Semiring Provenance for Logic and Games

Erich Grädel

joint work with Val Tannen

HVW 2018, Obertauern, January 2018

Oxford English Dictionary: provenance, n

The fact of coming from some particular source or quarter; origin, derivation.

Provenance questions in databases

Consider a computational process applied to a complex input consisting of multiple input items, such as a database query.

- Which input items are actually used in the computation of the output?
- Can we derive the same output from different combinations of input items?
- In how many different ways can the same output be computed?

Provenance questions in databases

Consider a computational process applied to a complex input consisting of multiple input items, such as a database query.

- Which input items are actually used in the computation of the output?
- Can we derive the same output from different combinations of input items?
- In how many different ways can the same output be computed?

Interesting for refined questions, beyond the truth or falsity of a query:

Confidence: How much to trust the output, assuming different levels of trust to the input items?

Access control: What clearance level is required for computing the output, assuming we know required clearance levels for the input items?

Cost: How to minimize the cost for obtaining the output based on prizes attached to input items?

Provenance approach in databases

To answer such questions, annotate input items (the atomic facts in a database) not just by true/false, but by elements of some appropriate structure K.

Provenance approach in databases

To answer such questions, annotate input items (the atomic facts in a database) not just by true/false, but by elements of some appropriate structure K.

Propagate annotations through database operations, keeping track of whether pieces of information are used jointly or alternatively. The laws of how information is processed in database systems imply that K must have the structure of a commutative semiring.

Provenance approach in databases

To answer such questions, annotate input items (the atomic facts in a database) not just by true/false, but by elements of some appropriate structure K.

Propagate annotations through database operations, keeping track of whether pieces of information are used jointly or alternatively. The laws of how information is processed in database systems imply that K must have the structure of a commutative semiring.

Provenance analysis by means of interpretations in semirings has been quite successful and influential. Test of time award and invited Gems of PODS Talk at PODS 2017 by Val Tannen.

Commutative Semirings

$(K, +, \cdot, 0, 1)$ with $0 \neq 1$, is a commutative semiring when $(K, +, 0)$ and $(K, \cdot, 1)$ are commutative monoids, $\cdot$ distributes over $+$ and $0 \cdot a = a \cdot 0 = 0$.

Commutative Semirings

$(K, +, \cdot, 0, 1)$ with $0 \neq 1$, is a commutative semiring when $(K, +, 0)$ and $(K, \cdot, 1)$ are commutative monoids, $\cdot$ distributes over $+$ and $0 \cdot a = a \cdot 0 = 0$.

A semiring K is positive if it has no divisors of 0 (i.e., $a \cdot b = 0$ implies that $a = 0$ or $b = 0$) and if $a + b = 0$ implies that $a = b = 0$. This is the case if, and only if, the unique function $h : K \to \{0, 1\}$ with $h^{-1}(0) = \{0\}$ is a homomorphism from K into the Boolean semiring $\mathbb{B} = (\{0, 1\}, \vee, \wedge, 0, 1)$.

Commutative Semirings

$(K, +, \cdot, 0, 1)$ with $0 \neq 1$, is a commutative semiring when $(K, +, 0)$ and $(K, \cdot, 1)$ are commutative monoids, $\cdot$ distributes over $+$ and $0 \cdot a = a \cdot 0 = 0$.

A semiring K is positive if it has no divisors of 0 (i.e., $a \cdot b = 0$ implies that $a = 0$ or $b = 0$) and if $a + b = 0$ implies that $a = b = 0$. This is the case if, and only if, the unique function $h : K \to \{0, 1\}$ with $h^{-1}(0) = \{0\}$ is a homomorphism from K into the Boolean semiring $\mathbb{B} = (\{0, 1\}, \vee, \wedge, 0, 1)$.

Intuition.

- $+$ interprets alternative use of information (union, $\vee$, $\exists$)
- $\cdot$ interprets joint use of information (joins, $\wedge$, $\forall$)
- $0 \in K$ interprets false assertions and an element $a \in K, a \neq 0$ provides a "nuanced" interpretation for true assertions.

Examples of commutative semirings

The Boolean semiring $\mathbb{B} = (\{0,1\}, \vee, \wedge, 0, 1)$ is the standard habitat of logic.

Examples of commutative semirings

The Boolean semiring $\mathbb{B} = (\{0,1\}, \vee, \wedge, 0, 1)$ is the standard habitat of logic.

$\mathbb{N} = (\mathbb{N}, +, \cdot, 0, 1)$ is used for bag semantics in databases. Not idempotent.

Examples of commutative semirings

The Boolean semiring $\mathbb{B} = (\{0,1\}, \vee, \wedge, 0, 1)$ is the standard habitat of logic.

$\mathbb{N} = (\mathbb{N}, +, \cdot, 0, 1)$ is used for bag semantics in databases. Not idempotent.

The tropical semiring $\mathbb{T} = (\mathbb{R}_+^\infty, \min, +, \infty, 0)$, which is idempotent but not a distributive lattice, is used for min-cost interpretations in logic and games.

Examples of commutative semirings

The Boolean semiring $\mathbb{B} = (\{0,1\}, \vee, \wedge, 0, 1)$ is the standard habitat of logic.

$\mathbb{N} = (\mathbb{N}, +, \cdot, 0, 1)$ is used for bag semantics in databases. Not idempotent.

The tropical semiring $\mathbb{T} = (\mathbb{R}_+^\infty, \min, +, \infty, 0)$, which is idempotent but not a distributive lattice, is used for min-cost interpretations in logic and games.

The Viterbi semiring $\mathbb{V} = ([0,1], \max, \cdot, 0, 1)$ is isomorphic to $\mathbb{T}$ via $x \mapsto e^{-x}$. It is used here for confidence scores.

Examples of commutative semirings

The Boolean semiring $\mathbb{B} = (\{0,1\}, \vee, \wedge, 0, 1)$ is the standard habitat of logic.

$\mathbb{N} = (\mathbb{N}, +, \cdot, 0, 1)$ is used for bag semantics in databases. Not idempotent.

The tropical semiring $\mathbb{T} = (\mathbb{R}_+^\infty, \min, +, \infty, 0)$, which is idempotent but not a distributive lattice, is used for min-cost interpretations in logic and games.

The Viterbi semiring $\mathbb{V} = ([0,1], \max, \cdot, 0, 1)$ is isomorphic to $\mathbb{T}$ via $x \mapsto e^{-x}$. It is used here for confidence scores.

The access control semiring $\mathbb{A} = (\{\mathsf{P} < \mathsf{C} < \mathsf{S} < \mathsf{T} < 0\}, \min, \max, 0, \mathsf{P})$ is a distributive lattice where P is "public", C is "confidential", S is "secret", T is "top secret", and 0 is "so secret that nobody can access it!"

One semiring to rule them all

For any set X of provenance tokens, the semiring $\mathbb{N}[X] = (\mathbb{N}[X], +, \cdot, 0, 1)$ of multivariate polynomials in indeterminates from X and with coefficients from $\mathbb{N}$ is the commutative semiring freely generated by the set X.

It is used for a general form provenance.

One semiring to rule them all

For any set X of provenance tokens, the semiring $\mathbb{N}[X] = (\mathbb{N}[X], +, \cdot, 0, 1)$ of multivariate polynomials in indeterminates from X and with coefficients from $\mathbb{N}$ is the commutative semiring freely generated by the set X.

It is used for a general form provenance.

Any function $f : X \to K$ from provenance tokens into a semiring K extends uniquely to a semiring homomorphism $h : \mathbb{N}[X] \to K$.

Think of $h(2x^2 + xy + 3z^2)$ as evaluating $2x^2 + xy + 3z^2$ in K.

How to deal with negation?

The bulk of the work on provenance in databases has been done for positive query languages such as positive relational algebra, (unions of) conjunctive queries and, to a lesser extent, datalog.

Val Tannen : “Divergent approaches and unsatisfactory state of affairs for queries with negation or difference of relations."

Goal: Extend the provenance approach to full FO (and fixed-point logics).

How to deal with negation?

The bulk of the work on provenance in databases has been done for positive query languages such as positive relational algebra, (unions of) conjunctive queries and, to a lesser extent, datalog.
Val Tannen : “Divergent approaches and unsatisfactory state of affairs for queries with negation or difference of relations."

Goal: Extend the provenance approach to full FO (and fixed-point logics).

Idea: Define provenance for logics through their model-checking games.

This implies that we deal with negation syntactically, trough translations into negation normal form.

In fact, it turns out that provenance for games is of independent interest and provides interesting information on positions in games, far beyond the question which player wins.

Provenance for finite games

Acyclic two player-game $\mathscr{G} = (V, V_0, V_1, T, E)$ with $V = V_0 \cup V_1 \cup T$

V_σ: positions of Player σ, T: terminal positions, $E \subseteq V \times V$: moves

Provenance for finite games

Acyclic two player-game $\mathscr{G} = (V, V_0, V_1, T, E)$ with $V = V_0 \cup V_1 \cup T$

V_σ: positions of Player σ, T: terminal positions, $E \subseteq V \times V$: moves

Valuations $f_\sigma : T \to K$ of terminal positions and $h_\sigma : E \to K \setminus \{0\}$ of moves in a semiring K.

- $f_\sigma(t)$ describes the value of the terminal position v for Player σ.
  $f_\sigma(t) = 0$ means that t is a losing position
- $h_\sigma(vw)$ describes the value (or cost) for Player σ of a move from v to w.
  (Values of moves may be irrelevant. In that case, set $h_\sigma(vw) = 1$.)

Provenance for finite games

Acyclic two player-game $\mathscr{G} = (V, V_0, V_1, T, E)$ with $V = V_0 \cup V_1 \cup T$

V_σ: positions of Player σ, T: terminal positions, $E \subseteq V \times V$: moves

Valuations $f_\sigma : T \to K$ of terminal positions and $h_\sigma : E \to K \setminus \{0\}$ of moves in a semiring K.

- $f_\sigma(t)$ describes the value of the terminal position v for Player σ.
  $f_\sigma(t) = 0$ means that t is a losing position
- $h_\sigma(vw)$ describes the value (or cost) for Player σ of a move from v to w.
  (Values of moves may be irrelevant. In that case, set $h_\sigma(vw) = 1$.)

Extension to valuations $f_\sigma : V \to K$ for all positions. A move from v to w contributes to $f_\sigma(v)$ the value $h_\sigma(vw) \cdot f_\sigma(w)$.

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

Reachability games and contradictory valuations

For acyclic game graphs $\mathscr{G} = (V, V_0, V_1, T, E)$, and semiring valuations $f_\sigma : V \to K$, Player σ has a winning strategy for the reachability objective $T \setminus f_\sigma^{-1}(0)$ from all positions v with $f_\sigma(v) \neq 0$.

Reachability games and contradictory valuations

For acyclic game graphs $\mathscr{G} = (V, V_0, V_1, T, E)$, and semiring valuations $f_\sigma : V \to K$, Player σ has a winning strategy for the reachability objective $T \setminus f_\sigma^{-1}(0)$ from all positions v with $f_\sigma(v) \neq 0$.

On a set $U \subseteq V$ the valuations f_0, f_1 are

- contradictory if either $f_0(u) = 0$ or $f_1(u) = 0$ for all $u \in U$,
- weakly contradictory if just $f_0(u) \cdot f_1(u) = 0$,
- strongly contradictory if, in addition, $f_0(u) + f_1(u) \neq 0$.

Reachability games and contradictory valuations

For acyclic game graphs $\mathcal{G} = (V, V_0, V_1, T, E)$, and semiring valuations $f_\sigma : V \to K$, Player σ has a winning strategy for the reachability objective $T \setminus f_\sigma^{-1}(0)$ from all positions v with $f_\sigma(v) \neq 0$.

On a set $U \subseteq V$ the valuations f_0, f_1 are
- contradictory if either $f_0(u) = 0$ or $f_1(u) = 0$ for all $u \in U$,
- weakly contradictory if just $f_0(u) \cdot f_1(u) = 0$,
- strongly contradictory if, in addition, $f_0(u) + f_1(u) \neq 0$.

If f_0 and f_1 are (weakly) contradictory on the the terminal positions of $\mathcal{G}$, then they are (weakly) contradictory on all positions of $\mathcal{G}$.
For positive semirings, also strongly contradictory valuations on the terminal positions extend to strongly contradictory ones on all positions.

For the Boolean semiring $\mathbb{B} = (\{0,1\}, \vee, \wedge, 0, 1)$ this is just the determinacy of reachability games on well-founded game graphs.

Applications for different semirings

(1) The tropical semiring and the cost of strategies. On $\mathcal{G}$, let $f_0 : T \to \mathbb{R}_+$ and $h_0 : E \to \mathbb{R}_+$ be cost functions for Player 0 on the terminal positions and the moves.

The cost of a play $\pi = v_0 v_1 \dots v_m$ for Player 0 is defined as
$c(\pi) := \sum_{i=0}^{m-1} h_0(v_i v_{i+1}) + f_0(v_m)$.
The cost of a strategy from v is the sum of the costs of all plays from v that are admitted by the strategy.

Proposition. The cost of an optimal strategy from v in a game $\mathcal{G}$ with basic cost functions $f_0 : T \to \mathbb{R}_+$ and $h_0 : E \to \mathbb{R}_+$ is given by the valuation $f_0(v)$ computed in the tropical semiring $(\mathbb{R}_+^\infty, \min, +, \infty, 0)$.

Applications for different semirings

(2) The access control semiring $\mathbb{A} = (\{\mathsf{P} < \mathsf{C} < \mathsf{S} < \mathsf{T} < 0\}, \min, \max, 0, \mathsf{P})$.
Let $f_0 : T \to \mathbb{A}$ and $h_0 : E \to \mathbb{A} \setminus \{0\}$ define access levels for the terminal positions and the moves.

The valuation $f_0(v) \in \mathbb{A}$ then describes the minimal clearance level that Player 0 needs to win from position v.

(3) Confidence scores. Based on confidences $f_\sigma : T \to [0,1]$ that Player σ puts into t being a winning position for her, compute confidence scores $f_\sigma(v)$ to describe the confidence of Player σ that she can win from v, as semiring valuations in the Viterbi semiring $\mathbb{V} = ([0,1], \max, \cdot, 0, 1)$.

Counting winning strategies

Let $\mathbb{N}[T]$ be the semiring of polynomials over indeterminates $t \in T$.
For a game $\mathscr{G}$, let $f_\sigma : V \to \mathbb{N}[T]$ be the valuation induced by $f_\sigma(t) = t$.
We can write $f_\sigma(v)$ as a sum of monomials $t_1^{j_1} \cdots t_k^{j_k}$.

Counting winning strategies

Let $\mathbb{N}[T]$ be the semiring of polynomials over indeterminates $t \in T$.
For a game $\mathscr{G}$, let $f_\sigma : V \to \mathbb{N}[T]$ be the valuation induced by $f_\sigma(t) = t$.
We can write $f_\sigma(v)$ as a sum of monomials $t_1^{j_1} \cdots t_k^{j_k}$.

Each monomial $t_1^{j_1} \cdots t_k^{j_k}$ in $f_\sigma(v)$ indicates a strategy of Player σ from v whose set of possible outcomes is precisely $\{t_1, \ldots, t_k\}$, and precisely j_i plays that are compatible with that strategy have the outcome t_i.

Counting winning strategies

Let $\mathbb{N}[T]$ be the semiring of polynomials over indeterminates $t \in T$.
For a game $\mathscr{G}$, let $f_\sigma : V \to \mathbb{N}[T]$ be the valuation induced by $f_\sigma(t) = t$.
We can write $f_\sigma(v)$ as a sum of monomials $t_1^{j_1} \cdots t_k^{j_k}$.

Each monomial $t_1^{j_1} \cdots t_k^{j_k}$ in $f_\sigma(v)$ indicates a strategy of Player σ from v whose set of possible outcomes is precisely $\{t_1, \ldots, t_k\}$, and precisely j_i plays that are compatible with that strategy have the outcome t_i.

Fix any reachability objective $W \subseteq T$. Let $f_\sigma(v) = f_\sigma^W(v) + g_\sigma^W(v)$ where $f_\sigma^W(v)$ is the sum of those monomials that only contain indeterminates in W.

Counting winning strategies

Let $\mathbb{N}[T]$ be the semiring of polynomials over indeterminates $t \in T$.
For a game $\mathscr{G}$, let $f_\sigma : V \to \mathbb{N}[T]$ be the valuation induced by $f_\sigma(t) = t$.
We can write $f_\sigma(v)$ as a sum of monomials $t_1^{j_1} \cdots t_k^{j_k}$.

Each monomial $t_1^{j_1} \cdots t_k^{j_k}$ in $f_\sigma(v)$ indicates a strategy of Player σ from v whose set of possible outcomes is precisely $\{t_1, \ldots, t_k\}$, and precisely j_i plays that are compatible with that strategy have the outcome t_i.

Fix any reachability objective $W \subseteq T$. Let $f_\sigma(v) = f_\sigma^W(v) + g_\sigma^W(v)$ where $f_\sigma^W(v)$ is the sum of those monomials that only contain indeterminates in W.

Theorem. Player σ has a strategy to reach W from v if, and only if, $f_\sigma^W(v) \neq 0$. Moreover, if $f_\sigma^W(v) = \sum_{j \in J} c_j M_j$ (where M_j are monomials with indeterminates in W), then $\sum_{j \in J} c_j$ is the number of distinct strategies from v that Player σ has for the reachability objective W.

Dualities

Games coming from logical formulae usually have (by negation) a bijection between winning and losing terminal positions.

A duality on the terminal positions of $\mathcal{G} = (V, V_0, V_1, T, E)$ is a self-inverse bijection $(\cdot)^* : t \mapsto t^*$ on T.
A valuation $f_0 : T \to K$ is sound for $(\cdot)^*$ if $f_0(t) f_0(t^*) = 0$ for all $t \in T$.

Dualities

Games coming from logical formulae usually have (by negation) a bijection between winning and losing terminal positions.

A duality on the terminal positions of $\mathcal{G} = (V, V_0, V_1, T, E)$ is a self-inverse bijection $(\cdot)^* : t \mapsto t^*$ on T.
A valuation $f_0 : T \to K$ is sound for $(\cdot)^*$ if $f_0(t) f_0(t^*) = 0$ for all $t \in T$.

Let $\mathbb{N}[T]/tt^*$ be the quotient semiring of $\mathbb{N}[T]$ with respect to the ideal generated by tt^* for all $t \in T$. We can describe $\mathbb{N}[T]/tt^*$ as the semiring of all polynomials $p \in \mathbb{N}[T]$ such that no monomial in p contains both t and t^*, for any $t \in T$.

Universality. For any sound valuation $f : T \to K$ in a semiring K there is a unique semiring homomorphism $h : \mathbb{N}[T]/tt^* \to K$ with $h(t) = f(t)$ for $t \in T$.

Provenance analysis for first-order logic

Let A be a finite universe and τ a finite relational vocabulary.

$\mathrm{Lit}_A(\tau) := \mathrm{Atoms}_A(\tau) \cup \mathrm{NegAtoms}_A(\tau) \cup \{a \stackrel{\neq}{=} b : a, b \in A\}$

A K-interpretation for A and τ is a function $\pi : \mathrm{Lit}_A(\tau) \to K$ that maps equalities and inequalities to their truth values.

If, for all atoms $R\bar{a}$, either $\pi(R\bar{a}) = 0$ or $\pi(\neg R\bar{a}) = 0$, ("consistency"), and, moreover, $\pi(R\bar{a}) + \pi(\neg R\bar{a}) \neq 0$ ("completeness"), then π specifies (provenance information for) a unique structure $\mathfrak{A}_\pi$.
Otherwise, π gives provenance information for a whole class of structures.

Provenance analysis for first-order logic

Let A be a finite universe and τ a finite relational vocabulary.

$\mathrm{Lit}_A(\tau) := \mathrm{Atoms}_A(\tau) \cup \mathrm{NegAtoms}_A(\tau) \cup \{a \stackrel{\neq}{=} b : a, b \in A\}$

A K-interpretation for A and τ is a function $\pi : \mathrm{Lit}_A(\tau) \to K$ that maps equalities and inequalities to their truth values.

If, for all atoms $R\bar{a}$, either $\pi(R\bar{a}) = 0$ or $\pi(\neg R\bar{a}) = 0$, ("consistency"), and, moreover, $\pi(R\bar{a}) + \pi(\neg R\bar{a}) \neq 0$ ("completeness"), then π specifies (provenance information for) a unique structure $\mathfrak{A}_\pi$.
Otherwise, π gives provenance information for a whole class of structures.

Extend π to a K-interpretation $\pi : \mathrm{FO}(\tau) \to K$:

$$\pi[\![\varphi \vee \psi]\!] := \pi[\![\varphi]\!] + \pi[\![\psi]\!] \qquad \pi[\![\varphi \wedge \psi]\!] := \pi[\![\varphi]\!] \cdot \pi[\![\psi]\!]$$
$$\pi[\![\exists x \varphi(x)]\!] := \sum_{a \in A} \pi[\![\varphi(a)]\!] \qquad \pi[\![\forall x \varphi(x)]\!] := \prod_{a \in A} \pi[\![\varphi(a)]\!]$$
$$\pi[\![\neg\varphi]\!] := \pi[\![\mathrm{nnf}(\neg\varphi)]\!].$$

This extension can also be understood in game-theoretic terms.

Provenance analysis via model-checking games

The standard model-checking game for $\psi \in \mathrm{FO}(\tau)$ and a finite structure $\mathfrak{A}$, has a game graph $\mathscr{G}(A, \psi)$ that only depends on ψ and the universe A.

Provenance analysis via model-checking games

The standard model-checking game for $\psi \in \mathrm{FO}(\tau)$ and a finite structure $\mathfrak{A}$, has a game graph $\mathscr{G}(A, \psi)$ that only depends on ψ and the universe A.

A K-interpretation $\pi : \mathrm{Lit}_A(\tau) \to K$ provides valuations $f_\sigma : T \to K$ of the terminal positions of $\mathscr{G}(A, \psi)$.

Provenance analysis via model-checking games

The standard model-checking game for $\psi \in \mathrm{FO}(\tau)$ and a finite structure $\mathfrak{A}$, has a game graph $\mathscr{G}(A, \psi)$ that only depends on ψ and the universe A.

A K-interpretation $\pi : \mathrm{Lit}_A(\tau) \to K$ provides valuations $f_\sigma : T \to K$ of the terminal positions of $\mathscr{G}(A, \psi)$.

These valuations extend to valuations $f_\sigma : V \to K$ of all positions of $\mathscr{G}(A, \psi)$, and in particular of the initial position ψ itself.

Provenance analysis via model-checking games

The standard model-checking game for $\psi \in \mathrm{FO}(\tau)$ and a finite structure $\mathfrak{A}$, has a game graph $\mathcal{G}(A, \psi)$ that only depends on ψ and the universe A.

A K-interpretation $\pi : \mathrm{Lit}_A(\tau) \to K$ provides valuations $f_\sigma : T \to K$ of the terminal positions of $\mathcal{G}(A, \psi)$.

These valuations extend to valuations $f_\sigma : V \to K$ of all positions of $\mathcal{G}(A, \psi)$, and in particular of the initial position ψ itself.

Proposition. For all positions φ of the game $\mathcal{G}(A, \psi)$,

$$\pi[\![\varphi]\!] = f_0(\varphi) \quad \text{and} \quad \pi[\![\neg\varphi]\!] = f_1(\varphi).$$

Provenance analysis via model-checking games

The standard model-checking game for $\psi \in \mathrm{FO}(\tau)$ and a finite structure $\mathfrak{A}$, has a game graph $\mathcal{G}(A, \psi)$ that only depends on ψ and the universe A.

A K-interpretation $\pi : \mathrm{Lit}_A(\tau) \to K$ provides valuations $f_\sigma : T \to K$ of the terminal positions of $\mathcal{G}(A, \psi)$.

These valuations extend to valuations $f_\sigma : V \to K$ of all positions of $\mathcal{G}(A, \psi)$, and in particular of the initial position ψ itself.

Proposition. For all positions φ of the game $\mathcal{G}(A, \psi)$,

$$\pi[\![\varphi]\!] = f_0(\varphi) \quad \text{and} \quad \pi[\![\neg\varphi]\!] = f_1(\varphi).$$

In particular, if the K-interpretation π defines a unique structure $\mathfrak{A}_\pi$, then $\mathfrak{A}_\pi \models \psi \iff f_0(\psi) \neq 0$, and the provenance information $f_0(\psi)$ reveals information about the number and properties of the strategies of Verifier to establish the the truth of ψ in $\mathfrak{A}_\pi$.

Provenance for reachability games with cycles

Let $\mathcal{G} = (V, V_0, V_1, T, E)$ be a finite, not necessarily acyclic, game graph.

Given a valuation $f_\sigma : T \to K$ in a semiring K for the terminal nodes, the rules defining valuations for the other nodes have now to be read as an equation system in indeterminates X_v (for $v \in V$):

$$
\begin{aligned}
X_v &= f_\sigma(v) \quad \text{for } v \in T \\
X_v &= \sum_{w \in vE} h_\sigma(vw) \cdot X_w \quad \text{if } v \in V_\sigma \\
X_v &= \prod_{w \in vE} h_\sigma(vw) \cdot X_w \quad \text{if } v \in V_{1-\sigma}
\end{aligned}
$$

To make sure that a solution of such a system exists, we assume that the semiring K is naturally ordered and ω-continuous.

ω-continuous semirings

A semiring is naturally ordered if $a \leq b :\Leftrightarrow \exists x(a+x=b)$ is a partial order.

A semiring K is ω-continuous if it is naturally ordered and every ω-chain $a_0 < a_1 < \ldots$ has a supremum $\sup_{i<\omega} a_i$, such that the associated countable summation operator $\sum_{i<\omega} b_i := \sup_{i<\omega}(b_0 + \cdots + b_i)$ is compatible with the operations of K.

A formal power series $f \in K[\![X]\!]$ in variables $X = (X_1, \ldots, X_n)$ is a possibly infinite sum of monomials $c \cdot X_1^{e_1} \ldots X_n^{e_n}$.

Let $F = (f_1 \ldots f_n)$ be a system of formal power series $f_i \in K[\![X]\!]$. If K is ω-continuous, then by Kleene's Fixed-Point Theorem, the equation system $F(X) = X$ has a least fixed-point solution $\mathrm{lfp}(F)$ which is the supremum of the Kleene approximants F^k, defined by $F^0 = 0$, $F^{k+1} = F(F^k)$.

Semirings of power series

Notice that $(\mathbb{N}, +, \cdot, 0, 1)$ is not ω-continuous, but its completion $\mathbb{N}^\infty$ is.
The completion of $\mathbb{N}[X]$ is not $\mathbb{N}^\infty[X]$ but the semiring of (possibly infinite) formal power series, denoted $\mathbb{N}^\infty[\![X]\!]$.

Example.

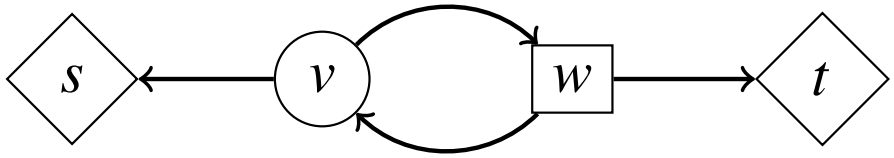

Equation system for valuation of Player 0: $X_v = s + X_w$ and $X_w = t \cdot X_v$

Semirings of power series

Notice that $(\mathbb{N}, +, \cdot, 0, 1)$ is not ω-continuous, but its completion $\mathbb{N}^\infty$ is.
The completion of $\mathbb{N}[X]$ is not $\mathbb{N}^\infty[X]$ but the semiring of (possibly infinite) formal power series, denoted $\mathbb{N}^\infty[\![X]\!]$.

Example.

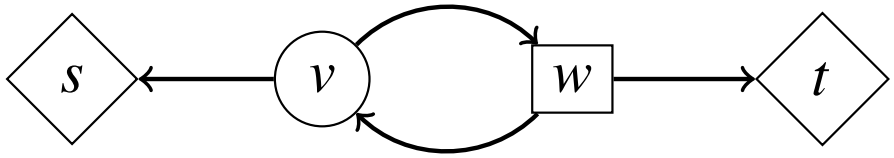

Equation system for valuation of Player 0: $X_v = s + X_w$ and $X_w = t \cdot X_v$

Solution in $\mathbb{N}^\infty[\![s,t]\!]$: $f(v) = s \cdot (1 + t + t^2 + \cdots)$ and $f(w) = s \cdot (t + t^2 + \cdots)$

Evaluation.

- $f(v)(0,t) = f(w)(0,t) = 0$

  Neither from v nor from w, Player 0 has a strategy to reach t.
- $f(v)(s,0) = s$ but $f(w)(s,0) = 0$:

  Player 0 has a strategy to reach s from v, but not from w.

Counting strategies

Again, valuations in $\mathbb{N}^\infty[\![X]\!]$ give more information than just who wins.

Example.

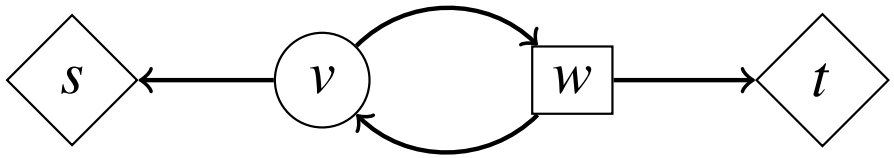

For every $n < \omega$, the monomial $s \cdot t^n$ in $f(v) = s \cdot (1 + t + t^2 + \cdots)$ tells us that Player 0 has precisely one strategy from v that admits $n+1$ consistent plays, one of which has outcome s, and the other n have outcome t.

Counting strategies

Again, valuations in $\mathbb{N}^\infty[\![X]\!]$ give more information than just who wins.

Example.

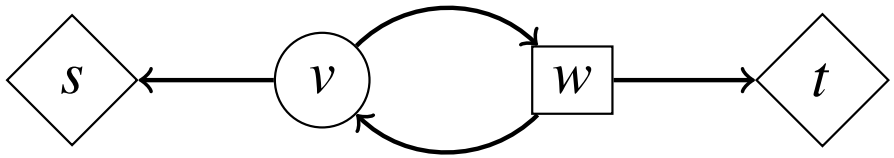

For every $n < \omega$, the monomial $s \cdot t^n$ in $f(v) = s \cdot (1 + t + t^2 + \cdots)$ tells us that Player 0 has precisely one strategy from v that admits $n+1$ consistent plays, one of which has outcome s, and the other n have outcome t.

By evaluating these formal power series in the tropical semiring, the Viterbi semiring, or the access control semiring, we obtain information about the cost of optimal strategies, and the confidence of winning or the required clearance levels for winning reachability games.

Least fixed-point logic

LFP extends FO by least and greatest fixed-points for monotone definable operators. It is a logic of great importance in finite model theory.

On ordered finite structures, LFP captures PTIME. (Immerman, Vardi)

Least fixed-point logic

LFP extends FO by least and greatest fixed-points for monotone definable operators. It is a logic of great importance in finite model theory.

On ordered finite structures, LFP captures PTIME. (Immerman, Vardi)

posLFP is the fragment of LFP that makes use of least fixed points only (which may appear only positively). posLFP is at the bottom level of the alternation hierarchy of LFP. In general, and for instance on $(\mathbb{N}, +, \cdot)$, this hierarchy is strict.

Least fixed-point logic

LFP extends FO by least and greatest fixed-points for monotone definable operators. It is a logic of great importance in finite model theory.

On ordered finite structures, LFP captures PTIME. (Immerman, Vardi)

posLFP is the fragment of LFP that makes use of least fixed points only (which may appear only positively). posLFP is at the bottom level of the alternation hierarchy of LFP. In general, and for instance on $(\mathbb{N}, +, \cdot)$, this hierarchy is strict.

Theorem (Immerman) On finite structures, LFP $\equiv$ posLFP.

The model checking games for general LFP-formulae are parity games, which are not known to solvable in polynomial time. However, the model-checking games for posLFP are reachability games.

Provenance for positive least fixed-point logic

For a finite universe A and a finite relational vocabulary, consider a K-interpretation $\pi : \mathrm{Lit}_A(\tau) \to K$ into an ω-continuous semiring K.

Goal: Extend π to a K-interpretation $\pi : \mathrm{posLFP}(\tau) \to K$

Provenance for positive least fixed-point logic

For a finite universe A and a finite relational vocabulary, consider a K-interpretation $\pi : \mathrm{Lit}_A(\tau) \to K$ into an ω-continuous semiring K.

Goal: Extend π to a K-interpretation $\pi : \mathrm{posLFP}(\tau) \to K$

The model-checking game $\mathscr{G}(\mathfrak{A}, \psi)$ for a posLFP-sentence ψ and a structure $\mathfrak{A}$, is a reachability game whose game graph $\mathscr{G}$ only depends on ψ and the universe A.

Provenance for positive least fixed-point logic

For a finite universe A and a finite relational vocabulary, consider a K-interpretation $\pi : \mathrm{Lit}_A(\tau) \to K$ into an ω-continuous semiring K.

Goal: Extend π to a K-interpretation $\pi : \mathrm{posLFP}(\tau) \to K$

The model-checking game $\mathscr{G}(\mathfrak{A}, \psi)$ for a posLFP-sentence ψ and a structure $\mathfrak{A}$, is a reachability game whose game graph $\mathscr{G}$ only depends on ψ and the universe A.

$\pi : \mathrm{Lit}_A(\tau) \to K$ provides a valuation $f_0 : T \to K$ of the terminal positions of $\mathscr{G}$. It extends to a least fixed-point solution $f_0 : V \to K$ of the equation system describing the game valuation for Player 0 of all positions of $\mathscr{G}(\mathfrak{A}, \psi)$, and in particular of the initial position ψ itself. Now set $\pi(\psi) := f_0(\psi)$.

Provenance for positive least fixed-point logic

For a finite universe A and a finite relational vocabulary, consider a K-interpretation $\pi : \mathrm{Lit}_A(\tau) \to K$ into an ω-continuous semiring K.

Goal: Extend π to a K-interpretation $\pi : \mathrm{posLFP}(\tau) \to K$

The model-checking game $\mathscr{G}(\mathfrak{A}, \psi)$ for a posLFP-sentence ψ and a structure $\mathfrak{A}$, is a reachability game whose game graph $\mathscr{G}$ only depends on ψ and the universe A.

$\pi : \mathrm{Lit}_A(\tau) \to K$ provides a valuation $f_0 : T \to K$ of the terminal positions of $\mathscr{G}$. It extends to a least fixed-point solution $f_0 : V \to K$ of the equation system describing the game valuation for Player 0 of all positions of $\mathscr{G}(\mathfrak{A}, \psi)$, and in particular of the initial position ψ itself. Now set $\pi(\psi) := f_0(\psi)$.

For a general form of provenance for posLFP, use the semirings $\mathbb{N}^\infty[\![X, \overline{X}]\!]$.

Beyond pos LFP

How to deal with full LFP? By moving to formulae in negation normal form, we have to take care of greatest fixed points. However, their existence (and meaning) is unclear in arbitrary ω-continuous semirings. (At least to me!)

Beyond pos LFP

How to deal with full LFP? By moving to formulae in negation normal form, we have to take care of greatest fixed points. However, their existence (and meaning) is unclear in arbitrary ω-continuous semirings. (At least to me!)

A semiring is absorptive if $a + ab = a$ for all a, b. Hence $a < 1$ and $ab < a$.

Consider monomials over a finite set X of provenance tokens, with exponents in $\mathbb{N}^\infty$. Absorption ordering: $x_1^{i_1} \cdots x_m^{i_m} \leq x_1^{j_1} \cdots x_m^{j_m} \iff i_k \geq j_k$ for all k.

Absorptive polynomials over X are antichains of monomials (which are always finite). They form a semiring $\mathbb{S}^\infty[X]$.

Beyond pos LFP

How to deal with full LFP? By moving to formulae in negation normal form, we have to take care of greatest fixed points. However, their existence (and meaning) is unclear in arbitrary ω-continuous semirings. (At least to me!)

A semiring is absorptive if $a + ab = a$ for all a, b. Hence $a < 1$ and $ab < a$.

Consider monomials over a finite set X of provenance tokens, with exponents in $\mathbb{N}^\infty$. Absorption ordering: $x_1^{i_1} \cdots x_m^{i_m} \leq x_1^{j_1} \cdots x_m^{j_m} \iff i_k \geq j_k$ for all k.

Absorptive polynomials over X are antichains of monomials (which are always finite). They form a semiring $\mathbb{S}^\infty[X]$.

Proposition. $\mathbb{S}^\infty[X]$ is a complete lattice with respect to the natural order.

Hence the Tarski-Knaster fixed-point theory applies to $\mathbb{S}^\infty[X]$, and we can inductively define provenance values in $\mathbb{S}^\infty[X]$ for arbitrary LFP-formulae.

Absorptive strategies

We have seen that with any strategy $\mathscr{S}$, we can associate a monomial $M_{\mathscr{S}}$ over the set of terminal positions. The value of a strategy is the product over the values of the plays it admits. Nonterminating plays have value 0.

Absorption: $\mathscr{S} \succeq \mathscr{S}'$ if $M_{\mathscr{S}} \geq M_{\mathscr{S}'}$
This means: for any outcome t, $\mathscr{S}$ admits less plays with outcome t than $\mathscr{S}'$.

In a game $\mathscr{G}$, a strategy $\mathscr{S}$ from v is absorption-dominant if it is not absorbed by any other strategy from v (of the same player).

Lemma. Every absorption-dominant strategy is positional.

The converse is not true.

Absorptive strategies

We have seen that with any strategy $\mathscr{S}$, we can associate a monomial $M_{\mathscr{S}}$ over the set of terminal positions. The value of a strategy is the product over the values of the plays it admits. Nonterminating plays have value 0.

Absorption: $\mathscr{S} \succeq \mathscr{S}'$ if $M_{\mathscr{S}} \geq M_{\mathscr{S}'}$
This means: for any outcome t, $\mathscr{S}$ admits less plays with outcome t than $\mathscr{S}'$.

In a game $\mathscr{G}$, a strategy $\mathscr{S}$ from v is absorption-dominant if it is not absorbed by any other strategy from v (of the same player).

Lemma. Every absorption-dominant strategy is positional.

The converse is not true.

Theorem. Let $\mathscr{G}$ be a reachability game. The provenance values in $\mathbb{S}^{\infty}[T]$ at v give the values of all absorption-dominant strategies from v.

Reachability versus safety games

Example.

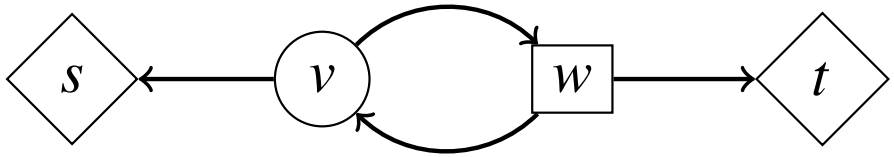

Equation system for valuation of Player 0: $X_v = s + X_w$ and $X_w = t \cdot X_v$

Reachability versus safety games

Example.

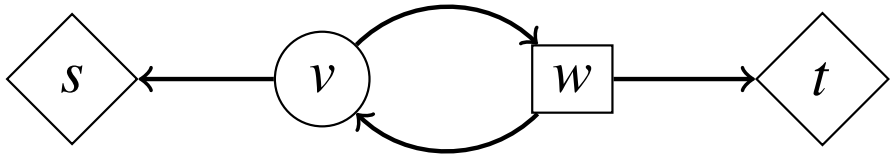

Equation system for valuation of Player 0: $X_v = s + X_w$ and $X_w = t \cdot X_v$

Least fixed point solution in $\mathbb{S}^{\infty}[s,t]$: $f(v) = s$ and $f(w) = st$.

The positional strategy to move always to s absorbs all other strategies!

Reachability versus safety games

Example.

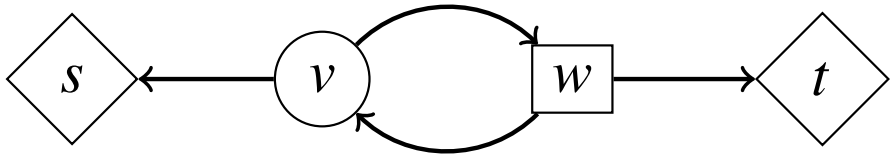

Equation system for valuation of Player 0: $X_v = s + X_w$ and $X_w = t \cdot X_v$

Least fixed point solution in $\mathbb{S}^\infty[s,t]$: $f(v) = s$ and $f(w) = st$.
The positional strategy to move always to s absorbs all other strategies!

But what if we analyse this game as a safety game:
- The value of a non-terminating play is 1, not 0.
- We have to compute greatest fixed-point solutions of the equation system.

Greatest fixed point solution in $\mathbb{S}^\infty[s,t]$: $f(v) = s + t^\infty$ and $f(w) = st + t^\infty$

For safety, Player 0 has two absorptive strategies: move to s, or move to w.
From v the first one admits a unique play with outcome s, the second one admits infinitely many plays with outcome t (and one non-terminating play).

Work in Progress

Provenance analysis for more general infinite games, in particular for parity games.

For such games, it does not suffice to track terminal positions.
Instead track the moves, to get provenance values that tell you which moves are used, and how often, by a strategy.

Hierarchical equations systems, with interleaving least and greatest fixed points, are used to compute provenance values for parity games.

Where are the limits of this approach?

Algorithmic questions