

Higher-Order Approximate Relational Refinement Types for Mechanism Design and Differential Privacy

Gilles Barthe, Marco Gaboardi, Emilio Jesús Gallego Arias,
Justin Hsu, Aaron Roth, Pierre-Yves Strub

UPenn-Mines ParisTech, IMDEA Software Institute, Dundee

Gallium Seminar, Nov 17th, 2014

Software Verification

- ▶ Reason *formally* about programs and their behavior.
- ▶ Increase trust in software, help programmers/designers.
- ▶ Has important practical and economical utility.
- ▶ Expressiveness? Automation?

Software Verification

- ▶ Reason *formally* about programs and their behavior.
- ▶ Increase trust in software, help programmers/designers.
- ▶ Has important practical and economical utility.
- ▶ Expressiveness? Automation?

Today:

- ▶ Verification of probabilistic programs.
- ▶ *Mechanisms*: inputs controlled by strategic agents.
- ▶ *Truthfulness*: An agent gets best utility when telling the truth.
- ▶ *Privacy*: An agent's information leak is bounded.

The Main Challenges

Relational Reasoning

Properties of interest are relational, that is, defined over *two runs* of the *same program*:

- ▶ *Truthfulness*: agent telling the truth vs not.
- ▶ *Privacy*: run including the agent vs not.

The Main Challenges

Relational Reasoning

Properties of interest are relational, that is, defined over *two runs* of the *same program*:

- ▶ *Truthfulness*: agent telling the truth vs not.
- ▶ *Privacy*: run including the agent vs not.

Probabilistic Reasoning

Interesting algorithms are randomized, properties rely on:

- ▶ Expected values.
- ▶ *Distance* on distributions.

Our Approach:

Related/Precursor Work:

- ▶ Relational logics.
- ▶ F^* , RF^* .
- ▶ CertiCrypt/CertiPriv.
- ▶ Fuzz/DFuzz.

Our Approach:

Related/Precursor Work:

- ▶ Relational logics.
- ▶ F^* , RF^* .
- ▶ CertiCrypt/CertiPriv.
- ▶ Fuzz/DFuzz.

Our Contributions

- ▶ Extended type system:
 - ▶ Support for Higher-Order refinements.
 - ▶ Embedding of logical relations! DFuzz soundness proof.
 - ▶ Probabilistic approximate types.
- ▶ New application domain and examples.
- ▶ Prototype implementation.

The System: Relational Refinement Types

Variables

Relational variables, $x \in \mathcal{X}_{\mathcal{R}}$; left/right instances $x_{\triangleleft}, x_{\triangleright} \in \mathcal{X}_{\mathcal{R}}^{\times}$.

Expressions

$$\begin{aligned} e^m &::= \mathbf{C} \mid x \in \mathcal{X}^m \mid e \ e \mid \lambda x. e \mid \text{case } e \text{ with } [\epsilon \Rightarrow e \mid x :: x \Rightarrow e] \\ &\mid \text{letrec}^{\uparrow} f \ x = e \mid \text{letrec}^{\downarrow} f \ x = e \\ &\mid e_{\uparrow} \mid \text{let}_{\uparrow} x = e \text{ in } e \mid \text{unit}_{\mathbf{M}} e \mid \text{bind}_{\mathbf{M}} x = e \text{ in } e \end{aligned}$$

Regular Types

$$\begin{aligned} \tilde{\tau}, \tilde{\sigma}, \dots \in \mathbf{CoreTy} &::= \bullet \mid \mathbb{B} \mid \mathbb{N} \mid \overline{\mathbb{R}} \mid \overline{\mathbb{R}}^+ \mid L[\tilde{\tau}] \\ \tau, \sigma, \dots \in \mathbf{Ty} &::= \tilde{\tau} \mid \mathfrak{M}[\tau] \mid \mathfrak{C}[\tau] \mid \tau \rightarrow \sigma \end{aligned}$$


Relational Refinement Types

$$\begin{aligned} T, U \in \mathcal{T} &::= \tilde{\tau} \mid \mathfrak{M}_{\epsilon, \delta}[T] \mid \mathfrak{C}[T] \mid \Pi(x :: T). T \mid \{x :: T \mid \phi\} \\ \phi, \psi \in \mathcal{A} &::= \mathcal{Q}(x : \tau). \phi \mid \mathcal{Q}(x :: T). \phi \\ &\mid \mathcal{C}(\phi_1, \dots, \phi_n) \mid e^{\times} = e^{\times} \mid e^{\times} \leq e^{\times} \\ \mathcal{C} &= \{\top/0, \perp/0, \neg/1, \vee/2, \wedge/2, \Rightarrow/2\} \end{aligned}$$

Relational Refinement Types: Example

Regular refinement types not enough to capture some properties.


k-sensitive function



Relational Refinement Types: Example

Regular refinement types not enough to capture some properties.


k-sensitive function



Relational Refinement Types: Example

Regular refinement types not enough to capture some properties.


k-sensitive function



Relational Refinement Types: Example

Regular refinement types no enough to capture some properties.

k-sensitive function




$$\forall x_1, x_2. |f(x_1) - f(x_2)| \leq k \cdot |x_1 - x_2|$$

Relational Refinement Types: Example

Regular refinement types no enough to capture some properties.

k-sensitive function



$$\forall x_1, x_2. |f(x_1) - f(x_2)| \leq k \cdot |x_1 - x_2|$$

What should the type for *f* be?

Relational Refinement Types: Example

For the property:

$$\forall x_1, x_2. |f(x_1) - f(x_2)| \leq k \cdot |x_1 - x_2|$$

Relational Refinement Types: Example

For the property:

$$\forall x_1, x_2. |f(x_1) - f(x_2)| \leq k \cdot |x_1 - x_2|$$

we can do a refinement at a higher type:

$$\{f : \mathbb{R} \rightarrow \mathbb{R} \mid \forall x :: \mathbb{R}. |f(x_\triangleleft) - f(x_\triangleright)| \leq k \cdot |x_\triangleleft - x_\triangleright|\}$$

Relational Refinement Types: Example

For the property:

$$\forall x_1, x_2. |f(x_1) - f(x_2)| \leq k \cdot |x_1 - x_2|$$

we can do a refinement at a higher type:

$$\{f : \mathbb{R} \rightarrow \mathbb{R} \mid \forall x :: \mathbb{R}. |f(x_{\triangleleft}) - f(x_{\triangleright})| \leq k \cdot |x_{\triangleleft} - x_{\triangleright}|\}$$

or we can refer to two copies of the input:

$$f : \Pi(x :: \mathbb{R}). \{r :: \mathbb{R} \mid k \cdot |r_{\triangleleft} - r_{\triangleright}| \leq |x_{\triangleleft} - x_{\triangleright}|\}$$

Both types are equivalent in our system, but the pre/post style more convenient for reasoning.

The System: Semantics

Semantic subtyping for non-relational types:

$$\frac{\vdash e : T \quad \Gamma \models \phi[x/e]}{\vdash e : \{x : T \mid \phi\}}$$

The System: Semantics

Semantic subtyping for non-relational types:

$$\frac{\vdash e : T \quad \Gamma \models \phi[x/e]}{\vdash e : \{x : T \mid \phi\}} \quad \vdash e : T \Rightarrow e \in \llbracket T \rrbracket$$

The System: Semantics

Semantic subtyping for non-relational types:

$$\frac{\vdash e : T \quad \Gamma \models \phi[x/e]}{\vdash e : \{x : T \mid \phi\}} \quad \vdash e : T \Rightarrow e \in \llbracket T \rrbracket \quad \frac{v \in \llbracket T \rrbracket \quad \models \phi(v)}{v \in \llbracket \{x : T \mid \phi(x)\} \rrbracket}$$

The System: Semantics

Semantic subtyping for non-relational types:

$$\frac{\vdash e : T \quad \Gamma \models \phi[x/e]}{\vdash e : \{x : T \mid \phi\}} \quad \vdash e : T \Rightarrow e \in \llbracket T \rrbracket \quad \frac{v \in \llbracket T \rrbracket \quad \models \phi(v)}{v \in \llbracket \{x : T \mid \phi(x)\} \rrbracket}$$

Semantic subtyping for HO relational types:

$$\langle T \rangle_\theta \subseteq \llbracket T \rrbracket \times \llbracket T \rrbracket$$

$$\frac{(d_1, d_2) \in \llbracket T \rrbracket \times \llbracket T \rrbracket}{(d_1, d_2) \in \langle T \rangle_\theta} \quad \frac{(d_1, d_2) \in \langle T \rangle_\theta \quad \llbracket \phi \rrbracket_{\theta \left\{ \begin{smallmatrix} x_{\triangleleft} \mapsto d_1 \\ x_{\triangleright} \mapsto d_2 \end{smallmatrix} \right\}}}{(d_1, d_2) \in \langle \{x :: T \mid \phi\} \rangle_\theta}$$

$$\frac{(f_1, f_2) \in \llbracket T \rightarrow U \rrbracket \quad \forall (d_1, d_2) \in \langle T \rangle_\theta. (f_1(d_1), f_2(d_2)) \in \langle U \rangle_{\theta \left\{ \begin{smallmatrix} x_{\triangleleft} \mapsto d_1 \\ x_{\triangleright} \mapsto d_2 \end{smallmatrix} \right\}}}{(f_1, f_2) \in \langle \Pi(x :: T). U \rangle_\theta}$$

SubTyping

$$\text{SUB-REFL} \frac{\mathcal{G} \vdash T}{\mathcal{G} \vdash T \preceq T}$$

$$\text{SUB-TRANS} \frac{\mathcal{G} \vdash T \preceq U \quad \mathcal{G} \vdash U \preceq V}{\mathcal{G} \vdash T \preceq V}$$

$$\text{SUB-LEFT} \frac{\mathcal{G} \vdash \{x :: T \mid \phi\}}{\mathcal{G} \vdash \{x :: T \mid \phi\} \preceq T}$$

$$\text{SUB-RIGHT} \frac{\mathcal{G} \vdash T \preceq U \quad \|\mathcal{G}, x :: U\| \vdash \phi \quad \forall \theta. \theta \vdash \mathcal{G}, x :: T \Rightarrow \llbracket \phi \rrbracket_{\theta}}{\mathcal{G} \vdash T \preceq \{x :: U \mid \phi\}}$$

$$\text{SUB-PROD} \frac{\mathcal{G} \vdash T_2 \preceq T_1 \quad \mathcal{G}, x :: T_2 \vdash U_1 \preceq U_2}{\mathcal{G} \vdash \Pi(x :: T_1). U_1 \preceq \Pi(x :: T_2). U_2}$$

The System: Typing

The typing judgment relates two programs to a type:

$$\mathcal{G} \vdash e_1 \sim e_2 :: T$$

The System: Typing

The typing judgment relates two programs to a type:

$$\mathcal{G} \vdash e_1 \sim e_2 :: T$$

Soundness

$$\mathcal{G} \vdash e_1 \sim e_2 :: T \Rightarrow \forall \mathcal{G} \vdash \theta, ([e_1]_\theta, [e_2]_\theta) \in (T)_\theta$$

The System: Typing

The typing judgment relates two programs to a type:

$$\mathcal{G} \vdash e_1 \sim e_2 :: T$$

Soundness

$$\mathcal{G} \vdash e_1 \sim e_2 :: T \Rightarrow \forall \mathcal{G} \vdash \theta, ([e_1]_{\theta}, [e_2]_{\theta}) \in (T)_{\theta}$$

Synchronicity

In most cases programs are synchronous, so we use:

$$\mathcal{G} \vdash e :: T \equiv \mathcal{G} \vdash e_{\triangleleft} \sim e_{\triangleright} :: T$$

with $e_{\triangleleft}, e_{\triangleright}$ projecting the variables in e .

Base Typing Rules

$$\text{VAR} \frac{x :: T \in \text{dom}(\mathcal{G})}{\mathcal{G} \vdash x :: T}$$

$$\text{ABS} \frac{\mathcal{G}, x :: T \vdash e :: U}{\mathcal{G} \vdash \lambda x. e :: \Pi(x :: T). U}$$

$$\text{APP} \frac{\mathcal{G} \vdash e_f :: \Pi(x :: T). U \quad \mathcal{G} \vdash e_a :: T}{\mathcal{G} \vdash e_f e_a :: U\{x \mapsto e_a\}}$$

Base Typing Rules

$$\text{VAR} \frac{x :: T \in \text{dom}(\mathcal{G})}{\mathcal{G} \vdash x :: T}$$

$$\text{ABS} \frac{\mathcal{G}, x :: T \vdash e :: U}{\mathcal{G} \vdash \lambda x. e :: \Pi(x :: T). U}$$

$$\text{APP} \frac{\mathcal{G} \vdash e_f :: \Pi(x :: T). U \quad \mathcal{G} \vdash e_a :: T}{\mathcal{G} \vdash e_f e_a :: U\{x \mapsto e_a\}}$$

$$\text{CASE} \frac{\begin{array}{l} \mathcal{G} \vdash e :: L[\tilde{\tau}] \quad \forall \theta. \theta \vdash \mathcal{G} \Rightarrow \text{skeleton}(e_{\triangleleft}, e_{\triangleright}) \\ \mathcal{G}, \{e_{\triangleleft} = e_{\triangleright} = \epsilon\} \vdash e_1 :: T \\ \mathcal{G}, x :: \tilde{\tau}, y :: L[\tilde{\tau}], \{e_{\triangleleft} = x_{\triangleleft} :: y_{\triangleleft} \wedge e_{\triangleright} = x_{\triangleright} :: y_{\triangleright}\} \vdash e_2 :: T \end{array}}{\mathcal{G} \vdash \text{case } e \text{ with } [\epsilon \Rightarrow e_1 \mid x :: y \Rightarrow e_2] :: T}$$

Typing Rules for Recursion

To ensure consistency at higher-types, we must embed non-terminating computations in the partiality monad:

$$\text{LETRECSN} \frac{\mathcal{G}, f :: \Pi(x :: T). U \vdash \lambda x. e :: \Pi(x :: T). U \quad \mathcal{G} \vdash \Pi(x :: T). U \quad \mathcal{SN}\text{-guard}}{\mathcal{G} \vdash \text{letrec}^\downarrow f x = e :: \Pi(x :: T). U}$$

Typing Rules for Recursion

To ensure consistency at higher-types, we must embed non-terminating computations in the partiality monad:

$$\text{LETRECSN} \frac{\mathcal{G}, f :: \Pi(x :: T). U \vdash \lambda x. e :: \Pi(x :: T). U \quad \mathcal{G} \vdash \Pi(x :: T). U \quad \text{SN-guard}}{\mathcal{G} \vdash \text{letrec}^\downarrow f x = e :: \Pi(x :: T). U}$$

$$\text{LETREC} \frac{\mathcal{G} \vdash \Pi(x :: T). \mathfrak{C}[U] \quad \mathcal{G}, f :: \Pi(x :: T). \mathfrak{C}[U] \vdash \lambda x. e :: \Pi(x :: T). \mathfrak{C}[U]}{\mathcal{G} \vdash \text{letrec } f x = e :: \Pi(x :: T). \mathfrak{C}[U]}$$

Typing Rules for Recursion

To ensure consistency at higher-types, we must embed non-terminating computations in the partiality monad:

$$\text{LETRECSN} \frac{\mathcal{G}, f :: \Pi(x :: T). U \vdash \lambda x. e :: \Pi(x :: T). U \quad \mathcal{G} \vdash \Pi(x :: T). U \quad \text{SN-guard}}{\mathcal{G} \vdash \text{letrec}^\downarrow f x = e :: \Pi(x :: T). U}$$

$$\text{LETREC} \frac{\mathcal{G} \vdash \Pi(x :: T). \mathfrak{C}[U] \quad \mathcal{G}, f :: \Pi(x :: T). \mathfrak{C}[U] \vdash \lambda x. e :: \Pi(x :: T). \mathfrak{C}[U]}{\mathcal{G} \vdash \text{letrec } f x = e :: \Pi(x :: T). \mathfrak{C}[U]}$$

$$\text{UNITC} \frac{\mathcal{G} \vdash e :: T}{\mathcal{G} \vdash e_\uparrow :: \mathfrak{C}[T]} \quad \text{BINDC} \frac{\mathcal{G} \vdash e_1 :: \mathfrak{C}[T_1] \quad \mathcal{G} \vdash \mathfrak{C}[T_2] \quad \mathcal{G}, x :: T_1 \vdash e_2 :: \mathfrak{C}[T_2]}{\mathcal{G} \vdash \text{let}_\uparrow x = e_1 \text{ in } e_2 :: \mathfrak{C}[T_2]}$$

Asynchronous Rules

$$\text{ASYM} \frac{\mathcal{G} \vdash e_1 \sim e_2 :: T}{\mathcal{G}^{\leftrightarrow} \vdash e_2^{\leftrightarrow} \sim e_1^{\leftrightarrow} :: T^{\leftrightarrow}}$$

$$\text{AREDLLEFT} \frac{e_1 \rightarrow e'_1 \quad \mathcal{G} \vdash e_1 \sim e_2 :: T}{\mathcal{G} \vdash e'_1 \sim e_2 :: T}$$

Asynchronous Rules

$$\text{ASYM} \frac{\mathcal{G} \vdash e_1 \sim e_2 :: T}{\mathcal{G}^{\leftrightarrow} \vdash e_2^{\leftrightarrow} \sim e_1^{\leftrightarrow} :: T^{\leftrightarrow}}$$

$$\text{AREDLLEFT} \frac{e_1 \rightarrow e'_1 \quad \mathcal{G} \vdash e_1 \sim e_2 :: T}{\mathcal{G} \vdash e'_1 \sim e_2 :: T}$$

$$\text{ACASE} \frac{\begin{array}{l} |\mathcal{G}| \vdash e : L[\tilde{\tau}] \quad |\mathcal{G}| \vdash e' : |T| \\ \mathcal{G}, \{e_{\triangleleft} = \epsilon\} \vdash e_1 \sim e' :: T \\ \mathcal{G}, x :: \tilde{\tau}, y :: L[\tilde{\tau}], \{e_{\triangleleft} = x_{\triangleleft} :: y_{\triangleleft}\} \vdash e_2 \sim e' :: T \end{array}}{\mathcal{G} \vdash \text{case } e \text{ with } [\epsilon \Rightarrow e_1 \mid x :: y \Rightarrow e_2] \sim e' :: T}$$

More on Mechanism Design

Mechanism design is the study of algorithm design where the inputs to the algorithm are controlled by strategic agents, who must be *incentivized* to faithfully report them.

More on Mechanism Design

Mechanism design is the study of algorithm design where the inputs to the algorithm are controlled by strategic agents, who must be *incentivized* to faithfully report them.

Formally

- ▶ n agents, with type for actions $A_i, i \in \{1, \dots, n\}$.
- ▶ A mechanism $M : A^n \rightarrow \mathcal{O}$.
- ▶ A payoff for every agent $P_i : \mathcal{O} \rightarrow R^+$.
- ▶ **Probabilistic algorithms are common!**
Payoff becomes *expected payoff*.

More on Mechanism Design

Mechanism design is the study of algorithm design where the inputs to the algorithm are controlled by strategic agents, who must be *incentivized* to faithfully report them.

Formally

- ▶ n agents, with type for actions $A_i, i \in \{1, \dots, n\}$.
- ▶ A mechanism $M : A^n \rightarrow \mathcal{O}$.
- ▶ A payoff for every agent $P_i : \mathcal{O} \rightarrow R^+$.
- ▶ **Probabilistic algorithms are common!**
Payoff becomes *expected payoff*.

Verification

Incentives are not enough, *the agents need to believe them*.

Verification is an attractive way to convince them.

Mechanism Examples

Auctions

- ▶ Buyers (agents), *bids* (actions), seller (mechanism).
- ▶ Outcome: price, goods assignment.
- ▶ An auction is *truthful* if the buyer gets maximal payoff when she reports her true valuation.

Mechanism Examples

Auctions

- ▶ Buyers (agents), *bids* (actions), seller (mechanism).
- ▶ Outcome: price, goods assignment.
- ▶ An auction is *truthful* if the buyer gets maximal payoff when she reports her true valuation.

Nash Equilibrium Computation

- ▶ n players, action type A .
- ▶ Payoff for i , $P_i : A^n \rightarrow R^+$, depends on others actions.
- ▶ The mechanism suggests an *action profile* (a_1, \dots, a_n) .
- ▶ If all the other players follow the suggestion, player i gets the best payoff by following too.

Digital Goods Auctions

- ▶ Price a good with infinite supply. (i.e: Digital goods)

Digital Goods Auctions

- ▶ Price a good with infinite supply. (i.e: Digital goods)
- ▶ Bidders and seller.

Digital Goods Auctions

- ▶ Price a good with infinite supply. (i.e: Digital goods)
- ▶ Bidders and seller.
- ▶ Bidders have a secret *true* value for the item v_j , and make a public *bid* b_j before the price is known.

Digital Goods Auctions

- ▶ Price a good with infinite supply. (i.e: Digital goods)
- ▶ Bidders and seller.
- ▶ Bidders have a secret *true* value for the item v_j , and make a public *bid* b_j before the price is known.
- ▶ The seller knows the bids, but not the real values. Sets the price p after the bids.

Digital Goods Auctions

- ▶ Price a good with infinite supply. (i.e: Digital goods)
- ▶ Bidders and seller.
- ▶ Bidders have a secret *true* value for the item v_i , and make a public *bid* b_i before the price is known.
- ▶ The seller knows the bids, but not the real values. Sets the price p after the bids.
- ▶ If $b_i \geq p$, the bidder i gets the item, with utility $v_i - p$. Otherwise she doesn't get it, and utility is 0.

Digital Goods Auctions

- ▶ Price a good with infinite supply. (i.e: Digital goods)
- ▶ Bidders and seller.
- ▶ Bidders have a secret *true* value for the item v_i , and make a public *bid* b_i before the price is known.
- ▶ The seller knows the bids, but not the real values. Sets the price p after the bids.
- ▶ If $b_i \geq p$, the bidder i gets the item, with utility $v_i - p$. Otherwise she doesn't get it, and utility is 0.

The auction is truthful if buyers have optimal utility when they reports the true value v_i as their bids b_i .

In general, an auction cannot be truthful if it depends on the bidder's price!

The Fixed Price Auction

Fixed Price Auctions

The simplest truthful auction is the *fixed price auction*. The seller will set p independently of the bid b for a seller with true value v . If $b \geq p$, then utility $v - p$, else 0. Note the bad revenue properties.

The Fixed Price Auction

Fixed Price Auctions

The simplest truthful auction is the *fixed price auction*. The seller will set p independently of the bid b for a seller with true value v . If $b \geq p$, then utility $v - p$, else 0. Note the bad revenue properties.

Informal proof of truthfulness

The price p is fixed, we compare $b_{\triangleleft} = v$ vs $b_{\triangleright} \neq v$. The interesting cases are when the bidder gets the item in one run and doesn't in the other:

- ▶ If b_{\triangleright} got the item, utility is negative, thus less than 0 for the b_{\triangleleft} case (remember b_{\triangleleft} didn't get the item).
- ▶ If b_{\triangleleft} got the item, utility will be greater or equal than 0, thus better or equal than b_{\triangleright} 's utility (0).

The Fixed Price Auction

We model the utility as a program:

```
let fp_utility (v : R) {b :: R} (p : R)
  : { u :: R } =
  if b >= p then v - p
  else 0.0
```

The Fixed Price Auction

We model the utility as a program:

```
let fp_utility (v : R) {b :: R | b◁ = v} (p : R)
    : { u :: R | u◁ >= u▷ } =
    if b >= p then v - p
    else 0.0
```

The Fixed Price Auction

We model the utility as a program:

```
let fp_utility (v : R) {b :: R | b◁ = v} (p : R)
      : { u :: R | u◁ >= u▷ } =
  if b >= p then v - p
  else 0.0
```

We use asynchronous reasoning. The interesting case is:

$$\{b_{\triangleleft} = v, b_{\triangleleft} \geq p, b_{\triangleright} < p\} \vdash v - p \sim 0.0 :: \{u :: \mathbb{R} \mid u_{\triangleleft} \geq u_{\triangleright}\}$$

substituting $[v - p/u_{\triangleleft}, 0.0/u_{\triangleright}]$ we get the proof obligation:

The Fixed Price Auction

We model the utility as a program:

```
let fp_utility (v : R) {b :: R | b◁ = v} (p : R)
      : { u :: R | u◁ >= u▷ } =
  if b >= p then v - p
  else 0.0
```

We use asynchronous reasoning. The interesting case is:

$$\{b_{\triangleleft} = v, b_{\triangleleft} \geq p, b_{\triangleright} < p\} \vdash v - p \sim 0.0 :: \{u :: \mathbb{R} \mid u_{\triangleleft} \geq u_{\triangleright}\}$$

substituting $[v - p/u_{\triangleleft}, 0.0/u_{\triangleright}]$ we get the proof obligation:

$$v \geq p \Rightarrow v - p \geq 0.0$$

The Distribution Type

We didn't specify the semantics of relational distribution types.

A first approach to lifting

$$\frac{(\mu_1, \mu_2) \in \mathfrak{M}[|T|] \times \mathfrak{M}[|T|]}{(\mu_1, \mu_2) \in (\mathfrak{M}[T])_\theta}$$

The Distribution Type

We didn't specify the semantics of relational distribution types.

A first approach to lifting

$$\frac{(d_1, d_2) \in \llbracket T \rrbracket_\theta \quad (\mu_1, \mu_2) \in \mathfrak{M}[\llbracket T \rrbracket] \times \mathfrak{M}[\llbracket T \rrbracket]}{(\mu_1, \mu_2) \in \llbracket \mathfrak{M}[T] \rrbracket_\theta}$$

The Distribution Type

We didn't specify the semantics of relational distribution types.

A first approach to lifting

$$\frac{?? \quad (d_1, d_2) \in \llbracket T \rrbracket_\theta \quad (\mu_1, \mu_2) \in \mathfrak{M}[\llbracket T \rrbracket] \times \mathfrak{M}[\llbracket T \rrbracket]}{(\mu_1, \mu_2) \in \llbracket \mathfrak{M}[T] \rrbracket_\theta}$$

The Distribution Type

We didn't specify the semantics of relational distribution types.

A first approach to lifting

$$\frac{?? \quad (d_1, d_2) \in \llbracket T \rrbracket_\theta \quad (\mu_1, \mu_2) \in \mathfrak{M}[\llbracket T \rrbracket] \times \mathfrak{M}[\llbracket T \rrbracket]}{(\mu_1, \mu_2) \in \llbracket \mathfrak{M}[T] \rrbracket_\theta}$$

We need to relate (d_1, d_2) to (μ_1, μ_2) !

The Distribution Type

We didn't specify the semantics of relational distribution types.

A first approach to lifting

$$\frac{?? \quad (d_1, d_2) \in \langle T \rangle_\theta \quad (\mu_1, \mu_2) \in \mathfrak{M}[|T|] \times \mathfrak{M}[|T|]}{(\mu_1, \mu_2) \in \langle \mathfrak{M}[T] \rangle_\theta}$$

We need to relate (d_1, d_2) to (μ_1, μ_2) !

Solution: define a lifting of the relation $\langle T \rangle_\theta$ through a witness distribution $\mu = \mathfrak{M}[|T| \times |T|]$, such that:

$$\Pr_{x \leftarrow \mu_1} x \in [|T|] = \sum_{y \in T} \Pr_{(x,y) \leftarrow \mu} (x, y) \in \langle T \rangle_\theta$$

More formally, for a relation $\Phi : T_1 \times T_2$, the predicate $\mathcal{L}(\Phi) \mu_1 \mu_2$ holds iff there exists a distribution $\mu \in \mathfrak{M}[T_1 \times T_2]$ such that for every $H \subseteq T_1$, we have

$$\Pr_{x \leftarrow \mu_1} [H(x)] = \sum_{y \in T_2} \Pr_{(x,y) \leftarrow \mu} [H(x) \wedge \Phi(x,y)]$$

and symmetrically for T_2 .

“Probability of events in $\mu_1 \mu_2$ must respect the relation”.

Examples of Lifting

As an example, for $\Phi \equiv \{(F, F), (F, T), (T, T)\}$ we have liftings:

$$\begin{array}{ll} \mu_1(F) = 2/3 & \mu(F, F) = 1/3 \\ \mu_1(T) = 1/3 & \mu(F, T) = 1/3 \\ \mu_2(F) = 1/3 & \mu(T, F) = 0 \\ \mu_2(T) = 2/3 & \mu(T, T) = 1/3 \end{array}$$

$$\begin{array}{ll} \mu_1(F) = 1 & \mu(F, F) = 1 \\ \mu_1(T) = 0 & \mu(F, T) = 0 \\ \mu_2(F) = 1 & \mu(T, F) = 0 \\ \mu_2(T) = 0 & \mu(T, T) = 0 \end{array}$$

Semantics of the Distribution Type

We can now interpret the relational distribution type as all the distributions satisfying the lifting:

$$\frac{\mu_1, \mu_2 \in \mathfrak{M}[|T|] \quad \mathcal{L}(|T|_\theta) \mu_1 \mu_2}{(\mu_1, \mu_2) \in (\mathfrak{M}[T])_\theta}$$

In particular, the type $\mathfrak{M}[\{x :: T \mid x_\triangleleft = x_\triangleright\}]$ forces equal distributions.

Expectation

Expectation of a function f over μ is:

$$E_{\mu} f := \sum_{x \in D} (f(x)) \cdot (\mu(x))$$

Expectation

Expectation of a function f over μ is:

$$E_{\mu} f := \sum_{x \in D} (f x) \cdot (\mu x)$$

We capture monotonicity of expectation as:

$$\begin{aligned} I &:= [0, 1] \\ IBF &:= \{f :: D \rightarrow I \mid \forall d : D. f_{\triangleleft} d \geq f_{\triangleright} d\} \\ E &: \Pi(\mu :: \mathfrak{M}[\{x :: D \mid x_{\triangleleft} = x_{\triangleright}\}]). \Pi(f :: IBF). \{e :: I \mid e_{\triangleleft} \geq e_{\triangleright}\} \end{aligned}$$

Sound as a primitive; other types are possible.

Randomized Auctions

- ▶ Using the probabilistic primitives, we can now define and verify randomized auctions, which have much better revenue properties than the fixed price one.
- ▶ The price a bidder gets won't still depend on her bid, however:
- ▶ we *randomly* split the bidders in two groups, g_a, g_b , we compute the revenue-maximizing price for each group, p_a, p_b , and sell to g_a using p_b and conversely.
- ▶ This auction is truthful on the *expected* utility.

Universal truthfulness:

A bidder will be never able to gain from lying, even knowing the random coins of the mechanism.

The Competitive Auction

```
let utility (v          : real)
            (bid        :: { b :: R | bΔ = v })
            (otherbids  : L[R])
            (g, groups) : (B * L[B])
            : { u :: real | uΔ >= u▷ } =
  match split g bid others otherbids with
  | (g1, g2) →
    if g then fixedprice v bid (prices g2)
    else fixedprice v bid (prices g1)

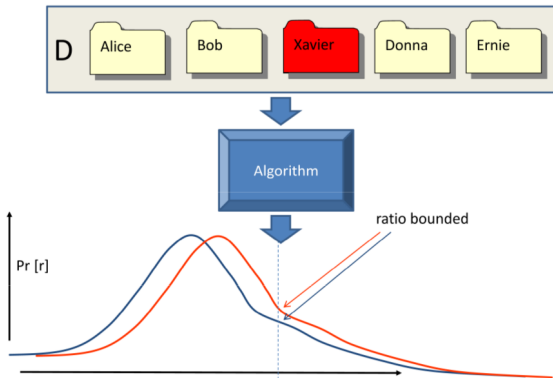
let auction (n : N) (v : R)
            (bid        :: { b  :: R | bΔ = v })
            (otherbids  : L[R])
            : { u :: real | uΔ >= u▷ } =
  let grouping :: M{ r :: (B * B list) | rΔ = r▷ } =
    mlet mycoin = flip      in
    mlet coins  = flipN n in
    munit (mycoin, coins)
  in E grouping (utility v bid otherbids)
```

The Competitive Auction

```
let E (mu : M[ r :  $\alpha$  |  $r_{\triangleleft} = r_{\triangleright}$  ])  
      (f :  $\alpha \rightarrow \text{real}$  |  $\forall x : \alpha, f_{\triangleleft} x \geq f_{\triangleright} x$ )  
      : { r :: real |  $r_{\triangleleft} \geq r_{\triangleright}$  } = ....  
  
let utility (v           : real)  
            (bid         :: { b :: R |  $b_{\triangleleft} = v$  })  
            (otherbids  : L[R])  
            (g, groups) : (B * L[B])  
            : { u :: real |  $u_{\triangleleft} \geq u_{\triangleright}$  } = ...  
  
let auction (n : N) (v : R)  
            (bid         :: { b :: R |  $b_{\triangleleft} = v$  })  
            (otherbids  : L[R])  
            : { u :: real |  $u_{\triangleleft} \geq u_{\triangleright}$  } =  
      let grouping :: M{ r :: (B * B list) |  $r_{\triangleleft} = r_{\triangleright}$  } = ...  
      in E grouping (utility v bid otherbids)
```

Differential Privacy

Contribution of a single individual to the output of a mechanism cannot be effectively distinguished by an attacker under worst-case assumptions.



Formal Definition

A probabilistic function $F : T \rightarrow S$ is (ϵ, δ) -Differentially Private if for all pairs of *adjacent* $t_1, t_2 \in T$ and for every $E \subseteq S$:

$$\Pr_{x \leftarrow F t_1} [x \in E] \leq \exp(\epsilon) \Pr_{x \leftarrow F t_2} [x \in E] + \delta$$

Formal Definition

A probabilistic function $F : T \rightarrow S$ is (ϵ, δ) -Differentially Private if for all pairs of *adjacent* $t_1, t_2 \in T$ and for every $E \subseteq S$:

$$\Pr_{x \leftarrow F t_1} [x \in E] \leq \exp(\epsilon) \Pr_{x \leftarrow F t_2} [x \in E] + \delta$$

Example: The Laplace Mechanism:

- ▶ Compute the *sensitivity* k of f .
- ▶ For input t , release $f(t) + \text{random noise}$, scaled by k .

Formal Definition

A probabilistic function $F : T \rightarrow S$ is (ϵ, δ) -Differentially Private if for all pairs of *adjacent* $t_1, t_2 \in T$ and for every $E \subseteq S$:

$$\Pr_{x \leftarrow F t_1} [x \in E] \leq \exp(\epsilon) \Pr_{x \leftarrow F t_2} [x \in E] + \delta$$

Example: The Laplace Mechanism:

- ▶ Compute the *sensitivity* k of f .
- ▶ For input t , release $f(t) + \text{random noise}$, scaled by k .

Many algorithms are DP: private database release, counters, analytics, **strong connection to Mechanism Design!**

Approximately Reasoning over Distributions

We can capture DP with a refinement over the type of probability distributions using the definition of Δ -distance:

$$\Delta_{\epsilon}(\mu_1, \mu_2) = \max_{E \subseteq U} \left(\Pr_{x \leftarrow \mu_2} [x \in E] - \exp(\epsilon) \Pr_{x \leftarrow \mu_1} [x \in E] \right)$$

Approximately Reasoning over Distributions

We can capture DP with a refinement over the type of probability distributions using the definition of Δ -distance:

$$\Delta_{\epsilon}(\mu_1, \mu_2) = \max_{E \subseteq U} \left(\Pr_{x \leftarrow \mu_2} [x \in E] - \exp(\epsilon) \Pr_{x \leftarrow \mu_1} [x \in E] \right)$$

Then, f is (ϵ, δ) differentially private if it has type:

$$\{d :: T \mid \text{Adj}(d_{\triangleleft}, d_{\triangleright})\} \rightarrow \{r :: \mathfrak{M}[\mathbb{R}] \mid \Delta_{\epsilon}(r_{\triangleleft}, r_{\triangleright}) \leq \delta\}$$

However, verification conditions involving Δ are quite hard.

The Relational Distribution Type

Our solution: Internalize distribution distance in the types:

$$\frac{\mu_1, \mu_2 \in \mathfrak{M}[|T|] \quad \mathcal{L}_{\epsilon, \delta}(|T|)_\theta \mu_1 \mu_2}{(\mu_1, \mu_2) \in (\mathfrak{M}_{\epsilon, \delta}[T])_\theta}$$

Lifting is extended from $p = p_1$ to $p \leq p_1 \leq \text{exp}(p) + \delta$.

The Relational Distribution Type

Our solution: Internalize distribution distance in the types:

$$\frac{\mu_1, \mu_2 \in \mathfrak{M}[|T|] \quad \mathcal{L}_{\epsilon, \delta}(|T|)_\theta \mu_1 \mu_2}{(\mu_1, \mu_2) \in (\mathfrak{M}_{\epsilon, \delta}[T])_\theta}$$

Lifting is extended from $p = p_1$ to $p \leq p_1 \leq \text{exp}(p) + \delta$.

Capturing DP

The interpretation of $\mathfrak{M}_{\epsilon, \delta}[\{r :: \mathbb{R} \mid r_{\triangleleft} = r_{\triangleright}\}]$ is the set of pairs of probability distributions that are (ϵ, δ) -apart, **capturing DP**.

The Relational Distribution Type

Our solution: Internalize distribution distance in the types:

$$\frac{\mu_1, \mu_2 \in \mathfrak{M}[|T|] \quad \mathcal{L}_{\epsilon, \delta}(|T|)_\theta \mu_1 \mu_2}{(\mu_1, \mu_2) \in (\mathfrak{M}_{\epsilon, \delta}[T])_\theta}$$

Lifting is extended from $p = p_1$ to $p \leq p_1 \leq \text{exp}(p) + \delta$.

Capturing DP

The interpretation of $\mathfrak{M}_{\epsilon, \delta}[\{r :: \mathbb{R} \mid r_{\triangleleft} = r_{\triangleright}\}]$ is the set of pairs of probability distributions that are (ϵ, δ) -apart, **capturing DP**.

DP algorithms are typed as:

$$f : \{d :: T \mid \text{Adj}(d_{\triangleleft}, d_{\triangleright})\} \rightarrow \mathfrak{M}_{\epsilon, \delta}[\{r :: \mathbb{R} \mid r_{\triangleleft} = r_{\triangleright}\}]$$

The Probability Polymonad

Reasoning about distance is compositional:

$$\text{SUB-M} \frac{\mathcal{G} \vdash T \preceq U \quad \forall \theta. \theta \vdash \mathcal{G}, x :: T \Rightarrow \llbracket \epsilon_1 \leq \epsilon_2 \wedge \delta_1 \leq \delta_2 \rrbracket_\theta}{\mathcal{G} \vdash \mathfrak{M}_{\epsilon_1, \delta_1}[T] \preceq \mathfrak{M}_{\epsilon_2, \delta_2}[U]}$$

$$\text{UNITM} \frac{\mathcal{G} \vdash e :: T}{\mathcal{G} \vdash \text{unit}_M e :: \mathfrak{M}_{\epsilon, \delta}[T]}$$

$$\text{BINDM} \frac{\mathcal{G} \vdash e_1 :: \mathfrak{M}_{\epsilon_1, \delta_1}[T_1] \quad \mathcal{G}, x :: T_1 \vdash e_2 :: \mathfrak{M}_{\epsilon_2, \delta_2}[T_2]}{\mathcal{G} \vdash \text{bind}_M x = e_1 \text{ in } e_2 :: \mathfrak{M}_{\epsilon_1 + \epsilon_2, \delta_1 + \delta_2}[T_2]}$$

Bind is distance-adjusting sampling.

Type for the Laplace Mechanism

Recall the Laplace Mechanism:

For a k -sensitive f , f plus k/ϵ -scaled Laplacian noise is DP. This is captured by the type:

Type for the Laplace Mechanism

Recall the Laplace Mechanism:

For a k -sensitive f , f plus k/ϵ -scaled Laplacian noise is DP. This is captured by the type:

$$\text{lap} : \Pi(\epsilon :: \mathbb{R}). \Pi(x :: \mathbb{R}). \mathfrak{M}_{\epsilon * |x_{\triangleleft} - x_{\triangleright}|, 0}[\{r :: \mathbb{R} \mid r_{\triangleleft} = r_{\triangleright}\}]$$

Note that the actual distance $\epsilon * |x_{\triangleleft} - x_{\triangleright}|$ depends on the distance of the inputs. This is a better alternative than using a precondition on x .

Type for the Laplace Mechanism

Recall the Laplace Mechanism:

For a k -sensitive f , f plus k/ϵ -scaled Laplacian noise is DP. This is captured by the type:

$$\text{lap} : \Pi(\epsilon :: \mathbb{R}). \Pi(x :: \mathbb{R}). \mathfrak{M}_{\epsilon * |x_{\triangleleft} - x_{\triangleright}|, 0}[\{r :: \mathbb{R} \mid r_{\triangleleft} = r_{\triangleright}\}]$$

Note that the actual distance $\epsilon * |x_{\triangleleft} - x_{\triangleright}|$ depends on the distance of the inputs. This is a better alternative than using a precondition on x .

Using the bind rule, we can sample from laplace and assume the sampled value equal in both runs.

Example: Private Histogram

We add noise to an histogram to make it private.

```
let rec histogram {l :: L(R) | Adj x◁ x▷) }  
  : M[e * d(l◁, l▷)] { r :: L(R) | r◁ = r▷ } =  
match l with  
| []      → unit []  
| x :: xs →  
  mlet y = lap eps x    in  
  mlet ys = histogram xs in  
  munit (y :: ys)
```

Example: Private Histogram

We add noise to an histogram to make it private.

```
let rec histogram {l :: L(R) | Adj x◁ x▷) }  
  : M[e * d(l◁, l▷)] { r :: L(R) | r◁ = r▷ } =  
match l with  
| []           → unit []  
| x :: xs →  
  mlet y = lap eps x    in  
  mlet ys = histogram xs in  
  munit (y :: ys)
```

The main proof obligation is:

$$e * d(x_{\triangleleft} :: xs_{\triangleleft}, x_{\triangleright} :: xs_{\triangleright}) \geq e * (d(x_{\triangleleft}, x_{\triangleright}) + d(xs_{\triangleleft}, xs_{\triangleright}))$$

which is implied by the adjacency precondition.

Combining MD and DP: Aggregative Games

- ▶ We verify the computation of an approximate Nash-equilibrium.
- ▶ n agents can choose over a space of actions $a_i \in A$.

Combining MD and DP: Aggregative Games

- ▶ We verify the computation of an approximate Nash-equilibrium.
- ▶ n agents can choose over a space of actions $a_i \in A$.
- ▶ (a_1, \dots, a_n) is an α -approximate Nash-equilibrium if no single agent i can gain more than α payoff by *unilateral* deviation: For all agents i and actions a'_i :

$$E[P_i(a_1, \dots, a_i, \dots, a_N)] \geq E[P_i(a_1, \dots, a'_i, \dots, a_N)] - \alpha.$$

Combining MD and DP: Aggregative Games

- ▶ We verify the computation of an approximate Nash-equilibrium.
- ▶ n agents can choose over a space of actions $a_i \in A$.
- ▶ (a_1, \dots, a_n) is an α -approximate Nash-equilibrium if no single agent i can gain more than α payoff by *unilateral* deviation: For all agents i and actions a'_i :

$$E[P_i(a_1, \dots, a_i, \dots, a_N)] \geq E[P_i(a_1, \dots, a'_i, \dots, a_N)] - \alpha.$$

- ▶ Assumption: Payoff for i depends only on a_i plus a *signal*, a positive (bounded) real number depending on the aggregated actions of all players.

Combining MD and DP: Aggregative Games

- ▶ The key: use differential privacy to compute the equilibria.
- ▶ Mediator: The mechanism suggests the equilibria action a_i .
- ▶ We prove that the player gets optimal utility if she does a_i .
- ▶ We reason over a deviation function dev_i for player i .

Combining MD and DP: Aggregative Games

- ▶ The key: use differential privacy to compute the equilibria.
- ▶ Mediator: The mechanism suggests the equilibria action a_i .
- ▶ We prove that the player gets optimal utility if she does a_i .
- ▶ We reason over a deviation function dev_i for player i .

In types:

```
let aggregative_utility ( ... )  
  { dev :: act → act |  $\forall a : \text{act}. \text{dev}_{\triangleleft} a = a$  }  
  : { u :: real |  $u_{\triangleleft} \geq u_{\triangleright} - \text{alpha}$  }
```

Combining MD and DP: Aggregative Games

- ▶ The key: use differential privacy to compute the equilibria.
- ▶ Mediator: The mechanism suggests the equilibria action a_i .
- ▶ We prove that the player gets optimal utility if she does a_i .
- ▶ We reason over a deviation function dev_i for player i .

In types:

```
let aggregative_utility ( ... )
  { dev :: act → act | ∀ a : act. dev_Δ a = a }
  : { u :: real | u_Δ ≥ u_▷ - alpha }
```

Relate expectation to distance on the distributions:

$$E : \Pi(\mu :: \mathfrak{M}_{\epsilon, \delta}[\{x :: I \mid x_{\Delta} \leq x_{\triangleright} + c\}]). \{e :: I \mid e_{\Delta} \leq e_{\triangleright} + \epsilon + c + \delta e^{-\epsilon}\}$$

The Implementation

- ▶ Hybrid SMT/Bidirectional type checking.
- ▶ Why3 as the SMT backend, multiple solvers required.
- ▶ Verification using top-level annotations (+2 cuts).
- ▶ Top-level types act as the specification.
- ▶ Support for debug of type-checking failures important.

Benchmarks

Example	# Lines	Verif. time
histogram	25	2.66 s.
dummysum	31	11.95 s.
noisysum	55	3.64 s.
two-level-a	38	2.55 s.
two-level-b	56	3.94 s.
binary	95	18.56 s.
idc	73	27.60 s.
dualquery	128	27.71 s.
competitive-b	81	2.80 s.
competitive	75	4.19 s.
fixedprice	10	0.90 s.
summarization	471	238.42 s.

Table : Benchmarks

Future work and Conclusions:

Future Work:

- ▶ More examples from the algorithms community.
- ▶ More examples from the security/cryptography domain.
- ▶ More properties: accuracy, fancier distributions.
- ▶ Extensions to the language.

Future work and Conclusions:

Future Work:

- ▶ More examples from the algorithms community.
- ▶ More examples from the security/cryptography domain.
- ▶ More properties: accuracy, fancier distributions.
- ▶ Extensions to the language.

Conclusions

- ▶ Higher-Order Approximate Probabilistic Relational Refinement Types: HOARe2
- ▶ Built-in support for approximate reasoning.
- ▶ Logic seems to capture many examples.
- ▶ Automatic verification worked reasonably well.
- ▶ SMT interaction is still a challenge.

Thank you

Questions?

More on the Aggregative Example:

Expected Payoff for the deviating agent

```
let expay br* dev* br dev =  
  E (mlet  sums = mkSums k br* br      in  
    let  s•  = search k br* br sums in  
    let  a*   = dev* (br* s•)      in  
    let  a    = dev  (br  s•)      in  
    let  p*   = pay* a* (sign a* a) in  
    munit p*) (λx. x)
```

s^\bullet is close to the true signal on the strategy profile.

More on the Aggregative Example:

Type for *expay*

$$\begin{aligned} & \{br^* :: \mathbb{R} \rightarrow \mathbf{A} \mid \forall s, a. pay^* (br^*_{\triangleleft} s) s \geq pay^* a s\} \\ \rightarrow & \{dev^* :: \mathbf{A} \rightarrow \mathbf{A} \mid \forall x. dev^*_{\triangleleft} x = x\} \\ \rightarrow & \{br :: \mathbb{R} \rightarrow \mathbf{A} \mid br_{\triangleleft} = br_{\triangleright}\} \\ \rightarrow & \{dev :: \mathbf{A} \rightarrow \mathbf{A} \mid \forall a. dev_{\triangleleft} a = dev_{\triangleright} a = a\} \\ \rightarrow & \{u :: \overline{\mathbb{R}}^+ \mid u_{\triangleleft} \geq u_{\triangleright} - \alpha\}. \end{aligned}$$

Extended type for Laplace

Lap with a refinement type capturing *accuracy*:

$$\Pi(x :: \mathbb{R}). \mathfrak{M}_{\epsilon, |x_{\triangleleft} - x_{\triangleright}|, \beta}[\{u :: \mathbb{R} \mid u_{\triangleleft} = u_{\triangleright} \wedge |x_{\triangleleft} - u_{\triangleleft}| < T\}]$$