

Semiring Semantics in the Infinite

Erich Grädel

(joint work with Sophie Brinke, Lovro Mrkonjić, and Matthias Naaf)

Finite and Algorithmic Model Theory, May 2025, Les Houches

Semiring provenance

originates in the seminal work of Green, Karvounarakis, and Tannen 2007.

Idea: Annotate the facts of a database by values of a commutative semiring $(S, +, \cdot, 0, 1)$, which gives us a **semiring interpretation** $\pi : \text{Facts}_A(\tau) \rightarrow S$.

Semiring provenance

originates in the seminal work of Green, Karvounarakis, and Tannen 2007.

Idea: Annotate the facts of a database by values of a commutative semiring $(S, +, \cdot, 0, 1)$, which gives us a **semiring interpretation** $\pi : \text{Facts}_A(\tau) \rightarrow S$.

Propagate these annotations through a query φ .

- $+$ interprets **alternative use** of information (\vee, \exists , unions)
- \cdot interprets **joint** use of information (\wedge, \forall , joins)
- $0 \in S$ interprets **false assertions** and elements $s \neq 0$ provide **annotations for true assertions**.

The resulting valuation $\pi[[\varphi]] \in S$ can give detailed insights about which combinations of facts are responsible for the truth of a statement and further information about **confidence scores, cost analysis, number of evaluation strategies, access levels,**

Example: the tropical semiring for cost computation

A \mathbb{T} -interpretation into the **tropical semiring** $\mathbb{T} = (\mathbb{R}_{\geq 0} \cup \{\infty\}, \min, +, \infty, 0)$ associates with every atomic fact α a value $\text{cost}(\alpha) \in \mathbb{R}_{\geq 0} \cup \{\infty\}$.

If $\text{cost}(\alpha) = \infty$, then α is false, or inaccessible for other reasons.

Such interpretation extend to valuations $\text{cost}[\![\psi]\!]$ that describe the minimal cost for evaluating ψ .

$$\text{cost}[\![\psi]\!] = \min \left\{ \text{cost}[\![S]\!] : S \text{ is an evaluation strategy for } \psi \right\}.$$

Here, the cost of a strategy S is $\text{cost}[\![S]\!] := \sum_{\alpha} \#_{\alpha}(S) \cdot \text{cost}(\alpha)$ where $\#_{\alpha}(S)$ denotes the number of occurrences of α in the strategy S .

Example: the tropical semiring for cost computation

A \mathbb{T} -interpretation into the **tropical semiring** $\mathbb{T} = (\mathbb{R}_{\geq 0} \cup \{\infty\}, \min, +, \infty, 0)$ associates with every atomic fact α a value $\text{cost}(\alpha) \in \mathbb{R}_{\geq 0} \cup \{\infty\}$.

If $\text{cost}(\alpha) = \infty$, then α is false, or inaccessible for other reasons.

Such interpretation extend to valuations $\text{cost}[\![\psi]\!]$ that describe the minimal cost for evaluating ψ .

$$\text{cost}[\![\psi]\!] = \min \left\{ \text{cost}[\![S]\!] : S \text{ is an evaluation strategy for } \psi \right\}.$$

Here, the cost of a strategy S is $\text{cost}[\![S]\!] := \sum_{\alpha} \#_{\alpha}(S) \cdot \text{cost}(\alpha)$ where $\#_{\alpha}(S)$ denotes the number of occurrences of α in the strategy S .

This is an instance of the so-called **Sum-of-Strategies-Theorem**.

Universal semirings

Let X be a **set of indeterminates**, which are used to label the facts that we want to track:

$\pi : \alpha \mapsto X_\alpha$ (untracked facts are mapped to their truth values 0 or 1).

$\mathbb{N}[X]$: semiring of multivariate polynomials in X with coefficients from \mathbb{N} .

Valuations $\pi[\![\varphi]\!] \in \mathbb{N}[X]$ give precise information about which combinations of atomic facts are responsible for the truth of φ , and how often a fact used in an evaluation.

Indeed, by the **Sum-of-Strategies-Theorem**, $\pi[\![\psi]\!] \in \mathbb{N}[X]$ is a sum of monomials $c x_1^{e_1} \cdots x_k^{e_k}$. Each such monomial tells us that there are precisely c distinct successful evaluation strategies for ψ that use e_i times the literal annotated by x_i , where $x_1, \dots, x_k \in X$. In particular, when $\pi[\![\psi]\!] = 0$ no such strategy exists, and hence there is no model of ψ that is compatible with π .

Universal semirings

Let X be a set of indeterminates, which are used to label the facts that we want to track:

$\pi : \alpha \mapsto X_\alpha$ (untracked facts are mapped to their truth values 0 or 1).

$\mathbb{N}[X]$: semiring of multivariate polynomials in X with coefficients from \mathbb{N} .

Valuations $\pi[\![\varphi]\!] \in \mathbb{N}[X]$ give precise information about which combinations of atomic facts are responsible for the truth of φ , and how often a fact used in an evaluation.

Indeed, by the **Sum-of-Strategies-Theorem**, $\pi[\![\psi]\!] \in \mathbb{N}[X]$ is a sum of monomials $c x_1^{e_1} \cdots x_k^{e_k}$. Each such monomial tells us that there are precisely c distinct successful evaluation strategies for ψ that use e_i times the literal annotated by x_i , where $x_1, \dots, x_k \in X$. In particular, when $\pi[\![\psi]\!] = 0$ no such strategy exists, and hence there is no model of ψ that is compatible with π .

Universality: Any function $f : X \rightarrow S$ into an arbitrary semiring S extends uniquely to a semiring homomorphism $h : \mathbb{N}[X] \rightarrow S$.

From Provenance to Semiring Semantics

Val Tannen (2016): Can we extend semiring provenance from positive database queries to first-order logic and other logical systems, giving us systematic and general **semiring semantics**?

Several issues arise:

From Provenance to Semiring Semantics

Val Tannen (2016): Can we extend semiring provenance from positive database queries to first-order logic and other logical systems, giving us systematic and general **semiring semantics**?

Several issues arise:

Negation: not an algebraic semiring operation, and in general not compositional. For a long time, semiring provenance had essentially been confined to positive query languages.

From Provenance to Semiring Semantics

Val Tannen (2016): Can we extend semiring provenance from positive database queries to first-order logic and other logical systems, giving us systematic and general **semiring semantics**?

Several issues arise:

Negation: not an algebraic semiring operation, and in general not compositional. For a long time, semiring provenance had essentially been confined to positive query languages.

Fixed points: It was known how to treat **Datalog**, using **ω -continuous semirings**. But it was unclear, how to deal with **greatest fixed points** (or interleavings of least fixed points and negation).

From Provenance to Semiring Semantics

Val Tannen (2016): Can we extend semiring provenance from positive database queries to first-order logic and other logical systems, giving us systematic and general **semiring semantics**?

Several issues arise:

Negation: not an algebraic semiring operation, and in general not compositional. For a long time, semiring provenance had essentially been confined to positive query languages.

Fixed points: It was known how to treat **Datalog**, using **ω -continuous semirings**. But it was unclear, how to deal with **greatest fixed points** (or interleavings of least fixed points and negation).

Infinity: Is semiring semantics confined to finite domains, or can it be extended to infinite ones? The obvious problem is the treatment of quantifiers

$$\pi[\![\exists x \varphi(x)]\!] := \sum_{a \in A} \pi[\![\varphi(a)]\!] \quad \text{and} \quad \pi[\![\forall x \varphi(x)]\!] := \prod_{a \in A} \pi[\![\varphi(a)]\!]$$

Semirings with Infinitary Operations?

2016-2022: New approaches for dealing with negation and fixed points, based on negation normal forms, semirings of polynomials with dual indeterminates, fully continuous and absorptive semirings, and generalised absorptive polynomials.

(joint work mainly with Val Tannen and Matthias Naaf)

Semirings with Infinitary Operations?

2016-2022: New approaches for dealing with negation and fixed points, based on negation normal forms, semirings of polynomials with dual indeterminates, fully continuous and absorptive semirings, and generalised absorptive polynomials.

(joint work mainly with Val Tannen and Matthias Naaf)

But until quite recently, provenance on infinite domains has not been systematically considered. It requires the expansion of semirings by **infinitary operations**

$$\sum_{i \in I} s_i \quad \text{and} \quad \prod_{i \in I} s_i \quad \text{for arbitrary index sets } I.$$

Semirings with Infinitary Operations?

2016-2022: New approaches for dealing with negation and fixed points, based on **negation normal forms**, **semirings of polynomials with dual indeterminates**, **fully continuous and absorptive semirings**, and **generalised absorptive polynomials**.

(joint work mainly with Val Tannen and Matthias Naaf)

But until quite recently, provenance on infinite domains has not been systematically considered. It requires the expansion of semirings by **infinitary operations**

$$\sum_{i \in I} s_i \quad \text{and} \quad \prod_{i \in I} s_i \quad \text{for arbitrary index sets } I.$$

Is there a reasonable algebraic notion of such **infinitary semirings** ?

Basic Requirements for an Infinitary Operation

Expand a commutative monoid $S = (S, +, 0)$ by operation $\sum_{i \in I} s_i$, for arbitrary index sets I .

Partition invariance (infinite associativity): For each partition $(I_j)_{j \in J}$ of I

$$\sum_{i \in I} s_i = \sum_{j \in J} \sum_{i \in I_j} s_i.$$

Bijection invariance (infinite commutativity): For every bijection $\sigma: J \rightarrow I$

$$\sum_{i \in I} s_i = \sum_{j \in J} s_{\sigma(j)}.$$

Compatibility with the finite: For each finite index set $I = \{i_0, \dots, i_n\}$

$$\sum_{i \in I} s_i = s_{i_0} + \dots + s_{i_n}.$$

The need for further requirements

The three basic conditions do not exclude pathological infinitary operations.

Example: $(\mathbb{N} \cup \{\infty\}, +, 0)$ with $\sum_{i \in I} s_i = \infty$ for all **infinite** I (and compatible with $+$ for finite I).

This violates, for instance, the following two natural properties.

Neutrality: \sum respects the neutral element if $\sum_{i \in I} s_i = \sum_{i \in I, s_i \neq 0} s_i$.

Idempotence: \sum respects idempotent elements if $s + s = s$ implies that $\sum_{i \in I} s = s$ for all $I \neq \emptyset$.

We can guarantee these, and other relevant conditions by imposing

Compactness properties: Informally, compactness means that the value of an infinitary sum only depends on the set of values of its finite subsums.

Compactness: pros and cons

Compactness has a number of important consequences:

- compact operators respect the neutral element and idempotent elements
- compactness implies that the monoid is aperiodic and naturally ordered

Problem. Compactness is sometimes difficult to verify, and some relevant infinitary semirings violate compactness. An important example is the **universal infinitary semiring** $\mathbb{N}^\infty[[X^\infty]]$

We therefore prefer to use a weaker condition. Notice that **bijection invariance** implies that $\sum_{i \in I} s = \sum_{j \in J} s$ whenever $|I| = |J|$. Compactness implies a stronger property.

Unique infinite powers: For every element $s \in S$ there exists a unique element $\infty \cdot s$ with $\sum_{i \in I} s = \infty \cdot s$ for all infinite I .

Distributive laws

The two algebraic operations in a semiring are related by the distributive law $s(r + t) = sr + st$.

The generalisation to infinitary operations comes in two variants:

Weak distributivity: $s \cdot \sum_{i \in I} s_i = \sum_{i \in I} (s \cdot s_i)$.

Strong distributivity: For every index set I and every collection $(J_i)_{i \in I}$ of index sets

$$\prod_{i \in I} \sum_{j \in J_i} s_j = \sum_{f \in F} \prod_{i \in I} s_{f(i)},$$

where F is the set of all choice functions $f : I \rightarrow \bigcup_{i \in I} J_i$ such that $f(i) \in J_i$ for all $i \in I$.

Weak distributivity is straightforward, but there are issues with strong distributivity:

- It is used to prove that **products are monotone** wrt. to the natural order.
- It is needed for certain results, such as the **Sum-of-Strategies Theorem**.
- In the **tropical semiring**, strong distributivity **only holds for countable products**.

Monotonicity

It is an important property of naturally ordered semirings that addition and multiplication are monotone. Since this is needed in many results, we also want to have it for infinitary operations.

Monotonicity: For all families $(s_i)_{i \in I}$ and $(t_i)_{i \in I}$ such that $s_i \leq t_i$, also

$$\sum_{i \in I} s_i \leq \sum_{i \in I} t_i \quad \text{and} \quad \prod_{i \in I} s_i \leq \prod_{i \in I} t_i.$$

Monotonicity for the infinitary sum is implied by partition invariance, but monotonicity of infinitary products does not seem to follow from weaker properties than strong distributivity.

Indeed there exist “pathological” semirings that satisfy all properties that we discussed except monotonicity and strong distributivity.

We can live without strong distributivity but not without monotonicity.

Infinitary Semirings

Definition. An **infinitary semiring** is a commutative, naturally ordered semiring $S = (S, +, \cdot, 0, 1)$, together with infinitary operations \sum and \prod that satisfy:

- partition invariance and hence also bijection invariance,
- compatibility with finite addition and multiplication,
- neutral elements are respected,
- there are unique infinite powers,
- weak distributivity, and
- monotonicity.

Examples of infinitary semirings

Finite semirings in which both operations are aperiodic expand to infinitary semirings.

Infinite lattice semirings $(S, \sqcup, \sqcap, 0, 1)$ expand to infinitary semirings if the underlying order is a complete lattice in which finite infima distribute over arbitrary suprema.

$\mathbb{N}^\infty = \mathbb{N} \cup \{\infty\}$, with the natural definitions of \sum and \prod is a strongly distributive infinitary semiring.

Infinitary absorptive semirings: Take an absorptive semiring S whose natural order (S, \leq) is a complete lattice, and which is (fully) continuous: suprema $\bigsqcup C$ and infima $\bigsqcap C$ of chains are compatible with addition and multiplication.

Define infinitary operations by taking suprema of finite subsums and infima of finite subproducts:

$$\sum_{i \in I} s_i := \bigsqcup_{\substack{I_0 \subseteq I \\ I_0 \text{ finite}}} \left(\sum_{i \in I_0} s_i \right) \quad \text{and} \quad \prod_{i \in I} s_i := \bigsqcap_{\substack{I_0 \subseteq I \\ I_0 \text{ finite}}} \left(\prod_{i \in I_0} s_i \right).$$

The semiring $\mathbb{N}^\infty[[X^\infty]]$ of generalised power series

Can we generalise $\mathbb{N}[X]$ to an infinitary semiring with an analogous universality property?

The semiring $\mathbb{N}^\infty[[X^\infty]]$ of generalised power series

Can we generalise $\mathbb{N}[X]$ to an **infinitary semiring** with an analogous universality property?

We must be able to

- add the same monomial infinitely often: $x + x + x + \dots$
allow coefficients in \mathbb{N}^∞
- add infinitely many different monomials: $x + x^2 + x^3 + \dots$
use formal power series instead of polynomials
- multiply the same variable infinitely often: $x \cdot x \cdot x \cdot \dots$
allow exponents in \mathbb{N}^∞ .

For **finite** sets X of indeterminates, and an appropriate definition for infinitary products, this indeed makes $\mathbb{N}^\infty[[X^\infty]]$ the free strongly distributive infinitary semiring.

The semiring $\mathbb{N}^\infty[[X^\infty]]$ of generalised power series

Can we generalise $\mathbb{N}[X]$ to an **infinitary semiring** with an analogous universality property?

We must be able to

- add the same monomial infinitely often: $x + x + x + \dots$
allow coefficients in \mathbb{N}^∞
- add infinitely many different monomials: $x + x^2 + x^3 + \dots$
use formal power series instead of polynomials
- multiply the same variable infinitely often: $x \cdot x \cdot x \cdot \dots$
allow exponents in \mathbb{N}^∞ .

For **finite** sets X of indeterminates, and an appropriate definition for infinitary products, this indeed makes $\mathbb{N}^\infty[[X^\infty]]$ the free strongly distributive infinitary semiring.

For infinite sets X , this does not work: $\mathbb{N}^\infty[[X^\infty]]$ is not an infinitary semiring.

Semiring provenance for first-order logic

Infinitary semirings admit to extend most of the results on semiring provenance for FO to infinite domains, sometimes with minor modifications.

An interesting example is the analysis of successful evaluation strategies for model checking problems $\mathfrak{A} \models \psi$, i.e. **winning strategies in the model-checking game**.

Sum-of-Strategies Theorem. Let A be domain of cardinality $< \kappa$, and let S be a κ -distributive infinitary semiring. For every interpretation $\pi : \text{Lit}_A(\tau) \rightarrow S$, and every $\psi \in \text{FO}$

$$\pi[\![\psi]\!] = \sum \left\{ \pi[\![\mathcal{S}]\!] : \mathcal{S} \text{ is a winning strategy in the model-checking game for } \psi \text{ and } \pi \right\}$$

Here, the valuation of a winning strategy \mathcal{S} is $\pi[\![\mathcal{S}]\!] := \prod_{\alpha \in \text{Lit}_A(\tau)} \pi(\alpha)^{\#_{\alpha}(\mathcal{S})}$ where $\#_{\alpha}(\mathcal{S})$ denotes the number of occurrences of the literal α in the strategy \mathcal{S} .

Semiring provenance for first-order logic

This nicely combines with universal semirings and interpretations, defining a class of compatible models.

Corollary. Let $\pi : \text{Lit}_A(\tau) \rightarrow \mathbb{N}^\infty[[X^\infty, \overline{X}^\infty]]$ be model-compatible and let $\psi \in \text{FO}(\tau)$. Then the power series $\pi[[\psi]]$ describes all strategies that verify ψ using premises from the literals that π does not map to 0, i.e. that are mapped to indeterminates or to 1.

Specifically, each monomial $c x_1^{e_1} \cdots x_k^{e_k}$ in $\pi[[\psi]]$ stands for c distinct strategies that use e_1 times the literal annotated by x_1, \dots , and e_k times the literal annotated by x_k , where $x_1, \dots, x_k \in X \cup \overline{X}$. In particular, when $\pi[[\psi]] = 0$ no such strategy exists, and hence there is no model of ψ that is compatible with π .