Provenance Analysis for First-Order Model Checking

Val Tannen, University of Pennsylvania

Joint work with Erich Grädel, RWTH Aachen University

And in part with Jane Xu, Waley Zhang, and Abdu Alawini, Penn

UW May 2018

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

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

FO Sentence: φ

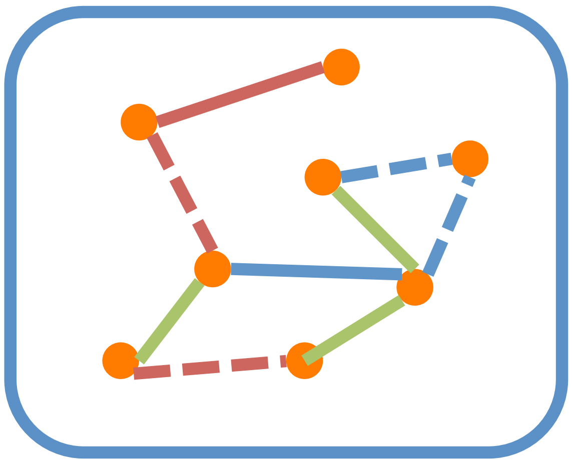

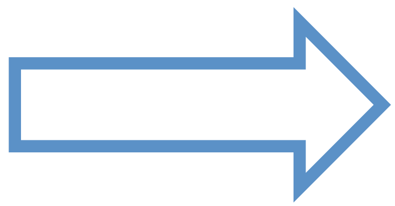

True or False

Provenance analysis of $\mathfrak{A} \models \varphi$ then $\mapsto \begin{cases} \text{cost} \\ \text{confidence} \\ \text{number of witnesses} \end{cases}$

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) $\mathfrak{A}$:

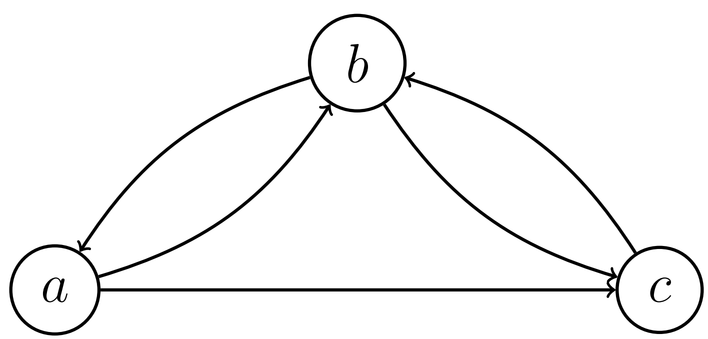

Witnesses for $\mathfrak{A} \models \varphi$: Proof Trees

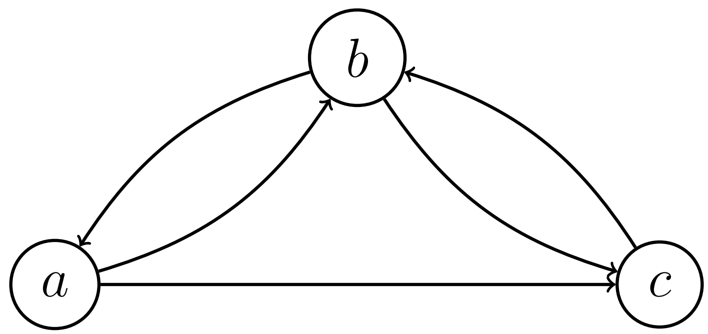

$$\dfrac{\dfrac{\dfrac{a \neq b \qquad \dfrac{E(b,a)}{\neg E(a,b) \vee E(b,a)}}{\mathsf{denydom}(a,b)}}{\exists y\, \mathsf{denydom}(a,y)} \qquad \dfrac{\dfrac{b \neq c \qquad \dfrac{E(c,b)}{\neg E(b,c) \vee E(c,b)}}{\mathsf{denydom}(b,c)}}{\exists y\, \mathsf{denydom}(b,y)} \qquad \dfrac{\dfrac{c \neq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

Non-Standard Semantics for Logical Truth

$(K, +, \cdot, 0, 1)$ commutative semiring

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

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

Examples:

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

   But in general, in $(K, +, \cdot, 0, 1)$:

   - $0 \in K$ interprets false assertions.
   - $a \in K, a \neq 0$ provides a “nuanced” interpretation for true assertions (“shades of truth”!).

2. $\mathbb{N} = (\mathbb{N}, +, \cdot, 0, 1)$ is used here for *counting proof trees*.

   Also used for *bag semantics* of positive queries. (Here *set-bag semantics* (kind of).)

3. $\mathbb{T} = (\mathbb{R}_+^\infty, \min, +, \infty, 0)$, the *tropical* semiring, used here for *cost*.

   Also used for *shortest paths*.

4. $\mathbb{V} = ([0,1], \max, \cdot, 0, 1)$ the *Viterbi* semiring, used here for *confidence scores*.

   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.

K-Interpretations (I)

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

Closed World Assumption.

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

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

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

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

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

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

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 $\mathfrak{A} \models \varphi$.

commutation with homomorphisms

Provenance? One Commutative Semiring to Rule Them All?

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

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 $f^* : \mathbb{N}[X] \to K$.

However, *not* appropriate for provenance of *negation*.

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 Provenance-Tracking Interpretation (I)

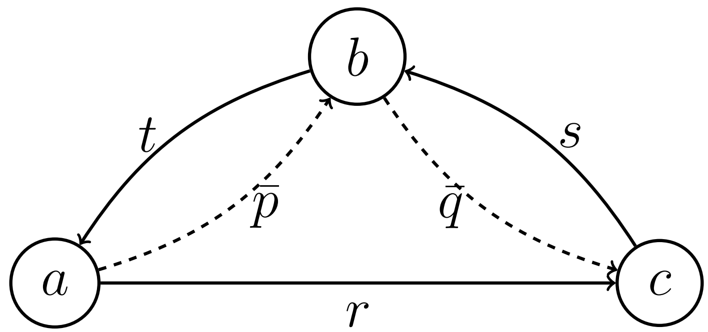

Define $\pi : \mathsf{Lit}_A \to \mathbb{N}[X, \bar{X}]$:

$\pi(E(a,b)) = 0 \qquad \pi(\neg E(a,b)) = \bar{p}$

$\pi(E(b,a)) = t \qquad \pi(\neg E(b,a)) = 0$

etc.

$\pi(L) = 0$ d for the other positive facts

$\pi(L) = 1$ d for the other negative facts

A Provenance-Tracking Interpretation (II)

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

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

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

We can track the **provenance of negative facts**!

This interpretation corresponds to a unique model.

It is not "flexible" enough for finding other models with desirable properties. (*reverse provenance analysis*).

Multi-Model Provenance-Tracking 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 $\pi(\neg R(\mathbf{a})) = 0$

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

$$\mathsf{Mod}_\pi = \{\mathfrak{A} \mid \forall L\ (\pi(L) = 1) \ \Rightarrow\ (\mathfrak{A} \models L)\}.$$

Example of Multi-Model Assumptions

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

$\pi(E(a,b)) = p \qquad \pi(\neg E(a,b)) = \bar{p}, \qquad \pi(E(b,a)) = t \qquad \pi(\neg E(b,a)) = \bar{t},$ etc.

$\pi(L) = 0$ d for the other positive facts
$\pi(L) = 1$ d for the other negative facts

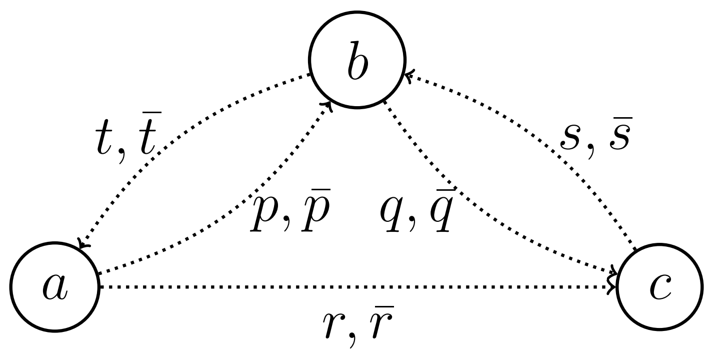

A Multi-Model Polynomial

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 the model checking.

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

Two monomials. They correspond to distinct models!

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 $L \in X \cup \bar{X} \cup \{1\}$:

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 $\pi[\![\neg\varphi]\!] = 0$

This kind of satisfiability/validity is decidable (compute the polynomial!).

What about Trakhtenbrot's Theorem?!

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

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

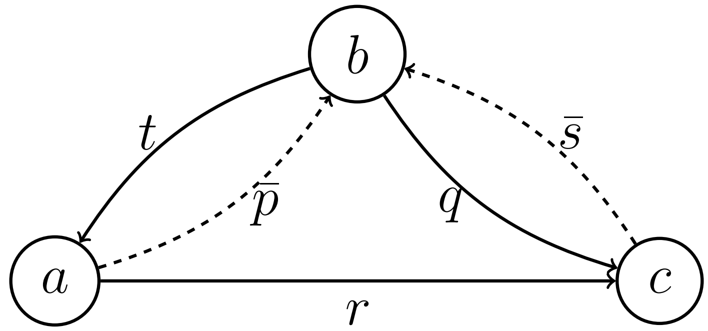

Query: $\quad \mathsf{dominant}(x) \;=\; \forall y\,(x = y) \vee [E(x,y) \wedge \neg E(y,x)]$

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 $\quad \bar{p} + t$.

Explanation:

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

Repair: $\quad \bar{p} = t = 0$ (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:

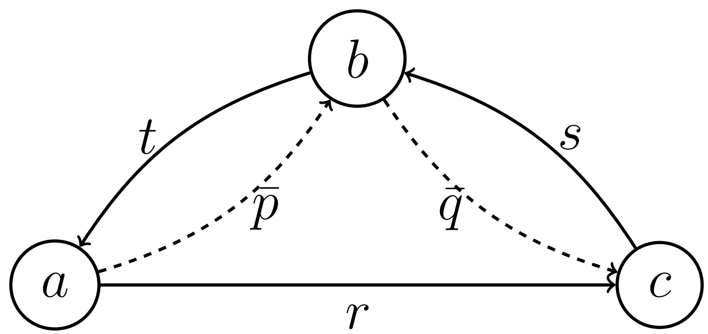

Integrity constraint (IC): "AT LEAST ONE VERTEX IS DOMINANT"

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

WHY IS THE IC FAILING? Itself it has provenance 0, not helpful.

Compute provenance $\mathfrak{p}$ of $\neg[\exists x\, \mathsf{dominant}(x)]$:

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

Integrity Constraint Failure: Explanations

Provenance of $\neg$IC is:

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

8 alternative explanations but 4 of them are “redundant”. We are left with:

- cause: $\bar{p}\bar{q}$ (absence of $E(a,b)$ and $E(b,c)$)
- alt-cause: $\bar{p}s$ (absence of $E(a,b)$ and presence of $E(c,b)$)
- alt-cause: $t\bar{q}$ (presence of $E(b,a)$ and absence of $E(b,c)$)
- alt-cause: ts (presence of $E(b,a)$ and presence of $E(c,b)$)

Integrity Constraint Failure: Repairs

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

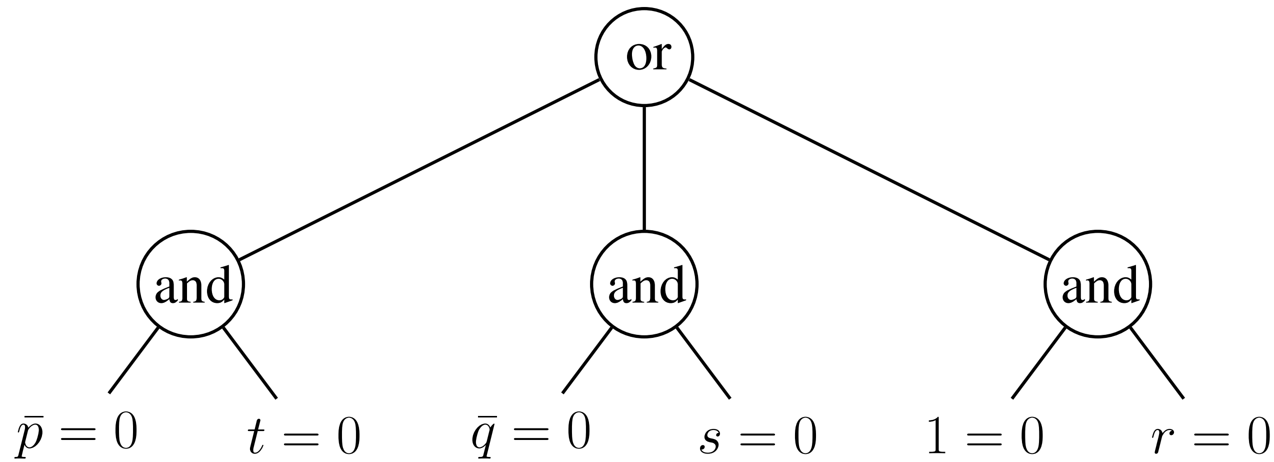

Each solution corresponds to a different **repair**: $\{\bar{p} = t = 0\}$ or $\{\bar{q} = s = 0\}$.

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.

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):

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

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

Evaluate polynomials in the *tropical semiring* $\mathbb{T}$.

Assumptions: cost of one insertion: 20 cost of one deletion: 15;
Cost of pos/neg facts in the model initially:

$\mathrm{cost}(\bar{p}) = \mathrm{cost}(\bar{q}) = 10 \qquad \mathrm{cost}(s) = \mathrm{cost}(t) = 5 \qquad \mathrm{cost}(r) = 10$

$\mathrm{cost}(pr\bar{t}) = 20 + 10 + 15 = 45 \qquad \mathrm{cost}(\bar{p}q\bar{s}t) = 10 + 20 + 15 + 5 = 50$

The first repair is cheaper.

"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, To appear in Data Eng. Bulletin.

What's next?

OWA.

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.