty[\![X]\!]$, $\mathbb{N}^\infty[X, \overline{X}]$ are fully continuous. ∎

By a straightforward generalisation of Kleene's Theorem, we obtain that every monotone function $f\colon S \to S$ on a fully chain-complete semiring has least and greatest fixed points $\mathbf{lfp}(f)$ and $\mathbf{gfp}(f)$. Hence semiring semantics for LFP is well-defined in fully chain-complete semirings (one has to show that the update operators F_π^φ are always monotone, even for nested fixed points).

Example 9.4. Existence of an infinite path from u:

$$\mathcal{G}: \quad \begin{tikzcd} \underset{u}{\bullet} \arrow[r] & \underset{v}{\bullet} \arrow[loop right] \end{tikzcd} \qquad\qquad \psi(u) := [\mathbf{gfp}\, Rx.\ \exists y(Exy \wedge Ry)](u).$$

The fixed-point induction gives $R^0 = \{u, v\}$, $R^1 = \{u, v\} = R^\infty$. Indeed, an infinite path exists from both u and v. What happens if we evaluate ψ in different semirings?

$\mathbb{B}$:

There is a unique $\mathbb{B}$-interpretation π that defines $\mathcal{G}$, and $\pi[\![\psi(u)]\!] = 1$.

$\mathbb{V}$:

For π with $\pi(Euv) = \pi(Evv) = 1$, we have $\pi[\![\psi(u)]\!] = 1$ as above. But if we set $\pi(Evv) = 1 - \varepsilon$ (for $\varepsilon > 0$), we obtain $\pi[\![\psi(u)]\!] = 0$ due to the fixed-point iteration $1, 1 - \varepsilon, (1 - \varepsilon)^2, \ldots$. So while $\mathfrak{A}_\pi = \mathcal{G}$, we have $\pi[\![\psi(u)]\!] = 0$ although $\mathcal{G} \models \psi(u)$.

The Viterbi semiring is thus *not* truth-preserving (but since the loop at v must occur infinitely often in an infinite path from u, the value $\pi[\![\psi(u)]\!]$ still makes sense as a confidence score).

$\mathbb{N}^\infty[\![X]\!]$:

Setting $\pi(Euv) = x$, $\pi(Evv) = y$, we get $\pi[\![\psi(u)]\!] = 0$ due to the fixed-point iteration $\top$, $y \cdot \top$, $y^2 \cdot \top$, $\ldots$ with infimum 0. (Here, $\top$ is the power series where all monomials have coefficient ∞, whereas $y^m \cdot \top$ contains only those monomials where y has an exponent $\geq m$.)

Thus $\mathbb{N}^\infty[\![X]\!]$ is not truth-preserving either.

$\mathbb{N}^\infty$:

Setting $\pi(Euv) = \pi(Evv) = 1$ gives $\pi[\![\psi(u)]\!] = \infty$. The fixed-point iteration is ∞, $1 \cdot \infty$, $\ldots$ stagnating immediately.

Problem: Multiplication with non-zero values is increasing. Greatest fixed-point iterations will almost always stay at infinity and do not provide informative values.

Notice that we cannot get the value in $\mathbb{N}^\infty$ from the computation in $\mathbb{N}^\infty[\![X]\!]$ by evaluation of the power series, so $\mathbb{N}^\infty[\![X]\!]$ are not the *right* universal provenance semirings for full LFP. ∎

The problem that in some semirings, the valuations of greatest fixed points are, although well-defined, not really informative and do not provide useful insights why a formula holds, can be tied to two separate problems:

- lack of symmetry between least and greatest fixed points in some semirings,
- semirings that are not truth-preserving (they may evaluate true statements to 0).

To deal with these problems, we work with semirings that are *absorptive* and *chain-positive*.

Absorptive semirings provide more symmetry: multiplication is decreasing (whereas addition is increasing). This avoids the problems (as in $\mathbb{N}^\infty$) that the **gfp**-induction remains stuck at the top element. Further, absorptive semirings give information about "reduced" proof and evaluation strategies.

We say that a semiring S is truth-preserving for a logic L if for every model-defining S-interpretation π we have that $\mathfrak{A}_\pi \models \varphi \Longleftrightarrow \pi[\![\varphi]\!] \neq 0$ for all $\varphi \in L$.

Definition 9.5. A fully chain-complete semiring S is *chain-positive* if for every non-empty chain $C \subseteq S$ of non-zero elements, $\bigsqcap C$ is non-zero as well.

Lemma 9.6. Every chain-positive, positive semiring is truth-preserving for LFP.

9.2 Generalised Absorptive Polynomials

We introduce the semirings $\mathbb{S}^\infty[X]$ (and $\mathbb{S}^\infty[X, \overline{X}]$). Let X be a finite set of indeterminates. We generalise the notion of a monomial to $m\colon X \to \mathbb{N}^\infty$ with $m = x_1^{m(x_1)} \cdots x_n^{m(x_n)}$. Multiplication of monomials adds exponents, and $x^n \cdot x^\infty = x^\infty$. We say that m_2 absorbs m_1 ($m_2 \succcurlyeq m_1$) if $m_2(x) \leq m_1(x)$ for all $x \in X$. The set of monomials is of course infinite, but

- every antichain of monomials is finite,
- while there are infinitely descending chains of monomials, such as $1 = x^0 \succ x^1 \succ x^2 \succ \dots$ (with infimum x^∞), there are no infinitely ascending chains.

Definition 9.7. $\mathbb{S}^\infty[X]$ is the set of antichains of monomials with indeterminates from X and exponents from $\mathbb{N}^\infty$. We write such antichains as formal sums of their monomials and call them *generalised absorptive polynomials*.

Addition and multiplication of polynomials are defined as usual, but we keep only $\succcurlyeq$-maximal monomials in the result and disregard coefficients.

Notice: There is no difference between polynomials and formal power series here, since antichains of monomials are finite.

Theorem 9.8 (Universality)**.** Every mapping $h\colon X \to S$ into an absorptive, fully continuous semiring S extends uniquely to a fully-continuous semiring homomorphism $\hat{h}\colon \mathbb{S}^\infty[X] \to S$.

In absorptive semirings, the powers of an element a form a descending chain $1 \geq a \geq a^2 \geq \dots$ with infimum a^∞. Hence $h(x) = a$ implies $\hat{h}(x^\infty) = \hat{h}(\bigsqcap_{n<\omega} x^n) = \bigsqcap_{n<\omega}(h(x))^n = a^\infty$ by continuity of $\hat{h}$. It is thus not difficult to see that $\hat{h}$ is uniquely defined. The non-trivial part of the proof is the one showing that $\hat{h}$ is fully continuous.

The semirings $\mathbb{S}^\infty[X]$ and $\mathbb{S}^\infty[X, \overline{X}]$ are the "right" most general semirings for LFP.

9.3 Case Study: Büchi Games

In a Büchi game $\mathcal{G} = (V, V_0, V_1, E, F)$, Player 0 wins plays that hit F infinitely often. Winning regions are LFP-definable by

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

$$\begin{aligned}\varphi(Y, Z, z) := &\Big(Fz \ \wedge\ ((V_0 z \wedge \exists u(Ezu \wedge Yu)) \vee (V_1 z \wedge \forall u(Ezu \to Yu)))\Big)\\ &\vee \Big(\neg Fz \ \wedge\ ((V_0 z \wedge \exists u(Ezu \wedge Zu)) \vee (V_1 z \wedge \forall u(Ezu \to Zu)))\Big).\end{aligned}$$

Idea: Compute $\pi[\![\mathrm{win}(v)]\!]$, where π is a semiring interpretation of a Büchi game (see Figure 9.1 for an example).

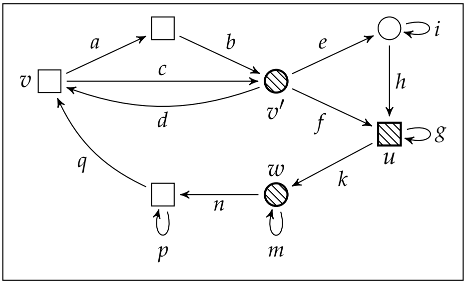

Figure 9.1. $\mathbb{S}^\infty[X]$-interpretation π_{strat} of a Büchi game (dashed nodes are in F).

Theorem 9.9 (Sum of Strategies)**.**

$$\pi[\![\mathrm{win}(v)]\!] = \sum \left\{ \pi[\![\mathcal{S}]\!] \;\middle|\; \begin{array}{l} \mathcal{S} \text{ is an absorption-dominant} \\ \text{winning strategy from } v \end{array} \right\},$$

where

$$\pi[\![\mathcal{S}]\!] := \prod_{e \in E} e^{\#_e(\mathcal{S})} \quad (\text{with } \#_e(\mathcal{S}) \in \mathbb{N}^\infty).$$

Here, $\pi\colon \mathcal{G} \to S$ maps a Büchi game into an absorptive, fully-continuous semiring S. For instance, $\pi_{\mathrm{strat}}\colon \mathcal{G} \mapsto \mathbb{S}^\infty[X]$, $e \mapsto x_e$ (we write e for x_e). In the example (see Figure 9.1), we get

$$\pi_{\mathrm{strat}}[\![\mathrm{win}(v)]\!] = \underbrace{(abcd)^\infty}_{\text{positional}} + \underbrace{abce^2h^2(gkm)^\infty}_{\text{positional}} + \underbrace{abcf^2(gkm)^\infty}_{\text{positional}}$$
$$+ \underbrace{abcefh(gkm)^\infty}_{\substack{\text{absortion-dominant,}\\ \text{but not positional}}}.$$

From $\pi_{\mathrm{strat}}[\![\mathrm{win}(v)]\!] \in \mathbb{S}^\infty[X]$ we can derive:

(1) Whether Player 0 wins from v (this holds iff $\pi_{\mathrm{strat}}[\![\mathrm{win}(v)]\!] \neq 0$).
(2) Edge profiles of all *absorption-dominant* winning strategies from v.

(3) The number and shapes of all *positional* winning strategies from v.
(4) Whether Player 0 can still win if a subset $X \subseteq E$ is forbidden.