

A Framework for the Automatic Formal Verification of Refinement from COGENT to C

Christine Rizkallah⁴, Japheth Lim¹, Yutaka Nagashima¹,
Thomas Sewell^{1,2}, Zilin Chen^{1,2}, Liam O'Connor^{1,2}, Toby
Murray^{1,3}, Gabriele Keller^{1,2}, and Gerwin Klein^{1,2}

¹ Data61 (formerly NICTA), CSIRO, Sydney, Australia

² University of New South Wales, Sydney, Australia

³ University of Melbourne, Australia

⁴ University of Pennsylvania, Philadelphia, PA, USA

ITP, Nancy, France

24th of August 2016

COGENT Project: Overall Story

Motivation

File systems are too important to remain unverified

COGENT Project: Overall Story

Motivation

File systems are too important to remain unverified

Problem

There are many file systems and they are huge (\approx seL4)

COGENT Project: Overall Story

Motivation

File systems are too important to remain unverified

Problem

There are many file systems and they are huge (\approx seL4)

Aim

Create method for feasibly verifying file systems

COGENT Project: Overall Story

Motivation

File systems are too important to remain unverified

Problem

There are many file systems and they are huge (\approx seL4)

Aim

Create method for feasibly verifying file systems

Idea

Automate large portions of the verification process

COGENT Project: Overall Story

Motivation

File systems are too important to remain unverified

Problem

There are many file systems and they are huge (\approx seL4)

Aim

Create method for feasibly verifying file systems

Idea

Automate large portions of the verification process

Sounds good, but **HOW?**

Languages, Type Systems, and Certifying Compilers

Languages, Type Systems, and Certifying Compilers

- ▶ Purely functional languages allow for an easier, more direct style of formal reasoning about code

Languages, Type Systems, and Certifying Compilers

- ▶ Purely functional languages allow for an easier, more direct style of formal reasoning about code
- ▶ Design a restricted purely functional language, expressive enough to describe most of the file system code

Languages, Type Systems, and Certifying Compilers

- ▶ Purely functional languages allow for an easier, more direct style of formal reasoning about code
- ▶ Design a restricted purely functional language, expressive enough to describe most of the file system code
- ▶ Use type systems techniques (linear types) to enforce properties such as memory safety at compile time

Languages, Type Systems, and Certifying Compilers

- ▶ Purely functional languages allow for an easier, more direct style of formal reasoning about code
- ▶ Design a restricted purely functional language, expressive enough to describe most of the file system code
- ▶ Use type systems techniques (linear types) to enforce properties such as memory safety at compile time
- ▶ Create a certifying compiler to automate the low level “boring” parts of the verification

COGENT Project: Results in a Nutshell

- ▶ COGENT is a functional language with a linear type system.

COGENT Project: Results in a Nutshell

- ▶ COGENT is a **functional language** with a **linear type system**.
- ▶ COGENT's **certifying compiler** automatically generates:
 1. efficient C code
 2. a formal model of the code that is easy to reason about
 3. a machine-checked proof linking the two (in Isabelle)

COGENT Project: Results in a Nutshell

- ▶ COGENT is a **functional language** with a **linear type system**.
- ▶ COGENT's **certifying compiler** automatically generates:
 1. efficient C code
 2. a formal model of the code that is easy to reason about
 3. a machine-checked proof linking the two (in Isabelle)
- ▶ implemented **BilbyFs** and **ext2** file systems in COGENT
- ▶ verified functional correctness of key **BilbyFs** operations

COGENT Project: Results in a Nutshell

- ▶ COGENT is a **functional language** with a **linear type system**.
- ▶ COGENT's **certifying compiler** automatically generates:
 1. efficient C code
 2. a formal model of the code that is easy to reason about
 3. a machine-checked proof linking the two (in Isabelle)
- ▶ implemented **BilbyFs** and **ext2** file systems in COGENT
- ▶ verified functional correctness of key **BilbyFs** operations
 - ▶ relied on **automating** large parts of the verification

How does the Certifying Compiler Automatically Generate Refinement Theorems and Proofs?

How does the Certifying Compiler Automatically Generate Refinement Theorems and Proofs?

But First Let's Get Familiar with COGENT

The COGENT Language

- ▶ linearly typed, restricted, polymorphic, higher-order, purely functional language
- ▶ has typical **let**, **if**, ..., for records **take**, **put** (linearity)
- ▶ COGENT programs rely on an external library of abstract data types (ADTs) for data structures like arrays, lists, red-black trees, etc.
- ▶ loops are implemented as iterators over ADTs

Linear Type System

- ▶ Variables with a linear type must be used exactly once.
- ▶ COGENT's linear type system:
 - ▶ allows generating efficient imperative code with in-place updates
 - ▶ assists memory management
 - ▶ prevent errors such as use-after-free, memory leaks, pointer mismanagement in error-handling, etc.

COGENT Example: Flip

```
flip :: {f :: U8} w → {f :: U8} w  
flip x =  
  take  $x'$  {f = y} = x  
  in if  $y == 0$   
    then put  $x'.f := 1$   
    else put  $x'.f := 0$ 
```

Linear Type System: Examples

Variables with a linear type must be used exactly once.

```
let  $x = allocData$   
and  $y = x$   
and  $_ = free\ x$   
in  $y$ 
```

Linear Type System: Examples

Variables with a linear type must be used exactly once.

```
let  $x = allocData$   
and  $y = \mathbf{x}$   
and  $_ = free \mathbf{x}$   
in  $y$ 
```

Linear Type System: Examples

Variables with a linear type must be used exactly once.

```
let  $x$  = allocData  
in ()
```

Linear Type System: Examples

Variables with a linear type must be used exactly once.

```
let  $x = allocData$   
in ()
```


Linear Type System: Examples

Variables with a linear type must be used exactly once.

```
let  $x = allocData$   
in if  $condition$  then  $Some\ x$   
      else  $None$ 
```

Variables with a linear type must be freed or returned.

Linear Type System: Examples

Variables with a linear type must be used exactly once.

```
let  $x = allocData$   
in if  $condition$  then  $Some\ x$   
      else  $None$ 
```

Variables with a linear type must be freed or returned.

Linear Type System: Examples

The linear type system enables the conversion from purely functional to imperative code with in-place memory update.

```
let  $x = \text{allocData}$   
and  $y = \text{updateData } x$   
in  $y$ 
```

Linear Type System: Examples

The linear type system enables the conversion from purely functional to imperative code with in-place memory update.

```
let  $x = \text{allocData}$   
and  $y = \text{updateData } x$   
in  $y$ 
```

The C code uses the same variables in memory for x and y .

Linear Type System: Examples

The linear type system enables the conversion from purely functional to imperative code with in-place memory update.

```
let  $x = \mathit{allocData}$   
and  $y = \mathit{updateData}\ x$   
in  $y$ 
```

The C code uses the same variables in memory for x and y .

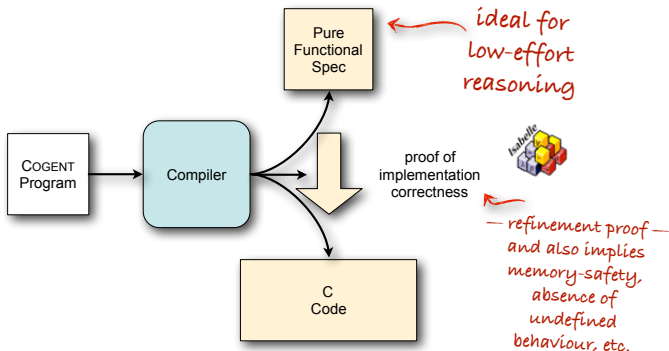
COGENT needs no garbage collector.

Why can we get away with type system and restrictions?

Idea

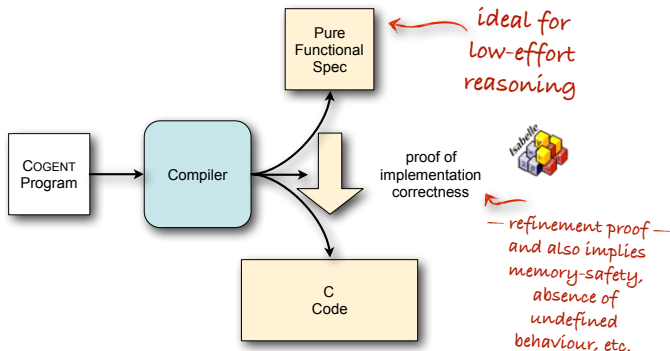
- ▶ Write most code in COGENT to simplify its verification
- ▶ Write a small part in C and verify it manually

COGENT: Certifying Compiler



- ▶ **Certifying compiler:** certifies correctness of its compilation.

COGENT: Certifying Compiler



- ▶ **Certifying compiler:** certifies correctness of its compilation.
- ▶ **COGENT's** certifying compiler **automatically** generates:
 1. efficient C code
 2. a formal HOL spec. of the code that is easy to reason about
 3. a machine-checked proof linking the two (in Isabelle)

COGENT: Certifying Compiler

Refinement

- ▶ The compiler uses an underlying routine to
 - ▶ discharge cumbersome C safety obligations and
 - ▶ provide a HOL emb. more amenable to verification.

COGENT: Certifying Compiler

Refinement

- ▶ The compiler uses an underlying routine to
 - ▶ discharge cumbersome C safety obligations and
 - ▶ provide a HOL emb. more amenable to verification.
- ▶ **Theorem:** When C executes to a value, HOL spec. evaluates similarly.

COGENT: Certifying Compiler

Refinement

- ▶ The compiler uses an underlying routine to
 - ▶ discharge cumbersome C safety obligations and
 - ▶ provide a HOL emb. more amenable to verification.
- ▶ **Theorem:** When C executes to a value, HOL spec. evaluates similarly.
- ▶ Hence, proofs about HOL spec. also hold for C code.

COGENT: Certifying Compiler

Refinement

- ▶ The compiler uses an underlying routine to
 - ▶ discharge cumbersome C safety obligations and
 - ▶ provide a HOL emb. more amenable to verification.
- ▶ **Theorem:** When C executes to a value, HOL spec. evaluates similarly.
- ▶ Hence, proofs about HOL spec. also hold for C code.
- ▶ This means we can now prove theorems using low-effort equational reasoning on HOL spec. rather than deal with tedious C, and the theorems also hold for the C code!

COGENT: Certifying Compiler

Refinement

- ▶ The compiler uses an underlying routine to
 - ▶ discharge cumbersome C safety obligations and
 - ▶ provide a HOL emb. more amenable to verification.
- ▶ **Theorem:** When C executes to a value, HOL spec. evaluates similarly.
- ▶ Hence, proofs about HOL spec. also hold for C code.
- ▶ This means we can now prove theorems using low-effort equational reasoning on HOL spec. rather than deal with tedious C, and the theorems also hold for the C code!

Performance

- ▶ The performance of the generated C is similar to that of hand written C.

COGENT Compiler uses Pre-Existing Tools

Simpl [Schirmer 2005]

- ▶ Imperative language embedded into Isabelle/HOL

C-Parser [Norrish 2012]

- ▶ straightforward translation from C to Simpl

AutoCorres [Greenaway et al. 2012]

- ▶ converts Simpl to monadic representation in Isabelle/HOL
- ▶ verified simplifications
- ▶ output meant for manual reasoning
- ▶ we tweak AutoCorres to make its output more predictable (for automation on top)

How does the Certifying Compiler Automatically Generate Refinement Theorems and Proofs?

COGENT Example: Back to Flip

```
1  flip :: {f :: U8} w → {f :: U8} w
2  flip x =
3    take  $x'$  {f = y} = x
4    in if  $y == 0$ 
5        then put  $x'.f := 1$ 
6        else put  $x'.f := 0$ 
```


COGENT Example: Generated HOL and C Flip

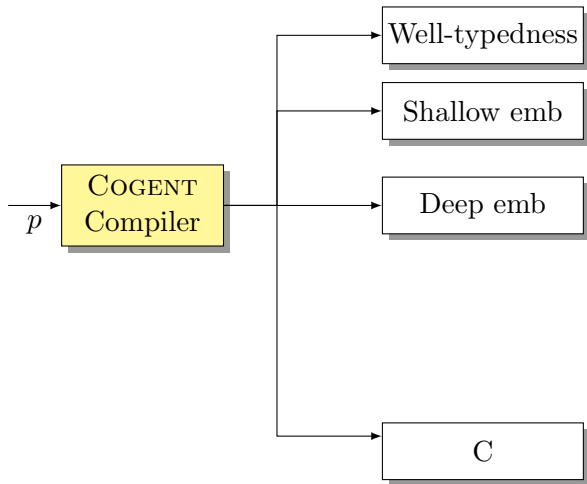
Functional Specification (HOL)

```
1 flip :: {f :: U8} → {f :: U8}
2 flip x =
3   take x' {f = y} = x
4   .
5   in let tmp1 = 0
6     and tmp2 = (y == tmp1)
7     in if tmp2
8       then let tmp3 = 1
9         and x'' = put x'.f := tmp3
10        .
11        in x''
12     else let tmp4 = 0
13       and x'' = put x'.f := tmp4
14        .
15        in x''
16     .
17     .
```

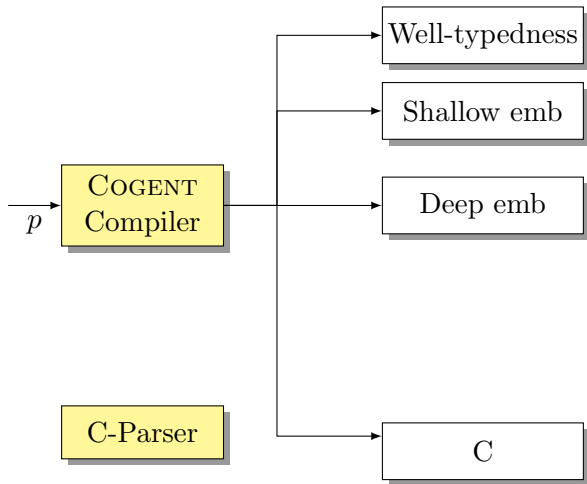
C code (monadic representation)

```
1 flipC :: rec1 ptr ⇒ (rec1 ptr, σ) nondet_monad
2 flipC x = do
3   guard (λσ. is-valid σ x);
4   y ← gets (λσ. σ[r].f);
5   tmp1 ← return 0;
6   tmp2 ← return bool (y = tmp1);
7   tmpresult ← condition (bool tmp2 ≠ 0)
8     (do tmp3 ← return 1;
9       guard (λσ. is-valid σ x);
10      modify (λσ. σ[x].f := tmp3);
11      return x od)
12   (do tmp4 ← return 0;
13     guard (λσ. is-valid σ x);
14     modify (λσ. σ[x].f := tmp4);
15     return x od);
16 return tmpresult
17 od
```

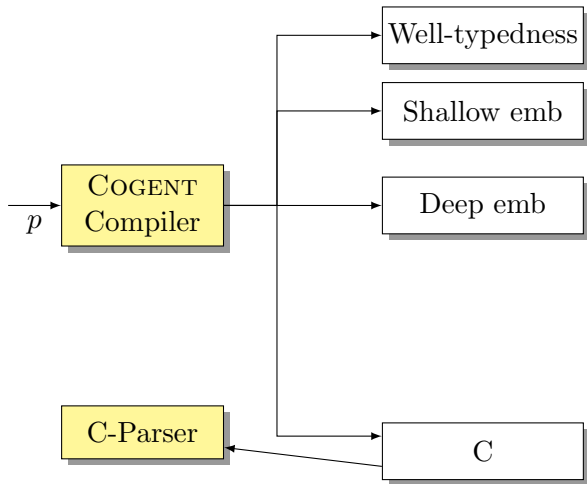
Overall Project: How did we structure refinement?



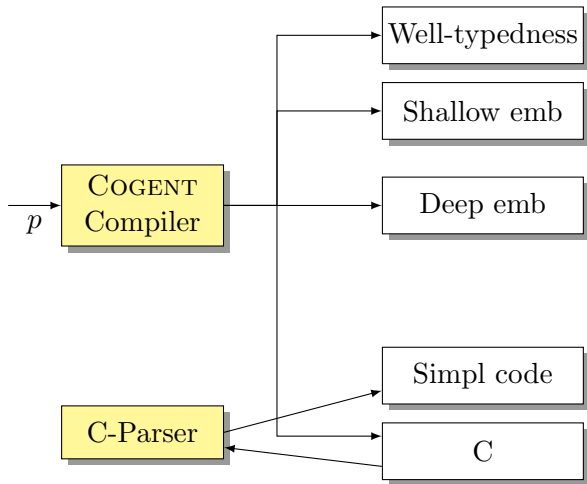
Overall Project: How did we structure refinement?



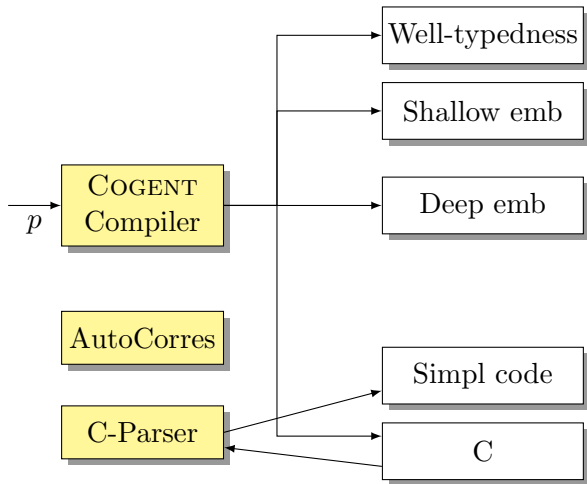
Overall Project: How did we structure refinement?



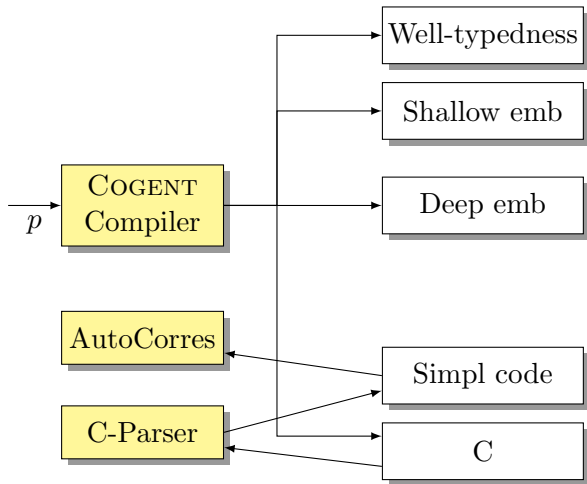
Overall Project: How did we structure refinement?



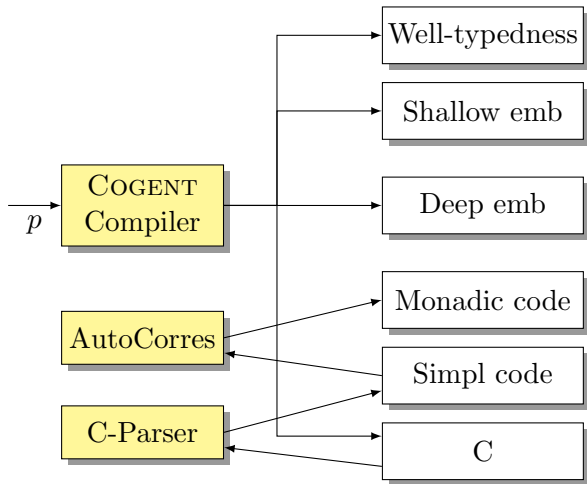
Overall Project: How did we structure refinement?



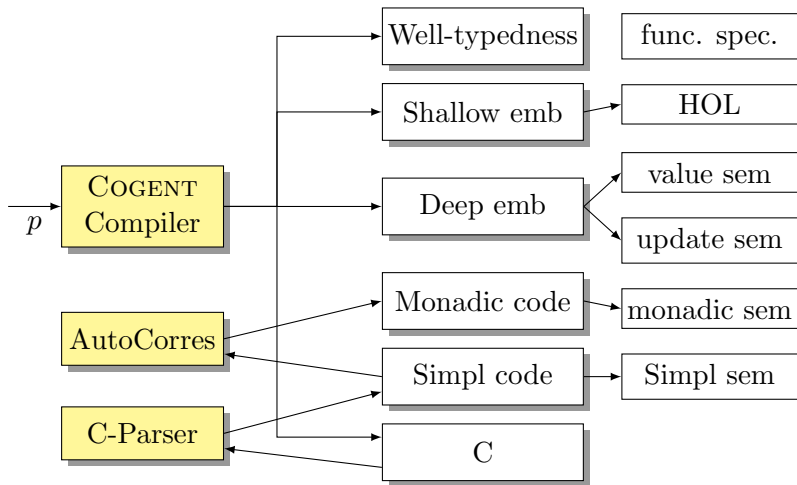
Overall Project: How did we structure refinement?



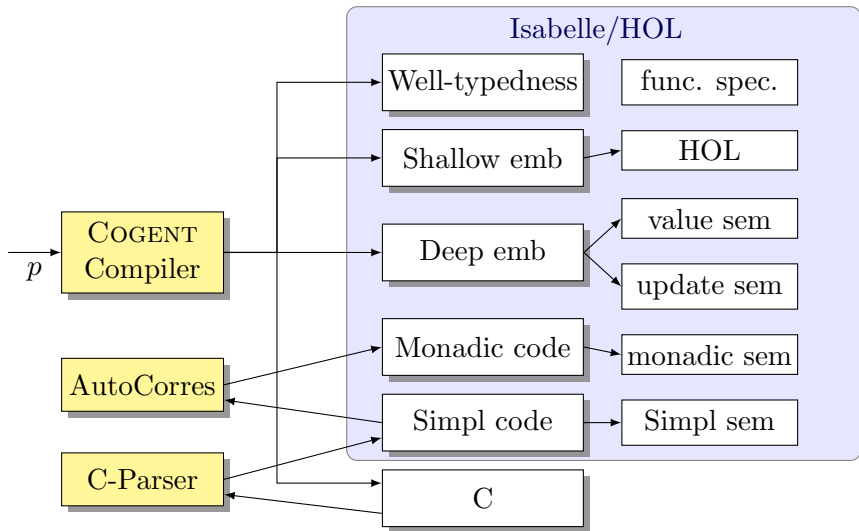
Overall Project: How did we structure refinement?



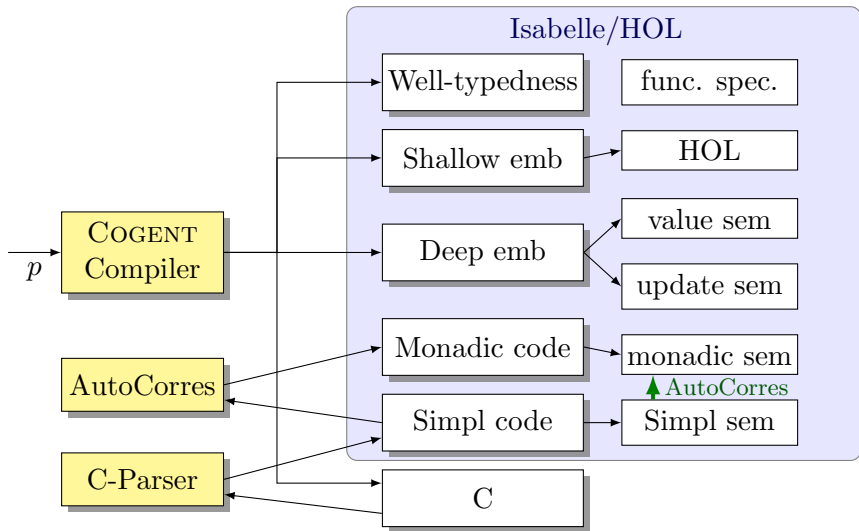
Overall Project: How did we structure refinement?



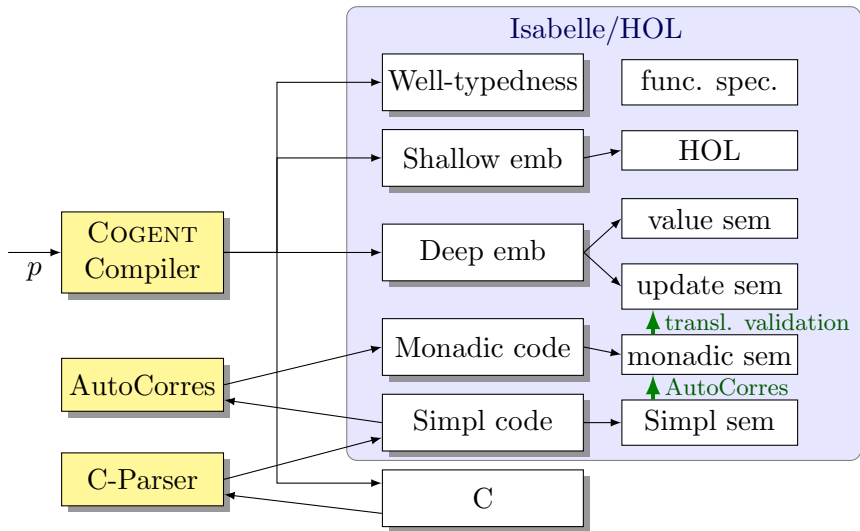
Overall Project: How did we structure refinement?



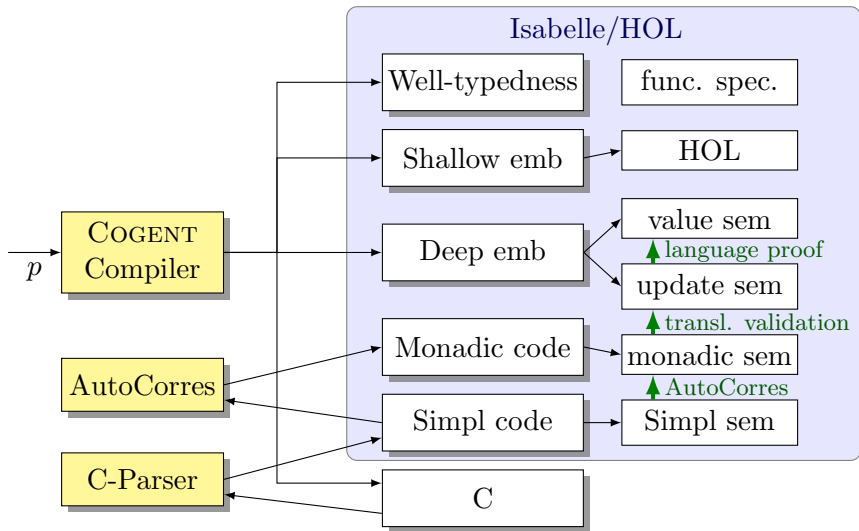
Overall Project: How did we structure refinement?



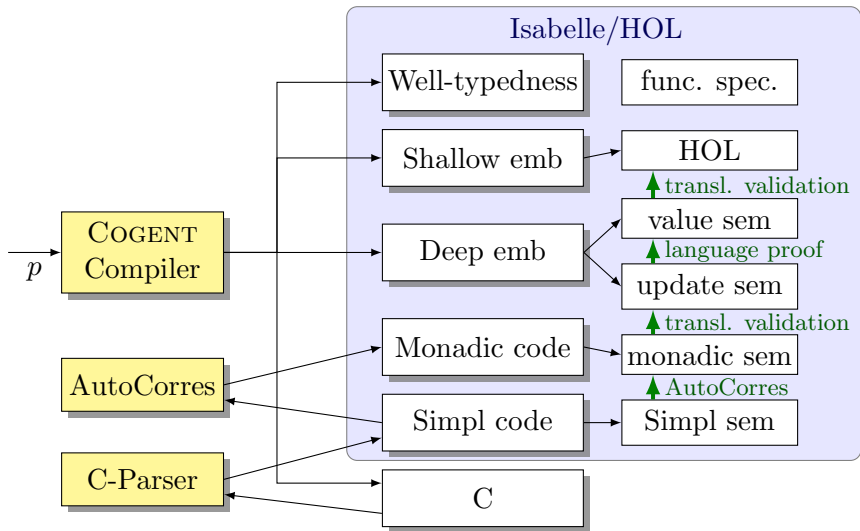
Overall Project: How did we structure refinement?



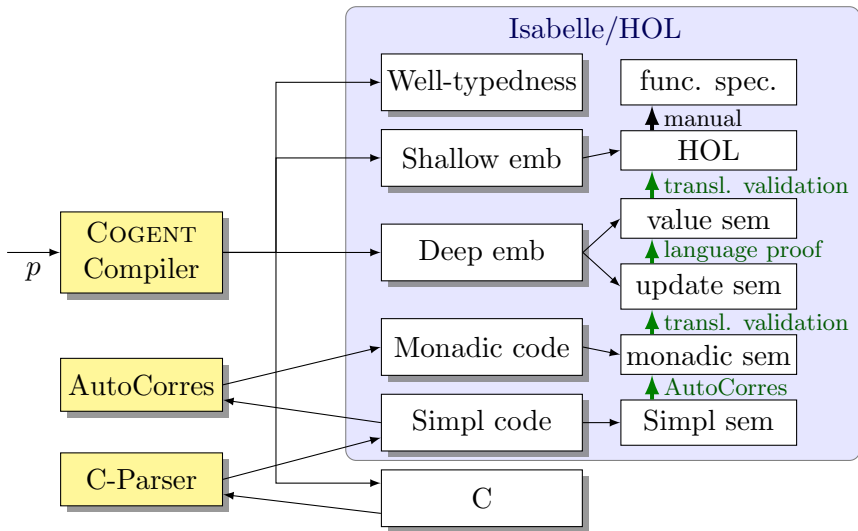
Overall Project: How did we structure refinement?



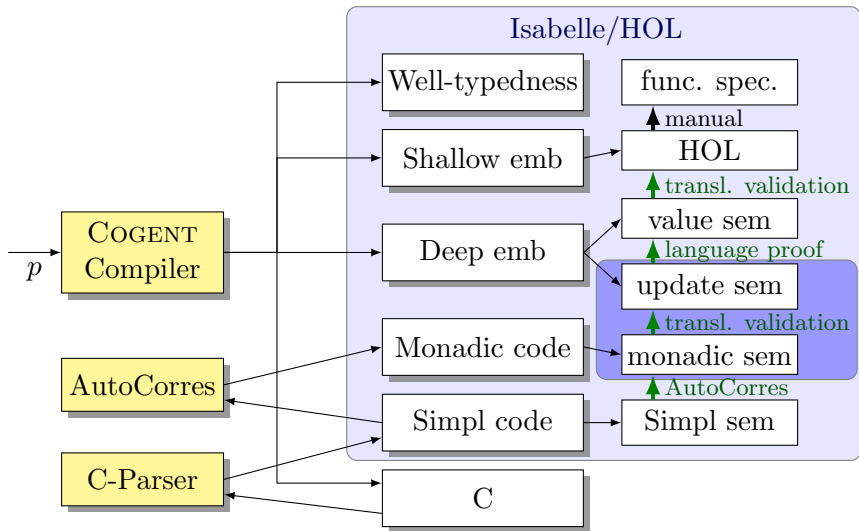
Overall Project: How did we structure refinement?



Overall Project: How did we structure refinement?



Overall Project: How did we structure refinement?



How did we prove refinement (update to monadic sem.)?

Data Relations and Refinement Predicate

- ▶ values
- ▶ types

- ▶ states
- ▶ statements

Theories and Automation

- ▶ refinement calculus
- ▶ well-typedness proof
- ▶ syntax directed proof automation

How did we prove refinement (update to monadic sem.)?

Data Relations and Refinement Predicate

- ▶ values
- ▶ types
 - ▶ several COGENT types correspond to the same C type
- ▶ states
- ▶ statements

Theories and Automation

- ▶ refinement calculus
- ▶ well-typedness proof
- ▶ syntax directed proof automation

How did we prove refinement (update to monadic sem.)?

Data Relations and Refinement Predicate

- ▶ values
- ▶ types
 - ▶ several COGENT types correspond to the same C type
 - ▶ partial type erasure removes linearity from COGENT type
- ▶ states
- ▶ statements ($\mathbf{corres} \ R \ e \ p_m \ U \ \Gamma \ \mu \ \sigma$)

Theories and Automation

- ▶ refinement calculus
- ▶ well-typedness proof
- ▶ syntax directed proof automation

Refinement Calculus: VAR rule

$$\frac{(x \mapsto v_u) \in U \quad \text{val-rel } v_u \ v_m}{\mathbf{corres } R \ x \ (\mathbf{return } v_m) \ U \ \Gamma \ \mu \ \sigma} \text{VAR}$$

Refinement Calculus: IF rule

$$\frac{
 \begin{array}{l}
 \Gamma_1 \vdash c : \mathbf{Bool} \quad (bool\ c' = 0 \vee bool\ c' = 1) \\
 c \text{ is a COGENT boolean equal to } (bool\ c' \neq 0) \\
 \mathbf{corres}\ R\ e_1\ e'_1\ U\ \Gamma_2\ \mu\ \sigma \quad \mathbf{corres}\ R\ e_2\ e'_2\ U\ \Gamma_2\ \mu\ \sigma
 \end{array}
 }{
 \begin{array}{l}
 \mathbf{corres}\ R\ (\mathbf{if}\ c\ \mathbf{then}\ e_1\ \mathbf{else}\ e_2) \\
 (\mathbf{do}\ x \leftarrow \mathbf{condition}\ (bool\ c' \neq 0)\ e'_1\ e'_2; \mathbf{return}\ x\ \mathbf{od})\ U\ (\Gamma_1\Gamma_2)\ \mu\ \sigma
 \end{array}
 } \text{IF}$$

Well-typedness

- ▶ COGENT compiler proves via an automated Isabelle tactic that the deep embedding of input program is well-typed

Well-typedness

- ▶ COGENT compiler proves via an automated Isabelle tactic that the deep embedding of input program is well-typed
- ▶ proving refinement requires access to typing judgements of subexpressions of the program

Well-typedness

- ▶ COGENT compiler proves via an automated Isabelle tactic that the deep embedding of input program is well-typed
- ▶ proving refinement requires access to typing judgements of subexpressions of the program
- ▶ due to linear types, it is not easy to statically infer that subexpressions of a well-typed program are well-typed

Well-typedness

- ▶ COGENT compiler proves via an automated Isabelle tactic that the deep embedding of input program is well-typed
- ▶ proving refinement requires access to typing judgements of subexpressions of the program
- ▶ due to linear types, it is not easy to statically infer that subexpressions of a well-typed program are well-typed
- ▶ our proof automation then uses well-typedness theorems to discharge typing assumptions in refinement calculus

Proof Automation (for concrete program)

- ▶ specialize refinement calculus
 - ▶ some of the rules in the calculus are pretty complicated
 - ▶ we specialize complex rules to one for each type (of record)

Proof Automation (for concrete program)

- ▶ specialize refinement calculus
 - ▶ some of the rules in the calculus are pretty complicated
 - ▶ we specialize complex rules to one for each type (of record)
- ▶ tactic compositionally applies syntax directed rules

Proof Automation (for concrete program)

- ▶ specialize refinement calculus
 - ▶ some of the rules in the calculus are pretty complicated
 - ▶ we specialize complex rules to one for each type (of record)
- ▶ tactic compositionally applies syntax directed rules
- ▶ tactic uses well-typedness theorems to discharge typing assumptions

Proof Automation (for concrete program)

- ▶ specialize refinement calculus
 - ▶ some of the rules in the calculus are pretty complicated
 - ▶ we specialize complex rules to one for each type (of record)
- ▶ tactic compositionally applies syntax directed rules
- ▶ tactic uses well-typedness theorems to discharge typing assumptions
- ▶ refinement for foreign functions is assumed

Conclusion

- ▶ We developed a compositional refinement calculus and proof rules to create a fully automatic refinement certificate from COGENT to C.
- ▶ Through co-generation of code and proofs our framework significantly reduces the cost of reasoning about efficient C.
- ▶ It does so by discharging cumbersome safety obligations and providing an embedding more amenable to verification.
- ▶ Our framework was applied to two real world file systems.

Future work

- ▶ Speed C-Parser and AutoCorres for our very large dev.
 - ▶ BilbyFs generated code is ≈ 1.8 x the size of `seL4`
- ▶ Add optimizations (additional to ones we get by using `gcc`)
- ▶ Verify a library of abstract data types

Meet the other Verification/PL Folks



Japheth
Lim




Yutaka
Nagashima



Thomas
Sewell



Zilin
Chen



Liam
O'Connor



Toby
Murray



Gabi Keller



Gerwin
Klein

Now that you know about COGENT,
consider trying it out!

COGENT: co-generation of code and proofs
significantly reduces the cost of reasoning about
efficient C.

Thanks for Listening, Questions?