Provenance Analysis for First-Order Model Checking

Val Tannen, University of Pennsylvania

Joint work with Erich Grädel, RWTH Aachen University

Logical Structures in Computation Reunion Workshop

Simons Institute, December 11, 2017

Model Checking $\mathfrak{A} \models \varphi$

Model: $\mathfrak{A}$, a structured collection of info items

Sentence: φ

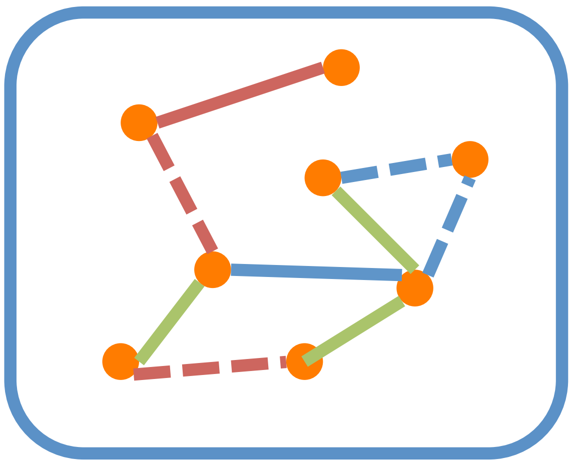

True or False

Which info items in the model are used in checking φ ? (not difficult)

Why (in terms of model info) is φ true ? (alternative reasons?)

How is the model info used to check the truth of φ ? (we will clarify)

These are **provenance** questions.

Application **confidence in** $\mathfrak{A} \models \varphi$

Model: $\mathfrak{A}$, a structured collection of info items

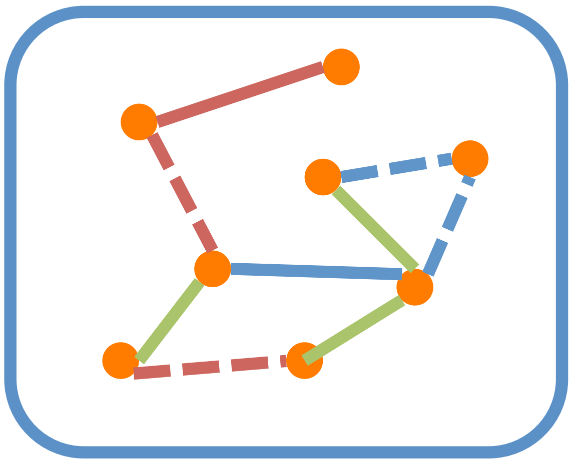

Sentence: φ

True with *confidence* score $\in (0, 1]$

Assuming confidence scores for the info items in the model.

Application **disclosure of** $\mathfrak{A} \models \varphi$

Model: $\mathfrak{A}$, a structured collection of info items

Sentence: φ

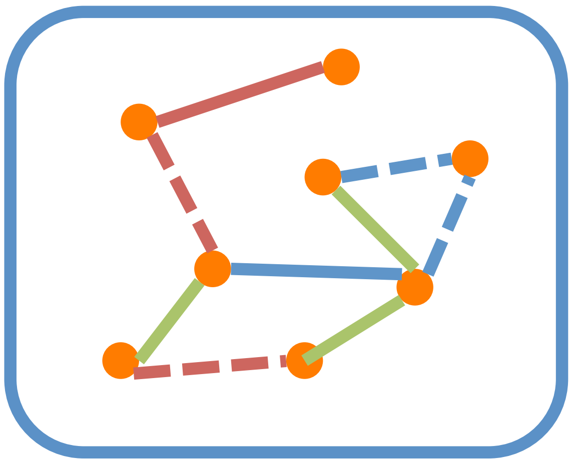

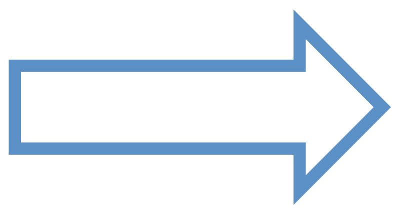

True with access level $\in \{\mathsf{P} < \mathsf{C} < \mathsf{S} < \mathsf{T}\}$

Assuming access levels for the info items in the model.

Application **how many witnesses for** $\mathfrak{A} \models \varphi$

Model: $\mathfrak{A}$, a structured collection of info items

Sentence: φ

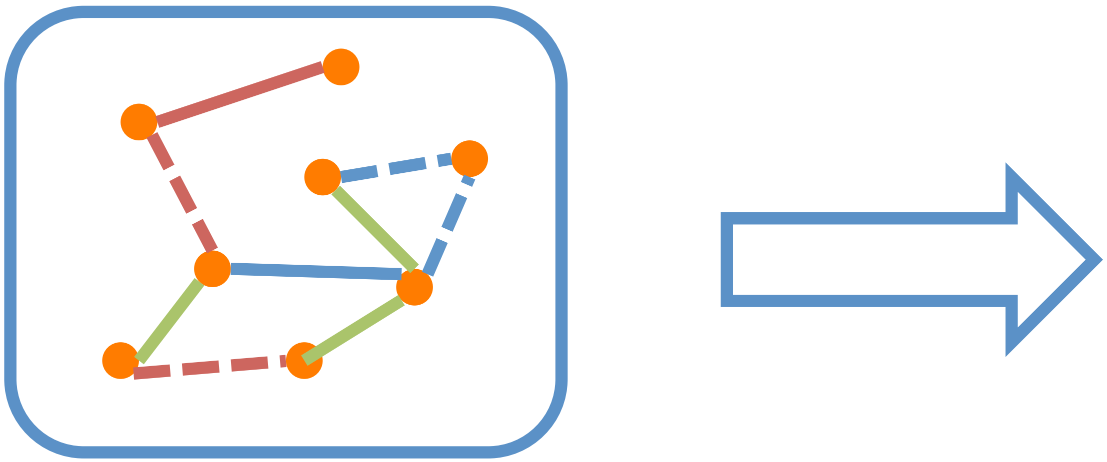

True witnessed by $n > 0$ (model-checking) proof trees

In all three applications we interpret model-checking as *shades of truth* in a specific **commutative semiring**.

Running Example of Model-Checking

In a digraph with edge relation E, the vertex x is "dominant":

$$\mathsf{dominant}(x) \;\equiv\; \forall y\, (x = y) \vee [E(x,y) \wedge \neg E(y,x)]$$

The digraph does not have a dominant vertex: $\varphi \;\equiv\; \forall x\, \neg\mathsf{dominant}(x)$

$$\varphi \;\equiv\; \forall x\, \exists y\, \boxed{\mathsf{denydom}(x,y)} \;\equiv\; \forall x\, \exists y\, \boxed{(x \neq y) \wedge [\neg E(x,y) \vee E(y,x)]} \;\text{ in NNF}$$

Model (digraph) $\mathfrak{A}$:

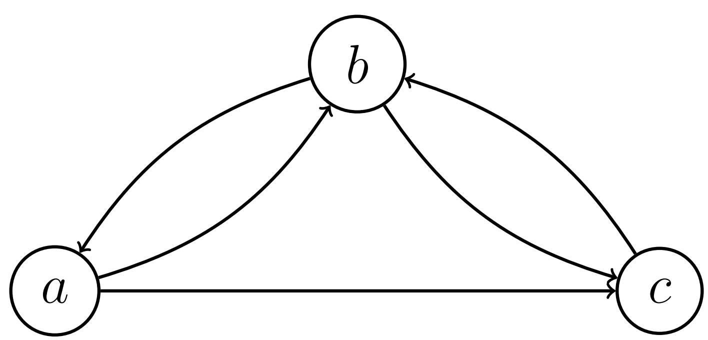

Witnesses for $\mathfrak{A} \models \varphi$ Proof Trees

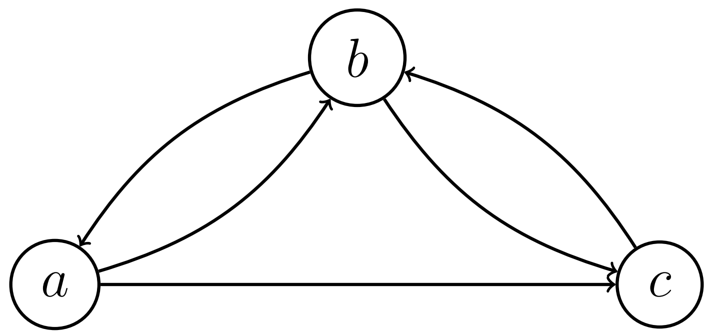

$$\dfrac{\dfrac{\dfrac{a \neq b \qquad \dfrac{E(b,a)}{\neg E(a,b) \vee E(b,a)}}{\mathsf{denydom}(a,b)}}{\exists y\, \mathsf{denydom}(a,y)} \qquad \dfrac{\dfrac{b \neq c \qquad \dfrac{E(c,b)}{\neg E(b,c) \vee E(c,b)}}{\mathsf{denydom}(b,c)}}{\exists y\, \mathsf{denydom}(b,y)} \qquad \dfrac{\dfrac{c \neq a \qquad \dfrac{E(a,c)}{\neg E(c,a) \vee E(a,c)}}{\mathsf{denydom}(c,a)}}{\exists y\, \mathsf{denydom}(c,y)}}{\forall x\, \exists y\, \mathsf{denydom}(x,y)}$$

Outline of the rest of the talk

1. First-order finite model checking interpreted in a commutative semiring.

2. Interpretations in a *provenance semiring*. Dual-indeterminate polynomials for FOL provenance.

3. Provenance tracking assumptions and reverse analysis for first-order models.

4. Missing/wrong answers and integrity constraint failure. Repairs.

Commutative Semirings

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

The semiring is **commutative** when $\cdot$ is commutative.
The semiring is **idempotent** when $+$ is idempotent.
Any distributive lattice is an idempotent commutative semiring.

$+$ interprets alternative use of information from a model.

$\cdot$ interprets joint use of information from a model.

Very roughly speaking:

- $0 \in K$ interprets false assertions.
- $a \in K, a \neq 0$ provides a "nuanced" interpretation for true assertions.

Examples of Commutative Semirings

1. $\mathbb{B} = (\mathbb{B}, \vee, \wedge, \bot, \top)$ is the standard habitat of logical truth.

2. $\mathbb{N} = (\mathbb{N}, +, \cdot, 0, 1)$ is used here for counting proof trees. Also used for *bag semantics* in databases. Not idempotent.

3. $\mathbb{T} = (\mathbb{R}_+^\infty, \min, +, \infty, 0)$, the *tropical* semiring, idempotent but not a distributive lattice. Used in *min-cost* interpretations (e.g., shortest paths).

4. $\mathbb{V} = ([0,1], \max, \cdot, 0, 1)$ the *Viterbi* semiring, isomorphic to $\mathbb{T}$ via $x \mapsto e^{-x}$ and $y \mapsto -\ln y$. Habitat for maximum likelihood trajectory calculations in HMM, also invoked in "possibilistic" uncertainty. Used here for *confidence scores*.

More Examples of Commutative Semirings

5. $\mathbb{A} = (\{\mathsf{P} < \mathsf{C} < \mathsf{S} < \mathsf{T} < 0\}, \min, \max, 0, \mathsf{P})$ is the *access control* semiring.

   P is "public" S is "secret"
   C is "confidential" T is "top secret"
   0 is "so secret that nobody can access it!"

   This is a distributive lattice (beware! the lattice order is the opposite of the one we used in the definition).

6. $\mathbb{F} = ([0,1], \max, \min, 0, 1)$, is called the *fuzzy* semiring. It is a distributive lattice.

One Commutative Semiring to Rule Them All

7. $\mathbb{N}[X] = (\mathbb{N}[X], +, \cdot, 0, 1)$

multivariate polynomials in indeterminates from X
and with coefficients from $\mathbb{N}$.

This is the commutative semiring **freely generated** by the set X.

It's used for a general form of **provenance** [Green, Karvounarakis & T. PODS'07].
We call the elements of X **provenance tokens**.

Proposition For any commutative semiring K, any $f : X \to K$ extends uniquely to a semiring homomorphism $f^* : \mathbb{N}[X] \to K$.

K-Interpretations (I)

Finite relational vocabulary. Finite set $A \neq \emptyset$ set of *ground values*.

Facts_A all ground relational atoms (facts) $R(\mathbf{a})$.
$\mathsf{NegFacts}_A$ all negated facts $\neg R(\mathbf{a})$.

$\mathsf{Lit}_A = \mathsf{Facts}_A \cup \mathsf{NegFacts}_A$

Definition K**-interpretation** where K commutative semiring:

starts with $\pi : \mathsf{Lit}_A \to K$
and is extended to all formulae/sentences $\pi : \mathrm{FOL} \to K$ as follows:

K-Interpretations (II)

valuation $\nu : \mathsf{Vars} \to A$

$$\pi[\![R(\mathbf{x})]\!]_\nu = \pi(R(\nu(\mathbf{x}))) \qquad \pi[\![\neg R(\mathbf{x})]\!]_\nu = \pi(\neg R(\nu(\mathbf{x})))$$

$$\pi[\![x \ \mathsf{op}\ y]\!]_\nu = \text{if } \nu(x) \ \mathsf{op}\ \nu(y) \text{ then } 1 \text{ else } 0 \qquad \pi[\![\varphi \wedge \psi]\!]_\nu = \pi[\![\varphi]\!]_\nu \cdot \pi[\![\psi]\!]_\nu$$

$$\pi[\![\varphi \vee \psi]\!]_\nu = \pi[\![\varphi]\!]_\nu + \pi[\![\psi]\!]_\nu \qquad \pi[\![\exists x\, \varphi]\!]_\nu = \sum_{a\in A} \pi[\![\varphi]\!]_{\nu[x\mapsto a]}$$

$$\pi[\![\forall x\, \varphi]\!]_\nu = \prod_{a\in A} \pi[\![\varphi]\!]_{\nu[x\mapsto a]} \qquad \pi[\![\neg\varphi]\!]_\nu = \pi[\![\mathsf{nnf}(\neg\varphi)]\!]_\nu$$

The symbol op stands for either $=$ or $\neq$.

Proposition It suffices to consider formulae in NNF: $\pi[\![\varphi]\!]_\nu = \pi[\![\mathsf{nnf}(\varphi)]\!]_\nu$.

Indeed…

Let $\mathfrak{A}$ be a finite FO model with universe A.

Define $\pi_{\mathfrak{A}} : \mathsf{Lit}_A \to \mathbb{B}$:

$$\pi_{\mathfrak{A}}(L) \;=\; \top \quad \text{iff} \quad \mathfrak{A} \models L$$

Proposition For any FO sentence φ

$$\pi_{\mathfrak{A}}[\![\varphi]\!] \;=\; \top \quad \text{iff} \quad \mathfrak{A} \models \varphi$$

Define $\pi_{\#\mathfrak{A}} : \mathsf{Lit}_A \to \mathbb{N}$:

$$\pi_{\#\mathfrak{A}}(L) \;=\; \begin{cases} 1 & \text{if } \mathfrak{A} \models L \\ 0 & \text{otherwise} \end{cases}$$

Proposition For any FO sentence φ, $\pi_{\#\mathfrak{A}}[\![\varphi]\!]$ is the number of (model-checking) proof trees that witness $\mathfrak{A} \models \varphi$.

A Provenance-Tracking Interpretation (I)

Previous example plus *annotation* of the edges:

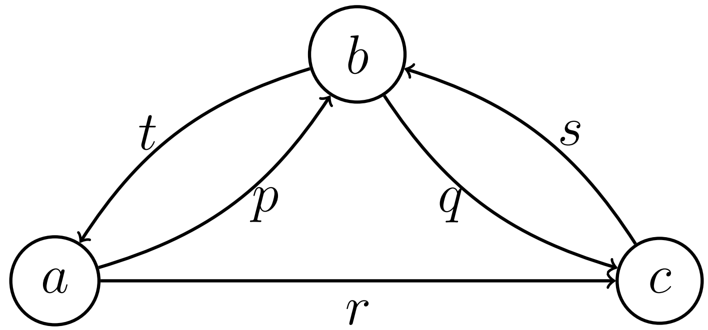

$X = \{r, s, t\}$ is a set of *provenance tokens*. Define $\pi : \mathsf{Lit}_A \to \mathbb{N}[X]$:

$$\pi(L) \;=\; \begin{cases} p & \text{if } L = E(a,b) \\ q & \text{if } L = E(b,c) \\ r & \text{if } L = E(a,c) \\ s & \text{if } L = E(c,b) \\ t & \text{if } L = E(b,a) \end{cases} \qquad = \; \begin{cases} 1 & \text{if } L = \neg E(c,a) \\ 0 & \text{otherwise} \end{cases}$$

Annotation is 1: assume always available without tracking!

A Provenance-Tracking Interpretation (II)

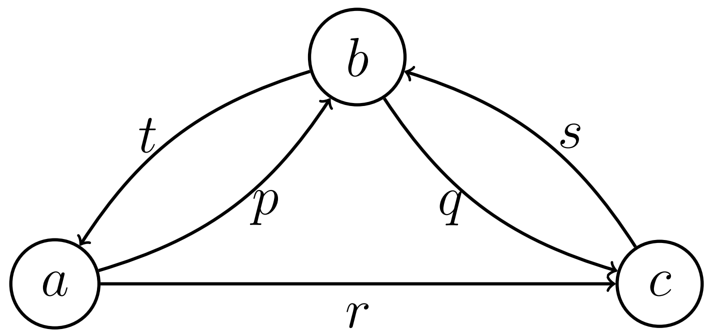

Compute $\pi[\![\forall x\, \neg \mathsf{dominant}(x)]\!] \;=\; \pi[\![\forall x\, \exists y\, (x \neq y) \wedge [\neg E(x,y) \vee E(y,x)]]\!] =$

$$= \;(0 + (0 + t) + (0 + 0)) \;\cdot\; ((0 + p) + 0 + (0 + s)) \;\cdot\; ((1 + r) + (0 + q) + 0)$$

$$t \;\cdot\; (p + s) \;\cdot\; (1 + r + q) \;=\; \boxed{pt + st + prt + rst + pqt + qst}$$

monomials $\sim$ proof trees that witnesses $\mathfrak{A} \models \varphi$. We saw rst before.
$\neg E(c,a)$ also holds, used in two proof trees, but we don't track it.
Difficulties tracking tokens through contradictions!

Positive and Negative Provenance Tokens

Use X to annotate Facts_A. Use $\bar{X}$ for $\mathsf{NegFacts}_A$. $\bar{X} \cap X = \emptyset$.
One-to-one correspondence $X \longleftrightarrow \bar{X}$; $p \longleftrightarrow \bar{p}$ *complementary* tokens.

Define $\mathbb{N}[X, \bar{X}]$ as the quotient of $\mathbb{N}[X \cup \bar{X}]$ by the congruence generated by the equalities $\boxed{p \cdot \bar{p} = 0}$.

Subset of the polynomials in $\mathbb{N}[X \cup \bar{X}]$, namely those such that no monomial contains complementary tokens: **dual(-indeterminate) polynomials**.

The following is the universality property of this construction:

Proposition For any commutative semiring K, any $f : X \cup \bar{X} \to K$ such that $\forall p \in X$ $\boxed{f(p) \cdot f(\bar{p}) = 0}$ extends uniquely to a semiring homomorphism $f^* : \mathbb{N}[X, \bar{X}] \to K$.

A Better Interpretation (I)

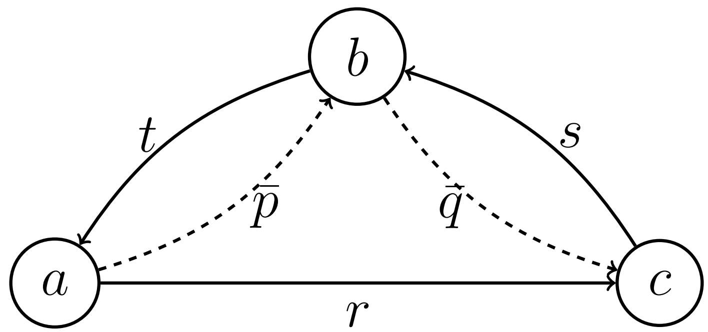

Define $\pi : \mathsf{Lit}_A \to \mathbb{N}[X, \bar{X}]$:

$$\pi(L) = \begin{cases} 0 & \text{if } L = E(a,b) \\ \bar{p} & \text{if } L = \neg E(a,b) \\ 0 & \text{if } L = E(b,c) \\ \bar{q} & \text{if } L = \neg E(b,c) \\ r & \text{if } L = E(a,c) \\ 0 & \text{if } L = \neg E(a,c) \end{cases} = \begin{cases} s & \text{if } L = E(c,b) \\ 0 & \text{if } L = \neg E(c,b) \\ t & \text{if } L = E(b,a) \\ 0 & \text{if } L = \neg E(b,a) \\ 0 & \text{for the other positive facts} \\ 1 & \text{for the other negative facts} \end{cases}$$

A Better Interpretation (II)

Compute $\pi[\![\forall x\, \neg \mathsf{dominant}(x)]\!] \;=\; \pi[\![\forall x\, \exists y\, (x \neq y) \wedge [\neg E(x,y) \vee E(y,x)]]\!] =$

$$= \;(0 + (\bar{p} + t) + (0+0)) \;\cdot\; ((0+0) + 0 + (\bar{q} + s)) \;\cdot\; ((1+r) + (0+0) + 0)$$

$$= \;(\bar{p} + t) \;\cdot\; (\bar{q} + s) \;\cdot\; (1+r) \;=\; \boxed{\bar{p}\bar{q} + \bar{p}s + \bar{q}t + st + \bar{p}\bar{q}r + \bar{p}rs + \bar{q}rt + rst}$$

Again monomials correspond to proof trees that witness $\mathfrak{A} \models \varphi$.

Finally, we can track the *provenance of negative facts*.

This interpretation defines a unique model. It is not “flexible” enough finding other models with desirable properties.

Multi-Model Interpretations

Definition An interpretation $\pi : \mathsf{Lit}_A \to \mathbb{N}[X, \bar{X}]$ is **model-compatible** if for any fact $R(\mathbf{a})$ one of the following three holds:

1. $\exists x \in X$ s.t. $\pi(R(\mathbf{a})) = x$ and $\pi(\neg R(\mathbf{a})) = \bar{x}$, or
2. $\pi(R(\mathbf{a})) = 0$ and $\pi(\neg R(\mathbf{a})) = 1$, or
3. $\pi(R(\mathbf{a})) = 1$ and $\pi(\neg R(\mathbf{a})) = 0$

Specification of **provenance tracking assumptions**.

Such π is "compatible" with at least one model (hence the name), but, in general, with *multiple* models.

This is not a bug but a feature (!) that supports *reverse provenance analysis* as well as *model update*.

Example of Provenance Tracking Assumptions

Define $\pi : \mathsf{Lit}_A \to \mathbb{N}[X, \bar{X}]$:

$$\pi(L) \;=\; \begin{cases} p & \text{if } L = E(a,b) \\ \bar{p} & \text{if } L = \neg E(a,b) \\ q & \text{if } L = E(b,c) \\ \bar{q} & \text{if } L = \neg E(b,c) \\ r & \text{if } L = E(a,c) \\ \bar{r} & \text{if } L = \neg E(a,c) \end{cases} \;=\; \begin{cases} s & \text{if } L = E(c,b) \\ \bar{s} & \text{if } L = \neg E(c,b) \\ t & \text{if } L = E(b,a) \\ \bar{t} & \text{if } L = \neg E(b,a) \\ 0 & \text{for the other positive facts} \\ 1 & \text{for the other negative facts} \end{cases}$$

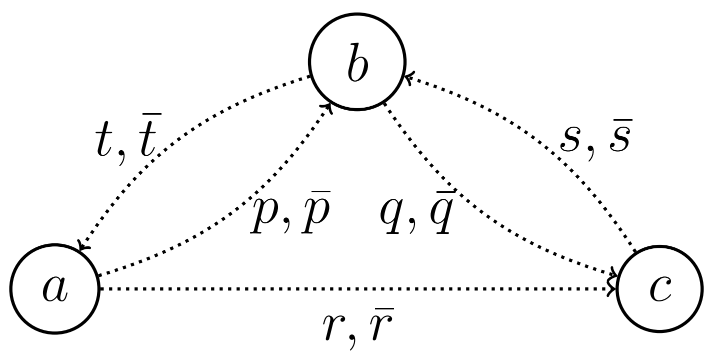

A Multi-Model Polynomial

This π is model-compatible.

Compute $\pi[\![\forall x\, \neg\mathsf{dominant}(x)]\!] \;=\; \pi[\![\forall x\, \exists y\, (x \neq y) \wedge [\neg E(x,y) \vee E(y,x)]]\!] =$

$$= \;(\bar{p} + \bar{r} + t) \cdot (p + \bar{q} + s + \bar{t}) \cdot (1 + q + r + \bar{s})$$

The resulting polynomial has $48 - 4 - 3 - 3 - 4 = 34$ monomials.
It describes the 34 distinct proof trees that witness …what?

Compute $\pi[\![\exists x\, \mathsf{dominant}(x)]\!] \;=\; pr\bar{t} + \bar{p}q\bar{s}t$

Two monomials. They correspond to distinct models!

What a Model-Compatible Interpretation Wants

Definition (again) An interpretation $\pi : \mathsf{Lit}_A \to \mathbb{N}[X, \bar{X}]$ is **truth-compatible** if for any fact $R(\mathbf{a})$ one of the following three holds:

1. $\exists z \in X \cup \bar{X}$ s.t. $\pi(R(\mathbf{a})) = z$ and $\pi(\neg R(\mathbf{a})) = \bar{z}$, or
2. $\pi(R(\mathbf{a})) = 0$ and $\pi(\neg R(\mathbf{a})) = 1$, or
3. $\pi(R(\mathbf{a})) = 1$ and $\pi(\neg R(\mathbf{a})) = 0$

$\mathsf{Must}_\pi = \{L \in \mathsf{Lit}_A \mid \pi(L) = 1\}$

$\mathsf{Mod}_\pi = \{\mathfrak{A} \mid \mathfrak{A} \models \mathsf{Must}_\pi\}$ (When $\mathfrak{A} \in \mathsf{Mod}_\pi$ we say $\mathfrak{A}$ compatible with π.)

$\mathsf{May}_\pi = \{L \in \mathsf{Lit}_A \mid \pi(L) \in X \cup \bar{X}\}$

What Makes It All Work

$\pi : \mathsf{Lit}_A \to \mathbb{N}[X, \bar{X}]$ model-compatible $\quad \varphi \in$ FOL.

Proposition The provenance polynomial $\pi[\![\varphi]\!]$
describes all the proof trees that verify φ using premises from $\mathsf{Must}_\pi \cup \mathsf{May}_\pi$:

Monomial $m\, x_1^{m_1} \cdots x_k^{m_k}$ represents m distinct proof trees
that use m_i times L where $\pi(L) = x_i$.

In particular, the sum of the monomial coefficients in $\pi[\![\varphi]\!]$ counts the number of these proof trees.

Soundness and Completeness of Provenance Tracking

Corollary $\pi : \mathsf{Lit}_A \to \mathbb{N}[X, \bar{X}]$ truth-compatible and φ $\in$FOL. Then,

(i) φ is Mod_π-satisfiable iff $\pi[\![\varphi]\!] \neq 0$, and

(ii) φ is Mod_π-valid iff $\pi[\![\neg\varphi]\!] = 0$

Satisfiability and validity *restricted to the class* Mod_π *of models that agree with some provenance tracking assumptions*. In particular all the models have universe A.

This kind of satisfiability (hence validity) is decidable.

Back to a Single Model

Definition π model-compatible and $\mathfrak{A} \in \mathsf{Mod}_\pi$. The **specialization** of π wrt $\mathfrak{A}$:

$$\pi|_{\mathfrak{A}}(L) \;=\; \begin{cases} \pi(L) & \text{if } \mathfrak{A} \models L \\ 0 & \text{otherwise} \end{cases}$$

Corollary π model-compatible, $\mathfrak{A} \in \mathsf{Mod}_\pi$, φ $\in$FOL s.t. $\mathfrak{A} \models \varphi$.

Then, $\pi|_{\mathfrak{A}}[\![\varphi]\!] \neq 0$ and every monomial in $\pi|_{\mathfrak{A}}[\![\varphi]\!]$ also appears in $\pi[\![\varphi]\!]$, with the same coefficient.

Moreover, $\pi|_{\mathfrak{A}}[\![\varphi]\!]$ describes all the proof trees that witness $\mathfrak{A} \models \varphi$. In particular, the sum of all the monomial coefficients in $\pi|_{\mathfrak{A}}[\![\varphi]\!]$ counts the number of distinct such proof trees.

Model Update

Given $\mathfrak{A}$, update to $\mathfrak{A}'$.

Here is how we update the provenance:

1. Choose model-compatible π such that $\mathfrak{A} \in \mathsf{Mod}_\pi$.
   Make sure you annotate with tokens the literals that you aim to update.

2. Apply the update to π, setting provenance tokens to 0/1. Obtain π'.

3. Compute the specialization $\pi'|_{\mathfrak{A}'}$.

Missing Query Answers [with Jane Xu, Waley Zhang and Abdu Alawini; Penn]

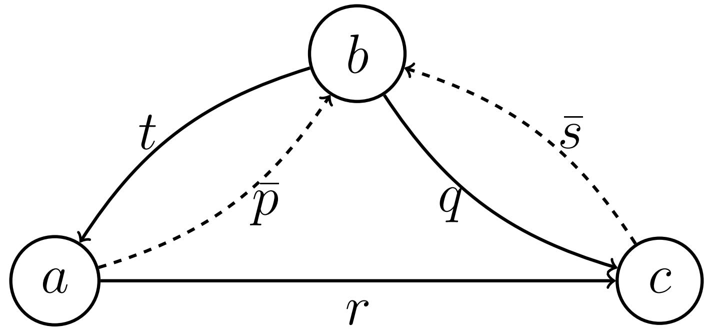

Query: $\quad \mathsf{dominant}(x) \;=\; \forall y\ (x = y) \lor [E(x,y) \land \neg E(y,x)]$

b is an answer for the query; provenance of $\mathsf{dominant}(b)$ is $\bar{p}q\bar{s}t$.

Missing answer: WHY IS a NOT AN ANSWER?

Provenance of $\mathsf{dominant}(a)$ is 0, no help.

Instead, compute the provenance of $\neg\mathsf{dominant}(a)$!

Missing Query Answers: Explanations and Repairs

$$\neg\mathsf{dominant}(a) \;=\; \exists y\,(a \neq y) \wedge [\neg E(a,y) \vee E(y,a)]$$

Has provenance $\bar{p} + t$.

Explanation:

- cause: $\bar{p} \neq 0$ (absence of edge $E(a,b)$)
- another cause: $t \neq 0$ (presence of edge $E(b,a)$)

Repair: $\bar{p} = t = 0$ (insert $E(a,b)$ and delete $E(b,a)$)

(Negative token set to 0: fact insertion.
Positive token set to 0: fact deletion.)

Integrity Constraint Failure [also with Jane Xu, Waley Zhang and Abdu Alawini; Penn]

Change things a bit:

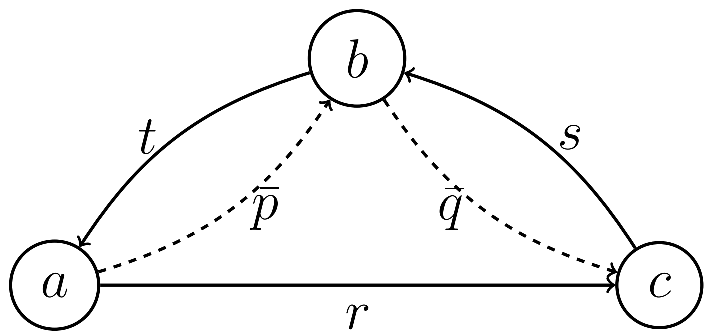

Integrity constraint (IC): "AT LEAST ONE VERTEX IS DOMINANT"

$$\exists x\, \mathsf{dominant}(x)$$

WHY IS THE IC FAILING? Has provenance 0, not helpful.

Compute provenance $\mathfrak{p}$ of $\neg[\exists x\, \mathsf{dominant}(x)]$ then "solve" $\mathfrak{p} = 0$.

$$\mathfrak{p} \;=\; (\bar{p} + t)\cdot(\bar{q} + s)\cdot(1 + r)$$

Integrity Constraint Failure: Repairs and Explanations (I)

and-or tree of solutions to $\quad \mathfrak{p} \ = \ (\bar{p} + t) \cdot (\bar{q} + s) \cdot (1 + r) \ = \ 0 \quad :$

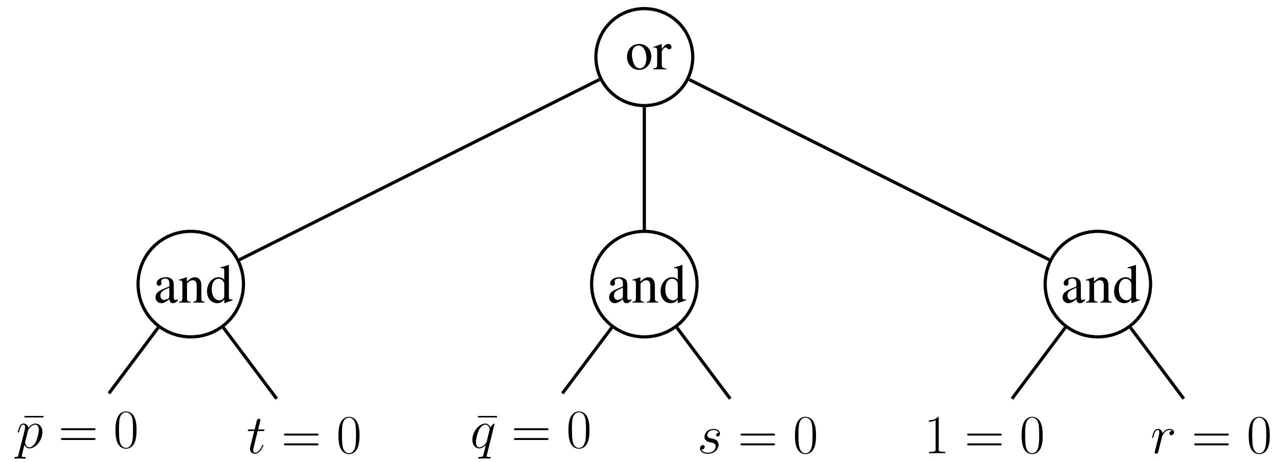

Each solution corresponds to a different **repair**: $\{\bar{p} = t = 0\}$ or $\{\bar{q} = s = 0\}$.

In general, exponential # of minimal repairs
however and-or tree is polysize (data complexity).

Proposition Any minimal repair is a subset of a repair represented in the tree.

Integrity Constraint Failure: Repairs and Explanations (II)

For explanations, *dualize* the tree:

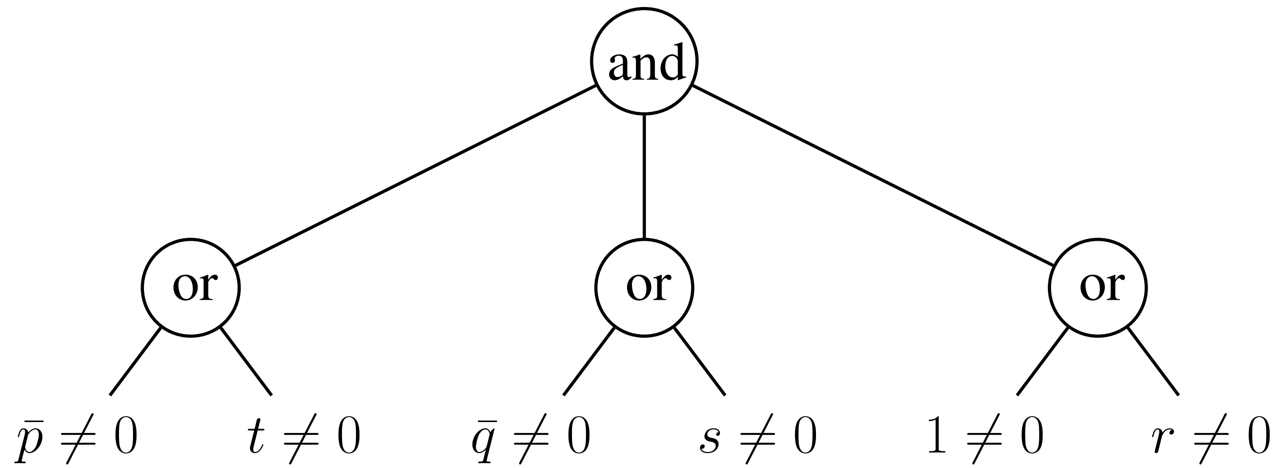

Four minimal **explanations**:

$$\{\bar{p} \neq 0, \bar{q} \neq 0\} \quad \{\bar{p} \neq 0, s \neq 0\} \quad \{t \neq 0, \bar{q} \neq 0\} \quad \{t \neq 0, s \neq 0\}$$

Choose Among Repairs Based on Cost

Update, for each repair, the provenance $\mathfrak{q}$ of IC $\pi[\![\exists x\, \mathsf{dominant}(x)]\!]$
(use a model-compatible interpretation that includes all tokens in all repairs)

$$\mathfrak{q} \;=\; pr\bar{t} + \bar{p}q\bar{s}t$$

Apply each repair (specialize wrt corresponding models):

$\{\bar{p} = t = 0\} \quad \mapsto \quad pr\bar{t}$

$\{\bar{q} = s = 0\} \quad \mapsto \quad \bar{p}q\bar{s}t$

Assumptions: cost of one insertion: α cost of one deletion: β;
Cost of pos/neg facts in the model initially:

$\mathrm{cost}(\bar{p}) = \mathrm{cost}(\bar{q}) = \gamma \qquad \mathrm{cost}(s) = \mathrm{cost}(t) = \delta \qquad \mathrm{cost}(r) = \epsilon$

$\mathrm{cost}(pr\bar{t}) = \alpha + \epsilon + \beta \qquad\qquad \mathrm{cost}(\bar{p}q\bar{s}t) = \gamma + \alpha + \beta + \delta$

If $\epsilon < \gamma + \delta$ the first repair is cheaper.
In general we evaluate polynomials in the *tropical semring* $\mathbb{T}$.

“Semiring Provenance for First-Order Model Checking”, Erich Grädel and Val Tannen, arXiv:1712.01980 [cs.LO], Dec. 2017.

“Provenance Analysis for Missing Answers and Integrity Repairs”, Jane Xu, Waley Zhang, Abdu Alawini, and Val Tannen, submitted.

What’s next?

Extensions to games, and to fixed-point logics, and henceforth to verification logics. Joint work ongoing with Erich Grädel.

Computational question: finding minimal cost repairs. NP-hard problem, looking for approximation techniques.

Other applications (**networks and databases**, workflows, verification). Work ongoing at Penn.