

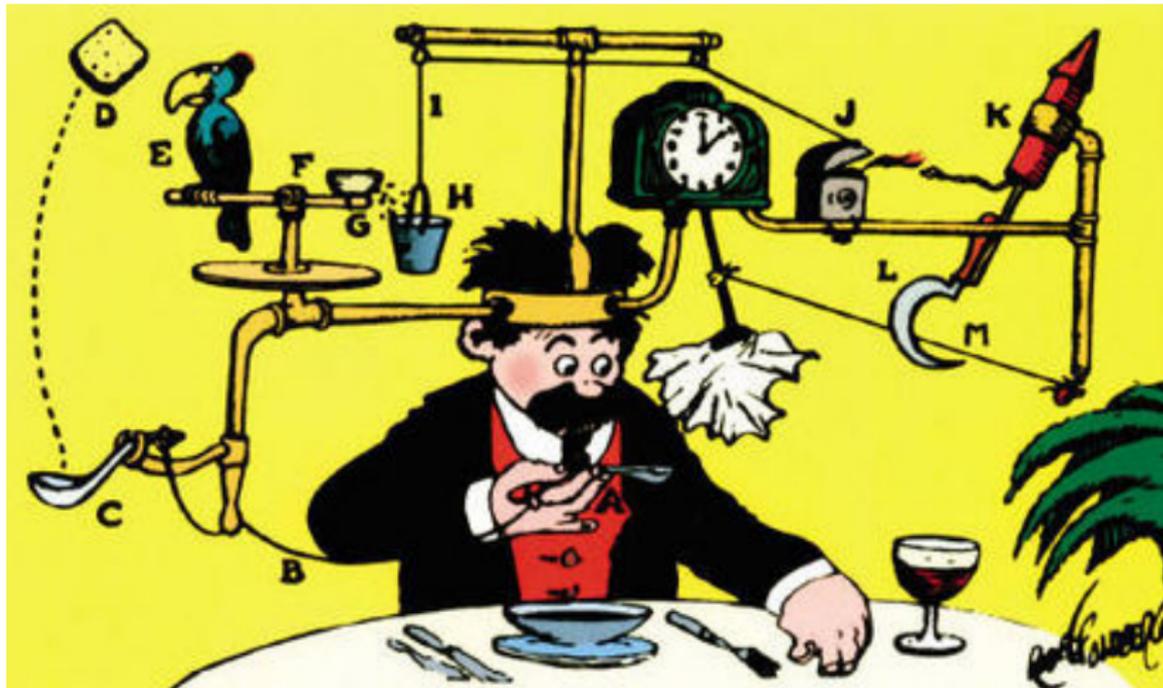
Higher-Order Relational Refinement Types for Mechanism Design and Differential Privacy

Gilles Barthe¹, Marco Gaboardi²,
Emilio Jesús Gallego Arias^{3,4}, Justin Hsu⁴,
Aaron Roth⁴, Pierre-Yves Strub¹

¹IMDEA Software, ²University of Dundee,
³CRI Mines–ParisTech, ⁴University of Pennsylvania

January 15th, 2015

The Application



Mechanism Design

One painting for sale



One painting for sale



How much will you pay?



One painting for sale



How much will you pay?



\$10 million!



\$50 million!



\$3

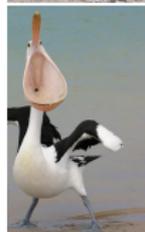
One painting for sale



How much will you pay?



\$10 million!



\$50 million!



\$3

Who wins, and for how much?

How much will you pay?

Top bid pays top price?

- Simple rule
- Can encourage manipulation...



\$10 million!



\$50 million!



\$3

How much will you pay?

Top bid pays top price?

- Simple rule
- Can encourage manipulation...



\$10 million!



~~\$50 million!~~

\$10.1 million?



\$3

Algorithm design with strategic inputs

Algorithm design with strategic inputs

Rational agents

- Report data
- Care about output
- May lie, strategize



Algorithm design with strategic inputs

Rational agents

- Report data
- Care about output
- May lie, strategize



Goal: encourage “good” behavior

Designing auctions

- Bidders each have personal value $v : \mathbb{R}$ for the item

Designing auctions

- Bidders each have personal value $v : \mathbb{R}$ for the item
- Bidder's **happiness** is function of price, v , whether they win

Designing auctions

- Bidders each have personal value $v : \mathbb{R}$ for the item
- Bidder's **happiness** is function of price, v , whether they win
- Bidder reports a bid $b : \mathbb{R}$ to the mechanism

Designing auctions

- Bidders each have personal value $v : \mathbb{R}$ for the item
- Bidder's **happiness** is function of price, v , whether they win
- Bidder reports a bid $b : \mathbb{R}$ to the mechanism

Property: agent always maximizes happiness with $b = v$

Fixed price auction

- Given a fixed price p
- Bidder bids b , buys item if $b > p$

Fixed price auction

- Given a fixed price $price$
- Bidder bids bid , buys item if higher than $price$

What is the happiness function for a bidder?

```
fixedprice price value bid =  
  if bid > price then  
    value - price  
  else  
    0
```

Consider bidder's happiness function...

- First run: bidder bids $b = v$ (honest)

Consider bidder's happiness function...

- First run: bidder bids $b = v$ (honest)
- Second run: bidder bids arbitrarily (maybe not honest)

Consider bidder's happiness function...

- First run: bidder bids $b = v$ (honest)
- Second run: bidder bids arbitrarily (maybe not honest)
- Verify: happiness in first run is higher than in second run

Consider bidder's happiness function...

- First run: bidder bids $b = v$ (honest)
- Second run: bidder bids arbitrarily (maybe not honest)
- Verify: happiness in first run is higher than in second run

Consider bidder's happiness function...

- First run: bidder bids $b = v$ (honest)
- Second run: bidder bids arbitrarily (maybe not honest)
- Verify: happiness in first run is higher than in second run

```
fixedprice p v v =  
  if v > p then  
    v - p  
  else  
    0
```



```
fixedprice p v b =  
  if b > p then  
    v - p  
  else  
    0
```

Consider bidder's happiness function...

- First run: bidder bids $b = v$ (honest)
- Second run: bidder bids arbitrarily (maybe not honest)
- Verify: happiness in first run is higher than in second run

```
fixedprice p v v =  
  if v > p then  
    v - p  
  else  
    0
```



```
fixedprice p v b =  
  if b > p then  
    v - p  
  else  
    0
```

This is a **relational** property

Introducing HOARe²



A type system with relational refinement types

Judgment

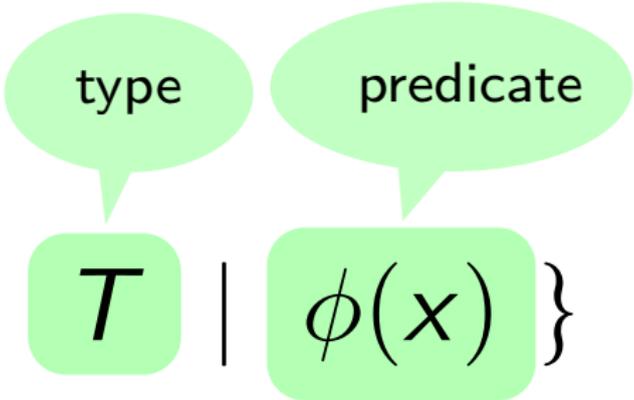
$$\Gamma \vdash e : \{x : T \mid \phi(x)\}$$

Judgment

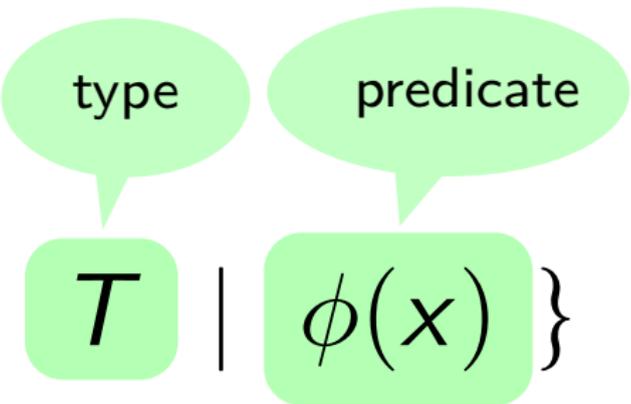
type

$$\Gamma \vdash e : \{x : \mathcal{T} \mid \phi(x)\}$$

Judgment

$$\Gamma \vdash e : \{x : \mathcal{T} \mid \phi(x)\}$$


Judgment



The diagram shows the refinement type judgment $\Gamma \vdash e : \{x : T \mid \phi(x)\}$. The symbol T is enclosed in a light green rounded rectangle, with a light green speech bubble above it containing the word "type". The expression $\phi(x)$ is enclosed in a light green rounded rectangle, with a light green speech bubble above it containing the word "predicate".

$$\Gamma \vdash e : \{x : T \mid \phi(x)\}$$

“ e is a program of type T such that $\phi(e)$ holds”

Example

$$\Gamma \vdash 3 : \{x : \mathbb{Z} \mid x \geq 0\}$$

Example

$$\Gamma \vdash 3 : \{x : \mathbb{Z} \mid x \geq 0\}$$

“3 is a non-negative integer”

Relational Judgment

$$\Gamma \vdash e :: \{x :: T \mid \phi(x_{\triangleleft}, x_{\triangleright})\}$$

Relational Judgment

$$\Gamma \vdash e :: \{x :: T \mid \phi(x_{\triangleleft}, x_{\triangleright})\}$$

Relational Judgment

$$\Gamma \vdash e :: \{x :: T \mid \phi(x_{\triangleleft}, x_{\triangleright})\}$$

ϕ mentions two runs of program e via x_{\triangleleft} and x_{\triangleright}

Relational Judgment

$$\Gamma \vdash e :: \{x :: T \mid \phi(x_{\triangleleft}, x_{\triangleright})\}$$

ϕ mentions two runs of program e via x_{\triangleleft} and x_{\triangleright}

Example

$$\{y :: \mathbb{Z} \mid y_{\triangleleft} \leq y_{\triangleright}\} \vdash e :: \{x :: \mathbb{Z} \mid x_{\triangleleft} \leq x_{\triangleright}\}$$

Relational Judgment

$$\Gamma \vdash e :: \{x :: T \mid \phi(x_{\triangleleft}, x_{\triangleright})\}$$

ϕ mentions two runs of program e via x_{\triangleleft} and x_{\triangleright}

Example

$$\{y :: \mathbb{Z} \mid y_{\triangleleft} \leq y_{\triangleright}\} \vdash e :: \{x :: \mathbb{Z} \mid x_{\triangleleft} \leq x_{\triangleright}\}$$

“If y increases, then e increases.”

Relational Judgment

$$\Gamma \vdash e :: \{x :: T \mid \phi(x_{\triangleleft}, x_{\triangleright})\}$$

ϕ mentions two runs of program e via x_{\triangleleft} and x_{\triangleright}

Example

$$\{y :: \mathbb{Z} \mid y_{\triangleleft} \leq y_{\triangleright}\} \vdash e :: \{x :: \mathbb{Z} \mid x_{\triangleleft} \leq x_{\triangleright}\}$$

“If y increases, then e increases.”

Background

- First used in the RF* language, POPL 2014

Happiness function

```
fixedprice price value bid =  
  if bid > price then  
    value - price  
  else  
    0
```

Happiness function

```
fixedprice price value bid =  
  if bid > price then  
    value - price  
  else  
    0
```

Truthfulness in a type

Happiness function

```
fixedprice price value bid =  
  if bid > price then  
    value - price  
  else  
    0
```

Truthfulness in a type

$$\{p :: \mathbb{R} \mid p_{\triangleleft} = p_{\triangleright}\}$$

(Fixed price)

Happiness function

```
fixedprice price value bid =  
  if bid > price then  
    value - price  
  else  
    0
```

Truthfulness in a type

$$\{p :: \mathbb{R} \mid p_{\triangleleft} = p_{\triangleright}\}$$
$$\rightarrow \{v :: \mathbb{R} \mid v_{\triangleleft} = v_{\triangleright}\}$$

(Fixed price)

(Bidder value fixed)

Happiness function

```

fixedprice price value bid =
  if bid > price then
    value - price
  else
    0

```

Truthfulness in a type

$$\{p :: \mathbb{R} \mid p_{\triangleleft} = p_{\triangleright}\}$$

(Fixed price)

$$\rightarrow \{v :: \mathbb{R} \mid v_{\triangleleft} = v_{\triangleright}\}$$

(Bidder value fixed)

$$\rightarrow \{b :: \mathbb{R} \mid b_{\triangleleft} = v_{\triangleleft}\}$$

(Bid = value on \triangleleft run)

Happiness function

```

fixedprice price value bid =
  if bid > price then
    value - price
  else
    0
  
```

Truthfulness in a type

$\{p :: \mathbb{R} \mid p_{\triangleleft} = p_{\triangleright}\}$	(Fixed price)
$\rightarrow \{v :: \mathbb{R} \mid v_{\triangleleft} = v_{\triangleright}\}$	(Bidder value fixed)
$\rightarrow \{b :: \mathbb{R} \mid b_{\triangleleft} = v_{\triangleleft}\}$	(Bid = value on \triangleleft run)
$\rightarrow \{u :: \mathbb{R} \mid u_{\triangleleft} \geq u_{\triangleright}\}$	(Truthful)

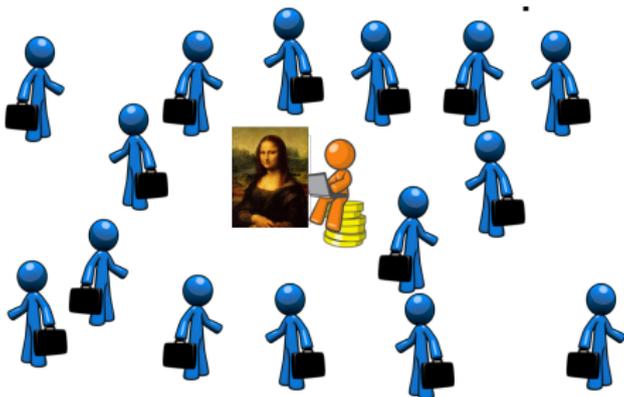
A more complex auction

- Unlimited supply of items (e.g., music files)
- Want to use `fixedprice`, but for what price?

A more complex auction

- Unlimited supply of items (e.g., music files)
- Want to use fixed price, but for what price?

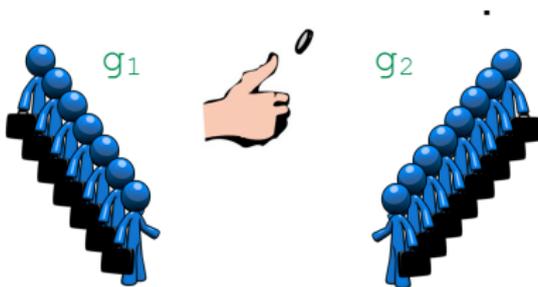
Randomize!



A more complex auction

- Unlimited supply of items (e.g., music files)
- Want to use fixed price, but for what price?

Randomize!



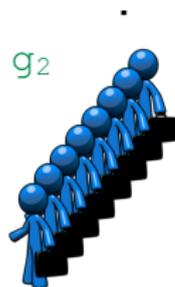
A more complex auction

- Unlimited supply of items (e.g., music files)
- Want to use fixed price, but for what price?

Randomize!



optimal
price

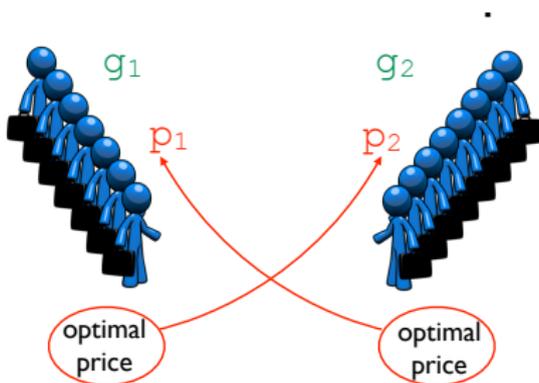


optimal
price

A more complex auction

- Unlimited supply of items (e.g., music files)
- Want to use fixed price, but for what price?

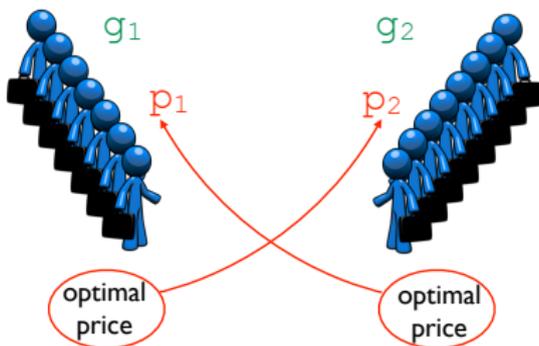
Randomize!



A more complex auction

- Unlimited supply of items (e.g., music files)
- Want to use fixed price, but for what price?

Randomize!

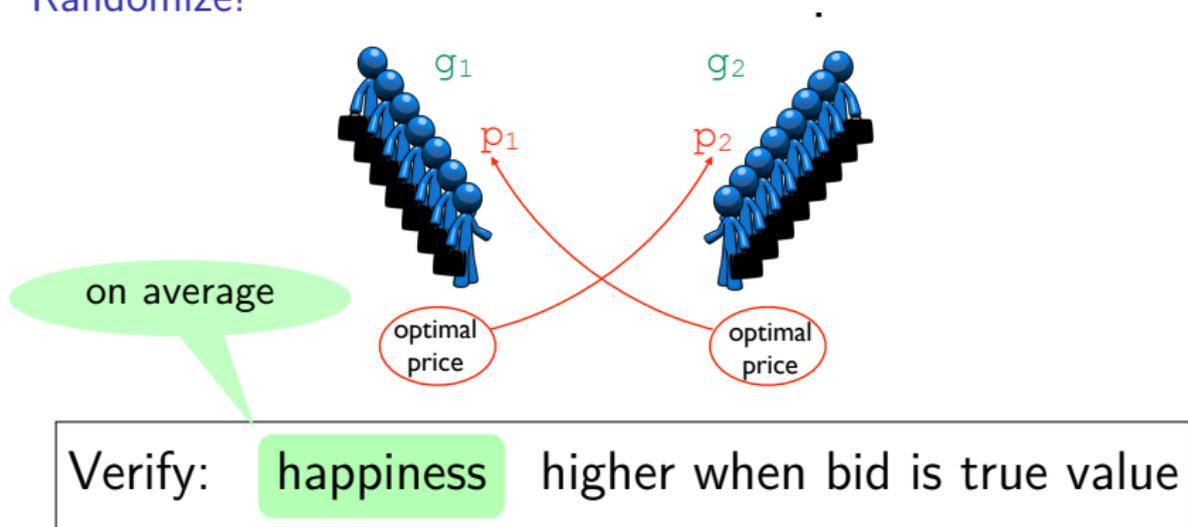


Verify: happiness higher when bid is true value

A more complex auction

- Unlimited supply of items (e.g., music files)
- Want to use fixed price, but for what price?

Randomize!



Monotonicity of expectation

- (One) Distribution μ over A



Monotonicity of expectation

- (One) Distribution μ over A
- Two functions $f_1, f_2 : A \rightarrow \mathbb{R}$ with

$$f_1(x) \geq f_2(x) \quad \text{for all } x \in A$$



Monotonicity of expectation

- (One) Distribution μ over A
- Two functions $f_1, f_2 : A \rightarrow \mathbb{R}$ with

$$f_1(x) \geq f_2(x) \quad \text{for all } x \in A$$

- Then, fact about expected values:

$$\mathbb{E}_\mu[f_1] \geq \mathbb{E}_\mu[f_2]$$



Monotonicity of expectation

- (One) Distribution μ over A
- Two functions $f_1, f_2 : A \rightarrow \mathbb{R}$ with

$$f_1(x) \geq f_2(x) \quad \text{for all } x \in A$$

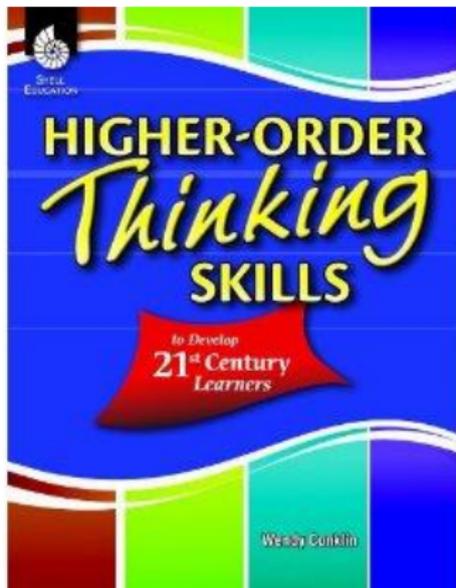
- Then, fact about expected values:

$$\mathbb{E}_\mu[f_1] \geq \mathbb{E}_\mu[f_2]$$

f_1 bigger than
 f_2 on average



Extending HOARe²



Distributions and Higher-order refinements

Probabilistic programs

- Reason about two runs of a probabilistic program
- Use type of probability distributions

Probabilistic programs

- Reason about two runs of a probabilistic program
- Use type of probability distributions

Typing distributions

$$\Gamma \vdash e :: \mathfrak{M}_{0,0}[\{x :: T \mid \phi(x_{\triangleleft}, x_{\triangleright})\}]$$

Probabilistic programs

- Reason about two runs of a probabilistic program
- Use type of probability distributions

Typing distributions

$$\Gamma \vdash e :: \mathfrak{M}_{0,0}[\{x :: T \mid \phi(x_{\triangleleft}, x_{\triangleright})\}]$$

“ e is a distribution over T , with two runs related by ϕ ”

Probabilistic programs

- Reason about two runs of a probabilistic program
- Use type of probability distributions

Typing distributions

$$\Gamma \vdash e :: \mathfrak{M}_{0,0}[\{x :: T \mid \phi(x_{\triangleleft}, x_{\triangleright})\}]$$

“e is a distribution over T , with two runs related by ϕ ”

???

$$\Gamma \vdash e :: \mathfrak{M}_{0,0}[\{x :: T \mid \phi(x_{\triangleleft}, x_{\triangleright})\}]$$

What does this mean?

- Convert relation ϕ to a relation $\phi^\#$ on **distributions** over T
- Two runs of e related by $\phi^\#$ (as distributions!)

Example

$$\Gamma \vdash e :: \mathcal{M}_{0,0}[\{x :: T \mid x_{\triangleleft} = x_{\triangleright}\}]$$

Example

$$\Gamma \vdash e :: \mathcal{M}_{0,0}[\{x :: T \mid x_{\triangleleft} = x_{\triangleright}\}]$$

Example

$$\Gamma \vdash e :: \mathcal{M}_{0,0}[\{x :: T \mid x_{\triangleleft} = x_{\triangleleft} \}]$$

“e is a distribution over T that is identical in both runs”

Example

$$\Gamma \vdash e :: \mathcal{M}_{0,0}[\{x :: T \mid x_{\triangleleft} = x_{\triangleright}\}]$$

“ e is a distribution over T that is identical in both runs”

Background

- Proposed by Barthe, Köpf, Olmedo, Zanella
- Generalizing $0, 0$ to ϵ, δ models differential privacy

Example

$$\Gamma \vdash e :: \mathcal{M}_{0,0}[\{x :: T \mid x_{\triangleleft} = x_{\triangleright}\}]$$

“ e is a distribution over T that is identical in both runs”

Background

- Proposed by Barthe, Köpf, Olmedo, Zanella
- Generalizing $0, 0$ to ϵ, δ models differential privacy

Our contribution

- Simplify and build into a type system

Refinements on functions

$$\Gamma \vdash e :: \{f :: T \rightarrow U \mid \phi\}$$

Refinements on functions

$$\Gamma \vdash e :: \{f :: T \rightarrow U \mid \phi\}$$

“ e is a function from T to U that satisfies ϕ ”

Refinements on functions

$$\Gamma \vdash e :: \{f :: T \rightarrow U \mid \phi\}$$

“ e is a function from T to U that satisfies ϕ ”

Our contribution

- Consistency by carefully handling termination

Refinements on functions

$$\Gamma \vdash e :: \{f :: T \rightarrow U \mid \phi\}$$

“ e is a function from T to U that satisfies ϕ ”

Our contribution

- Consistency by carefully handling termination
- Show naïve treatment leads to inconsistency

Expressing monotonicity of expectations

Want to show

$$\mathbb{E} \mu f_1 \geq \mathbb{E} \mu f_2$$

In HOARE², type \mathbb{E} as...

$$\mathfrak{M}_{0,0}[\{x :: A \mid x_{\triangleleft} = x_{\triangleright}\}]$$

(Same distributions)

Expressing monotonicity of expectations

Want to show

$$\mathbb{E} \mu f_1 \geq \mathbb{E} \mu f_2$$

In HOARE², type \mathbb{E} as...

$$\begin{aligned} & \mathfrak{M}_{0,0}[\{x :: A \mid x_{\triangleleft} = x_{\triangleright}\}] && \text{(Same distributions)} \\ & \rightarrow \{f :: A \rightarrow \mathbb{R} \mid \forall x. f_{\triangleleft} x \geq f_{\triangleright} x\} && \text{(Higher-order)} \end{aligned}$$

Expressing monotonicity of expectations

Want to show

$$\mathbb{E} \mu f_1 \geq \mathbb{E} \mu f_2$$

In HOARE², type \mathbb{E} as...

$$\begin{aligned} & \mathfrak{M}_{0,0}[\{x :: A \mid x_{\triangleleft} = x_{\triangleright}\}] && \text{(Same distributions)} \\ & \rightarrow \{f :: A \rightarrow \mathbb{R} \mid \forall x. f_{\triangleleft} x \geq f_{\triangleright} x\} && \text{(Higher-order)} \\ & \rightarrow \{e :: \mathbb{R} \mid e_{\triangleleft} \geq e_{\triangleright}\} && \text{(Monotonic)} \end{aligned}$$

Expressing monotonicity of expectations

Want to show

$$\mathbb{E} \mu f_1 \geq \mathbb{E} \mu f_2$$

In HOARE², type \mathbb{E} as...

$$\begin{aligned} \mathbb{E} &:: \mathfrak{M}_{0,0}[\{x :: A \mid x_{\triangleleft} = x_{\triangleright}\}] && \text{(Same distributions)} \\ &\rightarrow \{f :: A \rightarrow \mathbb{R} \mid \forall x. f_{\triangleleft} x \geq f_{\triangleright} x\} && \text{(Higher-order)} \\ &\rightarrow \{e :: \mathbb{R} \mid e_{\triangleleft} \geq e_{\triangleright}\} && \text{(Monotonic)} \end{aligned}$$

Semantics

- Soundness of the system
- Requires termination

Implementation

- Automated, low annotation burden
- Why3 and SMT solvers

Translation

- Embedding of DFuzz, a language for differential privacy

More complex examples

- Verify differential privacy
- Verify MD properties beyond truthfulness

Takeaway points



Four features, one system

- HOARe²: relational properties for randomized programs
- Combine features in a clean, usable way

Four features, one system

- HOARe²: relational properties for randomized programs
- Combine features in a clean, usable way

Formal verification for mechanism design!

- Exciting, under-explored area for verification
- Tons of interesting properties, mechanisms
- Strong motivation besides (mere) correctness

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

Gilles Barthe¹, Marco Gaboardi²,
Emilio Jesús Gallego Arias^{3,4}, Justin Hsu⁴,
Aaron Roth⁴, Pierre-Yves Strub¹

¹IMDEA Software, ²University of Dundee,
³CRI Mines–ParisTech, ⁴University of Pennsylvania

January 15th, 2015