

Lecture 23 – Refinement Types

AAA551: Programming Language Theory

Jihyeok Park



2026 Spring

- Dependent Types
 - Barendregt's λ -Cube: a **type that depends on a term**
 - Dependent Function Types (Π) and Pair Types (Σ)
 - The Calculus of Constructions (λC) and CIC
 - Curry-Howard, Reloaded: Π as \forall , Σ as \exists
 - Propositional Equality ($a =_A b$)
 - Dependent Types in Practice (Coq/Rocq, Agda, Idris, Lean, F^{*})
- Power vs. automation: type checking is **undecidable**
- Idea: restrict propositions to a **decidable** logic

1. Why Refinement Types?
2. Base Refinement Types
Syntax and Intuition
3. Subtyping as Implication
4. Refining Functions
5. Liquid Types
6. Refinement Types in Practice

1. Why Refinement Types?
2. Base Refinement Types
Syntax and Intuition
3. Subtyping as Implication
4. Refining Functions
5. Liquid Types
6. Refinement Types in Practice

Recall the **dependent** type from Lecture 22 – a number **bundled with a proof** of a proposition about it:

$$\text{Even} \triangleq \Sigma(n : \text{Nat}). (n \bmod 2 =_{\text{Nat}} 0)$$

Dependent types let that proposition be **arbitrary**, but the price is steep:

- the programmer must **write the proof term** by hand, and
- term equality is **undecidable**, so checking needs a termination checker.

Refinement types keep the idea – **a type bundled with a proposition** – but restrict the proposition to a **decidable logic**, so that an **SMT solver** discharges every proof obligation **automatically**.

Every type system trades **expressiveness** against **automation**:

system	expressiveness	checking
simple types (STLC)	low	decidable, fully automatic
refinement types	medium	decidable, SMT-automated
dependent types	high	undecidable, manual proofs

Refinement types sit in the **sweet spot**: far more precise than simple types – they can express “a positive integer”, “a non-empty list”, “an in-bounds index” – yet type checking stays **fully automatic**.

The bargain: give up arbitrary propositions, gain a **push-button** checker.

1. Why Refinement Types?
2. Base Refinement Types
Syntax and Intuition
3. Subtyping as Implication
4. Refining Functions
5. Liquid Types
6. Refinement Types in Practice

A **refinement type** refines a base type B with a **predicate** p – a boolean-valued expression that may mention the bound variable x :

$$\{x : B \mid p\}$$

read as “the values x of base type B **such that** p holds”.

Its inhabitants are exactly the values whose substitution makes p true:

$$\llbracket \{x : B \mid p\} \rrbracket = \{v \in \llbracket B \rrbracket \mid p[x \mapsto v] = \mathbf{true}\}$$

A plain base type is the **trivial** refinement, and \perp is the **empty** one:

$$B \triangleq \{x : B \mid \mathbf{true}\} \quad \perp \triangleq \{x : B \mid \mathbf{false}\}$$

$\text{Nat} \triangleq \{x : \text{num} \mid x \geq 0\}$	the non-negative integers
$\text{Pos} \triangleq \{x : \text{num} \mid x > 0\}$	the positive integers
$\text{Even} \triangleq \{x : \text{num} \mid x \bmod 2 = 0\}$	the even integers
$\text{Rng } n \triangleq \{x : \text{num} \mid 0 \leq x \wedge x < n\}$	an in-bounds index
$\{x : \text{num} \mid x = 42\}$	a singleton type

In **LiquidHaskell**, such types are written as ordinary signatures whose base types are decorated with predicates (here shown without the surrounding annotation syntax):

```
type Nat    = {v:Int | v >= 0}
type Pos    = {v:Int | v >  0}
type Rng N  = {v:Int | 0 <= v && v < N}
```

1. Why Refinement Types?
2. Base Refinement Types
Syntax and Intuition
3. Subtyping as Implication
4. Refining Functions
5. Liquid Types
6. Refinement Types in Practice

Recall the **subsumption** rule (Lecture 18): if $e : \tau$ and $\tau <: \tau'$, then $e : \tau'$.

For refinements, $\{x : B \mid p\}$ is a subtype of $\{x : B \mid q\}$ exactly when **every value satisfying p also satisfies q** :

$$\frac{\models \forall x. (p \implies q)}{\{x : B \mid p\} <: \{x : B \mid q\}} \text{SUB-BASE}$$

So subtyping reduces to checking the **validity of an implication** – a job for an **SMT solver**. For instance,

$$\text{Pos} <: \text{Nat} \quad \text{because} \quad \models \forall x. (x > 0 \implies x \geq 0).$$

The implications collected during type checking are called **verification conditions** (VCs). They live in **decidable theories**:

- linear arithmetic over integers and reals,
- equality with uninterpreted functions (EUF),
- arrays, bit-vectors, ...

An **SMT solver** (Z3, cvc5, ...) decides each VC. To check $\{x \mid x = 2 + 3\} <: \{x \mid x > 4\}$, it asks whether the **negation** of the VC is satisfiable:

```
(declare-const x Int)
(assert (not (=> (= x (+ 2 3)) (> x 4))))
(check-sat) ; unsat => the implication is valid
```

An **unsat** result means **no counterexample exists**, so the VC is **valid** and the subtyping holds.

1. Why Refinement Types?
2. Base Refinement Types
Syntax and Intuition
3. Subtyping as Implication
4. Refining Functions
5. Liquid Types
6. Refinement Types in Practice

Just like Π -types, a function argument can be **named** so the result type refers to it:

$$(x : \{y : B \mid p\}) \rightarrow \{z : B' \mid q\} \quad \text{where } q \text{ may mention } x.$$

Such refinements act as **pre-** and **postconditions** on functions:

```
abs :: Int -> {v:Int | v >= 0} -- return value is non-negative
div :: Int -> {v:Int | v /= 0} -> Int -- divisor is non-zero
(!) :: a:[b] -> {i:Int | 0 <= i && i < len a} -> b
```

- `div`'s precondition rules out **division by zero**;
- `(!)`'s precondition rules out **out-of-bounds** indexing.

Each is checked by the solver at **every call site**.

The type environment Γ now carries **logical facts**, not just bindings. Subtyping flows them into the VC:

$$\frac{\models \llbracket \Gamma \rrbracket \wedge p \implies q}{\Gamma \vdash \{x : B \mid p\} <: \{x : B \mid q\}} \text{SUB}$$

where $\llbracket \Gamma \rrbracket$ is the **conjunction of the refinement predicates** of all variables in scope.

A conditional **strengthens** the context with the branch condition – this is **path-sensitivity**:

$$\frac{\Gamma \vdash e_1 : \text{bool} \quad \Gamma, e_1 \vdash e_2 : \tau \quad \Gamma, \neg e_1 \vdash e_3 : \tau}{\Gamma \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \tau} \text{IF}$$

Inside the THEN-branch we may **assume** e_1 holds; inside the ELSE-branch, that it does not.

Example: Checking abs

We claim `abs` always returns a `Nat`, i.e. a value ≥ 0 :

```
abs :: Int -> Nat    -- Nat = {v:Int | v >= 0}
abs x = if x >= 0 then x else 0 - x
```

Two branches give **two verification conditions**, each demanding the result satisfy $v \geq 0$:

branch	context	VC
then : x	$x \geq 0$	$x \geq 0 \implies x \geq 0$ ✓
else : $0 - x$	$\neg(x \geq 0)$	$\neg(x \geq 0) \implies (0 - x) \geq 0$ ✓

Both VCs are **valid** in linear arithmetic, so `abs` type-checks – with **no manual proof**. Drop the `else` guard and the second VC becomes $\text{true} \implies (0 - x) \geq 0$, which the solver **refutes** with the counterexample $x = 1$.

1. Why Refinement Types?
2. Base Refinement Types
Syntax and Intuition
3. Subtyping as Implication
4. Refining Functions
5. Liquid Types
6. Refinement Types in Practice

Annotating **every** type by hand is tedious. **Liquid types** (Logically Qualified Data types) **infer** the refinements automatically:

- 1 Fix a finite set of **qualifiers** – predicate templates such as $\{0 \leq \star, \star < \star, \star = \star\}$.
- 2 Attach an **unknown** refinement variable κ to each type to be inferred.
- 3 Type checking emits **Horn constraints** (implications) over the κ 's.
- 4 Solve them by **predicate abstraction** and **least-fixed-point** iteration: pick the strongest conjunction of qualifiers satisfying every constraint.

The refinements are as expressive as the qualifiers allow, yet found **automatically** – the very **fixed-point** machinery we used for static analysis, now over a logical lattice.

Take abs again, but **drop** the result annotation: put an **unknown** refinement κ in its place, and fix a qualifier set Q .

$$\text{abs} : \text{num} \rightarrow \{ v : \text{num} \mid \kappa \} \qquad Q = \{ 0 \leq v, v \leq 0, v = 0 \}$$

Checking the body `if $x \geq 0$ then x else $0 - x$` emits two **Horn constraints** on κ – one per branch, with its path condition:

$$\begin{aligned} \text{then } (x) : \quad & x \geq 0 \wedge v = x \implies \kappa \\ \text{else } (0 - x) : \quad & x < 0 \wedge v = 0 - x \implies \kappa \end{aligned}$$

Solve for the **strongest** conjunction of qualifiers in Q entailed by **both** constraints. Only $0 \leq v$ survives ($v \leq 0$ and $v = 0$ each fail a branch):

$$\kappa = (0 \leq v) \implies \text{abs} : \text{num} \rightarrow \{ v : \text{num} \mid 0 \leq v \} = \text{num} \rightarrow \text{Nat}$$

The annotation `Nat` from the previous section was **inferred** – not written.

1. Why Refinement Types?
2. Base Refinement Types
Syntax and Intuition
3. Subtyping as Implication
4. Refining Functions
5. Liquid Types
6. Refinement Types in Practice

Many modern verifiers reduce types to **SMT queries**:

Tool	Host	Notes
LiquidHaskell	Haskell	liquid-type inference over Haskell
F*	ML-like	refinements and full dependent types
Flux	Rust	refinements layered on ownership
Stainless	Scala	contract verification
Dafny	imperative	SMT-backed pre/postconditions

All share one engine: **elaborate types into verification conditions and hand them to a solver**. The solver is the trusted, automated proof-finder.

Two points on the same Curry-Howard landscape:

	Refinement Types	Dependent Types
proposition	decidable logic	arbitrary
proof	SMT, automatic	manual proof term
checking	decidable	undecidable
expressiveness	medium	maximal

They are **complementary**, not rivals. F^* combines both: it uses refinements for everything an SMT solver can settle, and falls back to full dependent-type proofs only where automation gives up.

Next, we follow the **path-sensitive** idea from the IF rule further, to types that **change as control flows**.

- Flow-Sensitive Types

Jihyeok Park
jihyeok_park@korea.ac.kr
<https://plrg.korea.ac.kr>