

A Layered Certifying Compiler Architecture

Jacco Krijnen

PhD advisors: Wouter Swierstra, Manuel Chakravarty, Gabriele Keller

IFIP WG 2.8 (24 April 2024)

Utrecht University



How do you structure the verification of a compiler?

- Classic approach to compiler verification (CompCert, CakeML)

- Classic approach to compiler verification (CompCert, CakeML)
 - Compiler is implemented in a proof assistant

- Classic approach to compiler verification (CompCert, CakeML)
 - Compiler is implemented in a proof assistant
 - Target fixed, well-understood languages

Motivation

- Classic approach to compiler verification (CompCert, CakeML)
 - Compiler is implemented in a proof assistant
 - Target fixed, well-understood languages
- Many compilers and languages are:

- Classic approach to compiler verification (CompCert, CakeML)
 - Compiler is implemented in a proof assistant
 - Target fixed, well-understood languages
- Many compilers and languages are:
 - Implemented in language not suitable for verification

- Classic approach to compiler verification (CompCert, CakeML)
 - Compiler is implemented in a proof assistant
 - Target fixed, well-understood languages
- Many compilers and languages are:
 - Implemented in language not suitable for verification
 - Community-developed, open-source: in constant flux

Motivation

- Classic approach to compiler verification (CompCert, CakeML)
 - Compiler is implemented in a proof assistant
 - Target fixed, well-understood languages
- Many compilers and languages are:
 - Implemented in language not suitable for verification
 - Community-developed, open-source: in constant flux
- Instead of proving the compiler code correct, prove individual runs correct (translation validation)

Approach

Approach

We take a “gray-box” approach



- Tool that runs simultaneously with the compiler

Approach

We take a “gray-box” approach



- Tool that runs simultaneously with the compiler
- Compiler dumps the program between every pass

Approach

We take a “gray-box” approach



- Tool that runs simultaneously with the compiler
- Compiler dumps the program between every pass
- Tool checks if each pass was performed correctly

Approach

We take a “gray-box” approach



- Tool that runs simultaneously with the compiler
- Compiler dumps the program between every pass
- Tool checks if each pass was performed correctly
- The tool should produce some checkable evidence

Layered Architecture

- For each compiler pass, develop four layers:
 1. Specification layer
 2. Interface layer
 3. Automation layer
 4. Verification layer
- Implemented in a proof assistant (here Coq)
- Each layer gradually improves the trustworthiness of the overall system

$$\begin{array}{l} t ::= x \\ \quad | \lambda x. t \\ \quad | t t \\ \quad | \text{let } x = t \text{ in } t \end{array}$$

Inliner behaviour

“pre-term”

```
let x = 3
```

```
in ... (2 * x) ... x ...
```

Inliner behaviour

“pre-term”

```
let x = 3
```

```
in ... (2 * x) ... x ...
```

↳ “post-term”

```
let x = 3
```

```
in ... (2 * 3) ... x ...
```

Inliner: specification layer

Specification layer (translation relation)

$$\boxed{\Gamma \vdash s \triangleright t}$$

Specification layer (translation relation)

$$\boxed{\Gamma \vdash s \triangleright t}$$

$$\Gamma ::= \bullet \mid x \mapsto t, \Gamma \mid x, \Gamma$$

Specification layer (translation relation)

$$\boxed{\Gamma \vdash s \triangleright t}$$

$$\Gamma ::= \bullet \mid x \mapsto t, \Gamma \mid x, \Gamma$$

$$\frac{\Gamma(x) = t}{\Gamma \vdash x \triangleright t} \text{ INLINE-VAR}_1$$

$$\frac{}{\Gamma \vdash x \triangleright x} \text{ INLINE-VAR}_2$$

Specification layer (translation relation)

$$\boxed{\Gamma \vdash s \triangleright t}$$

$$\Gamma ::= \bullet \mid x \mapsto t, \Gamma \mid x, \Gamma$$

$$\frac{\Gamma(x) = t}{\Gamma \vdash x \triangleright t} \text{ INLINE-VAR}_1 \qquad \frac{}{\Gamma \vdash x \triangleright x} \text{ INLINE-VAR}_2$$

$$\frac{\Gamma \vdash s \triangleright s' \quad x \mapsto s, \Gamma \vdash t \triangleright t'}{\Gamma \vdash \text{let } x = s \text{ in } t \triangleright \text{let } x = s' \text{ in } t'} \text{ INLINE-LET}$$

Specification layer (translation relation)

$$\boxed{\Gamma \vdash s \triangleright t}$$

$$\Gamma ::= \bullet \mid x \mapsto t, \Gamma \mid x, \Gamma$$

$$\frac{\Gamma(x) = t}{\Gamma \vdash x \triangleright t} \text{ INLINE-VAR}_1 \qquad \frac{}{\Gamma \vdash x \triangleright x} \text{ INLINE-VAR}_2$$

$$\frac{\Gamma \vdash s \triangleright s' \quad x \mapsto s, \Gamma \vdash t \triangleright t'}{\Gamma \vdash \text{let } x = s \text{ in } t \triangleright \text{let } x = s' \text{ in } t'} \text{ INLINE-LET}$$

$$\frac{\Gamma \vdash s \triangleright s' \quad \Gamma \vdash t \triangleright t'}{\Gamma \vdash s \ t \triangleright s' \ t'} \text{ INLINE-APP}$$

Specification layer (translation relation)

$$\boxed{\Gamma \vdash s \triangleright t}$$

$$\Gamma ::= \bullet \mid x \mapsto t, \Gamma \mid x, \Gamma$$

$$\frac{\Gamma(x) = t}{\Gamma \vdash x \triangleright t} \text{ INLINE-VAR}_1 \qquad \frac{}{\Gamma \vdash x \triangleright x} \text{ INLINE-VAR}_2$$

$$\frac{\Gamma \vdash s \triangleright s' \quad x \mapsto s, \Gamma \vdash t \triangleright t'}{\Gamma \vdash \text{let } x = s \text{ in } t \triangleright \text{let } x = s' \text{ in } t'} \text{ INLINE-LET}$$

$$\frac{\Gamma \vdash s \triangleright s' \quad \Gamma \vdash t \triangleright t'}{\Gamma \vdash s \ t \triangleright s' \ t'} \text{ INLINE-APP}$$

$$\frac{x, \Gamma \vdash t \triangleright t'}{\Gamma \vdash \lambda x. t \triangleright \lambda x. t'} \text{ INLINE-LAM}$$

$$\Gamma \vdash s \triangleright t$$

- Serves as a *formal reference* for compiler implementors
- This relation does not prescribe *what* inlining strategy to use
- It captures all possible inlining strategies

Formalize in proof assistant

```
Inductive inline : list (string * binder_info) -> term -> term -> Prop
| inline_Var_1 : forall G x t t',
  lookup x G = let_bound t ->
  inline G t t' ->
  inline G (Var x) t
| inline_Var_2 : forall G x,
  inline G (Var x) (Var x)
| inline_Let : forall G x s t s' t',
  inline G s s' ->
  inline (let_bound x s :: G) t t' ->
  inline G (Let x s t) (Let x s' t')
| inline_Lam : forall G x t t',
  inline (lam_bound x :: G) t t' ->
  inline G (Lam x t) (Lam x t')
| inline_App : forall G s t s' t',
  inline G s s' ->
  inline G t t' ->
  inline G (App s t) (App s' t').
```

Inliner: interface layer

2. Interface Layer

- Bridges the gap between the compiler and the proof assistant
- Modify the compiler: dump intermediate ASTs (gray-box)
- Parse in the proof assistant

Definition `p1 := Let "x" (App (Lambda ...`

Definition `p2 := Let "x" (App (Lambda ...`

Theorem `valid : inline [] p1 p2.`

Infrastructure to formulate and prove that a pass was performed *according to specification*.

Inliner: automation layer

3. Automation Layer

- Writing proofs by hand gets tedious: large proof objects, proof search becomes slower
- Instead: prove specification decidable (reflection)

```
Fixpoint dec_inline :
```

```
  list binder_info -> AST -> AST -> bool.
```

```
...
```

```
Theorem iff_dec_inline : forall G p q,
```

```
  dec_inline G p q = true <-> inline G p q.
```

Inliner: verification layer

4. Verification Layer

- Prove that a specification is correct: related programs should be semantically equivalent
- Choose appropriate semantics, notion of program equivalence

Theorem correct : forall p q,
inline [] p q -> p \simeq q.

Scaling up: the Plutus compiler



Plutus is an optimizing compiler, used for compiling smart contracts on the Cardano blockchain.



Plutus is an optimizing compiler, used for compiling smart contracts on the Cardano blockchain.

- Smart contract: code that enforces a financial agreement, manages the flow of financial assets on a blockchain



Plutus is an optimizing compiler, used for compiling smart contracts on the Cardano blockchain.

- Smart contract: code that enforces a financial agreement, manages the flow of financial assets on a blockchain
- Compiled code is deployed on the the blockchain (once committed, immutable)



Plutus is an optimizing compiler, used for compiling smart contracts on the Cardano blockchain.

- Smart contract: code that enforces a financial agreement, manages the flow of financial assets on a blockchain
- Compiled code is deployed on the the blockchain (once committed, immutable)
- These programs have to be extremely well-behaved (DAO vulnerability: \sim \$50 million)




Plutus is an optimizing compiler, used for compiling smart contracts on the Cardano blockchain.

- Smart contract: code that enforces a financial agreement, manages the flow of financial assets on a blockchain
- Compiled code is deployed on the the blockchain (once committed, immutable)
- These programs have to be extremely well-behaved (DAO vulnerability: \sim \$50 million)
- Is the compiler correct?


The Plutus compiler pipeline

The Plutus compiler



The Plutus compiler pipeline


The Plutus compiler



- PIR = System F_{ω}^{μ} + let(rec) + ADTs

The Plutus compiler pipeline

The Plutus compiler



- PIR = System F_{ω}^{μ} + let(rec) + ADTs
- PLC = System F_{ω}^{μ}

Plutus pass: dead code elimination

1. Specification Layer (Dead Code elimination)

`let $x = t; bs$ in $s \triangleright$` DCE-ELIM

`let $x = t; bs$ in $s \triangleright$` DCE-KEEP

1. Specification Layer (Dead Code elimination)

$$\frac{}{\text{let } x = t; bs \text{ in } s \triangleright \text{let } bs' \text{ in } s'} \text{DCE-ELIM}$$

$$\frac{}{\text{let } x = t; bs \text{ in } s \triangleright \text{let } x = t'; bs' \text{ in } s'} \text{DCE-KEEP}$$

1. Specification Layer (Dead Code elimination)

$$x \notin \text{FV}(\text{let } bs' \text{ in } s')$$

$$\frac{}{\text{let } x = t; bs \text{ in } s \triangleright \text{let } bs' \text{ in } s'} \text{DCE-ELIM}$$

$$\frac{}{\text{let } x = t; bs \text{ in } s \triangleright \text{let } x = t'; bs' \text{ in } s'} \text{DCE-KEEP}$$

1. Specification Layer (Dead Code elimination)

$$\frac{x \notin \text{FV}(\text{let } bs' \text{ in } s') \quad t \in \text{value}}{\text{let } x = t; bs \text{ in } s \triangleright \text{let } bs' \text{ in } s'} \quad \text{DCE-ELIM}$$

$$\frac{}{\text{let } x = t; bs \text{ in } s \triangleright \text{let } x = t'; bs' \text{ in } s'} \quad \text{DCE-KEEP}$$

1. Specification Layer (Dead Code elimination)

$$\frac{x \notin \text{FV}(\text{let } bs' \text{ in } s') \quad t \in \text{value} \quad \text{let } bs \text{ in } s \triangleright \text{let } bs' \text{ in } s'}{\text{let } x = t; bs \text{ in } s \triangleright \text{let } bs' \text{ in } s'} \text{DCE-ELIM}$$

$$\frac{}{\text{let } x = t; bs \text{ in } s \triangleright \text{let } x = t'; bs' \text{ in } s'} \text{DCE-KEEP}$$

1. Specification Layer (Dead Code elimination)

$$\frac{x \notin \text{FV}(\text{let } bs' \text{ in } s') \quad t \in \text{value} \quad \text{let } bs \text{ in } s \triangleright \text{let } bs' \text{ in } s'}{\text{let } x = t; bs \text{ in } s \triangleright \text{let } bs' \text{ in } s'} \text{DCE-ELIM}$$

$$\frac{t \triangleright t' \quad \text{let } bs \text{ in } s \triangleright \text{let } bs \text{ in } s'}{\text{let } x = t; bs \text{ in } s \triangleright \text{let } x = t'; bs' \text{ in } s'} \text{DCE-KEEP}$$

1. Specification Layer (Dead Code elimination)

$$\frac{x \notin \text{FV}(\text{let } bs' \text{ in } s') \quad t \in \text{value} \quad \text{let } bs \text{ in } s \triangleright \text{let } bs' \text{ in } s'}{\text{let } x = t; bs \text{ in } s \triangleright \text{let } bs' \text{ in } s'} \text{DCE-ELIM}$$

$$\frac{t \triangleright t' \quad \text{let } bs \text{ in } s \triangleright \text{let } bs \text{ in } s'}{\text{let } x = t; bs \text{ in } s \triangleright \text{let } x = t'; bs' \text{ in } s'} \text{DCE-KEEP}$$

$$\frac{s \triangleright t \quad t \triangleright t'}{s \ t \triangleright s' \ t'} \text{DCE-APPLY} \quad \frac{t \triangleright t'}{\lambda x : \tau. t \triangleright \lambda x : \tau. t'} \text{DCE-LAMBDA} \quad \dots$$

Specification Layer: Scaling up

- More syntax \Rightarrow More rules

Specification Layer: Scaling up

- More syntax \Rightarrow More rules
- Scoping of letrec

$$\frac{x \notin \text{FV}(\text{let rec } bs'_0 \text{ in } s') \quad t \in \text{value} \quad \text{let rec } bs \text{ in } s \triangleright_{rec} \text{let rec } bs' \text{ in } s' \mid bs'_0}{\text{let rec } x = t; bs \text{ in } s \triangleright_{rec} \text{let rec } bs' \text{ in } s' \mid bs'_0} \text{DCE-ELIM-REC}$$

Specification Layer: Scaling up

- More syntax \Rightarrow More rules
- Scoping of letrec

$$\frac{x \notin \text{FV}(\text{let rec } bs'_0 \text{ in } s') \quad t \in \text{value} \quad \text{let rec } bs \text{ in } s \triangleright_{rec} \text{let rec } bs' \text{ in } s' \mid bs'_0}{\text{let rec } x = t; bs \text{ in } s \triangleright_{rec} \text{let rec } bs' \text{ in } s' \mid bs'_0} \text{DCE-ELIM-REC}$$

- Composite pass: inliner includes dead code elimination

$$t \triangleright t' := \exists t_1. t \triangleright_{inl} t_1 \wedge t_1 \triangleright_{dce} t'$$

2. Interface Layer

- Plutus implementation and process are open-source: fork + pull request on Github
- Dump the intermediate ASTs (and sometimes pass-specific information)

3. Automation Layer

Decision procedure dead-code elimination:

Fixpoint dec_dce : AST -> AST -> **bool**.

Theorem iff_dec_dce : **forall** p q,
dec_dce p q = **true** <-> dce p q.

3. Automation Layer

Composite pass: Inlining + Dead Code Elimination

$$t \triangleright t' := \exists t_1. t \triangleright_{inl} t_1 \wedge t_1 \triangleright_{dce} t'$$

Intermediate term t_1 is not dumped by the compiler...

...But we can reconstruct it!

```
Fixpoint make_intermediate :
```

```
  list Var -> AST -> AST -> option AST.
```

```
...
```


4. Verification Layer

We formalised PIR (PLC) semantics in Coq

4. Verification Layer

We formalised PIR (PLC) semantics in Coq

- Type system $\Delta; \Gamma \vdash t : \tau$ with type normalisation

4. Verification Layer

We formalised PIR (PLC) semantics in Coq

- Type system $\Delta; \Gamma \vdash t : \tau$ with type normalisation
- Static semantics preservation:

$$s \triangleright t \wedge \Delta; \Gamma \vdash s : \tau \Rightarrow \Delta; \Gamma \vdash t : \tau$$

4. Verification Layer

We formalised PIR (PLC) semantics in Coq

- Type system $\Delta; \Gamma \vdash t : \tau$ with type normalisation
- Static semantics preservation:

$$s \triangleright t \wedge \Delta; \Gamma \vdash s : \tau \Rightarrow \Delta; \Gamma \vdash t : \tau$$

- Big-step semantics with substitution: $t \Downarrow v$

4. Verification Layer

We formalised PIR (PLC) semantics in Coq

- Type system $\Delta; \Gamma \vdash t : \tau$ with type normalisation
- Static semantics preservation:

$$s \triangleright t \wedge \Delta; \Gamma \vdash s : \tau \Rightarrow \Delta; \Gamma \vdash t : \tau$$

- Big-step semantics with substitution: $t \Downarrow v$
- Contextual equivalence: $\Delta; \Gamma \vdash s \simeq_{ctx} t : \tau$

4. Verification Layer

We formalised PIR (PLC) semantics in Coq

- Type system $\Delta; \Gamma \vdash t : \tau$ with type normalisation
- Static semantics preservation:

$$s \triangleright t \wedge \Delta; \Gamma \vdash s : \tau \Rightarrow \Delta; \Gamma \vdash t : \tau$$

- Big-step semantics with substitution: $t \Downarrow v$
- Contextual equivalence: $\Delta; \Gamma \vdash s \simeq_{ctx} t : \tau$
- Step-indexed logical relation that implies contextual equivalence

4. Verification Layer

We formalised PIR (PLC) semantics in Coq

- Type system $\Delta; \Gamma \vdash t : \tau$ with type normalisation
- Static semantics preservation:

$$s \triangleright t \wedge \Delta; \Gamma \vdash s : \tau \Rightarrow \Delta; \Gamma \vdash t : \tau$$

- Big-step semantics with substitution: $t \Downarrow v$
- Contextual equivalence: $\Delta; \Gamma \vdash s \simeq_{ctx} t : \tau$
- Step-indexed logical relation that implies contextual equivalence
- Correctness of translation relation

$$s \triangleright t \wedge \Delta; \Gamma \vdash s : \tau \Rightarrow \Delta; \Gamma \vdash s \simeq_{ctx} t : \tau$$

Certificates

A Certificate



```
0101  
0101  
0101
```

A Certificate



A Certificate



We want a *verifiable link* (certificate) that

A Certificate



We want a *verifiable link* (certificate) that

1. Relates the source code to the compiled code

A Certificate



We want a *verifiable link* (certificate) that

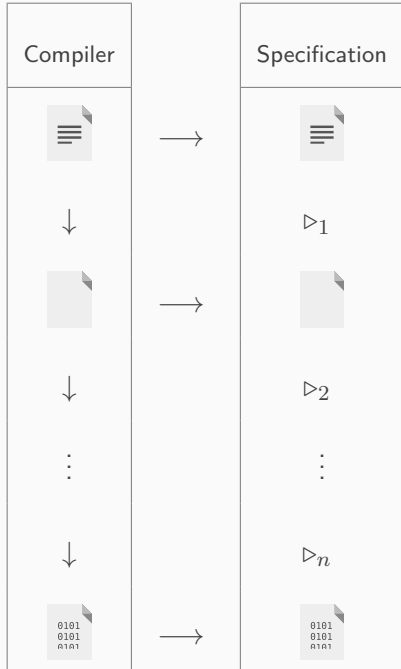
1. Relates the source code to the compiled code
2. The compiler did not introduce any bugs


Compiler

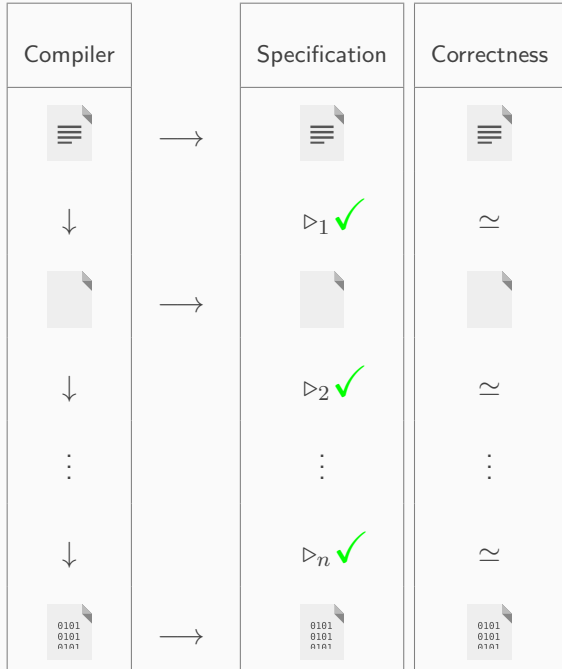


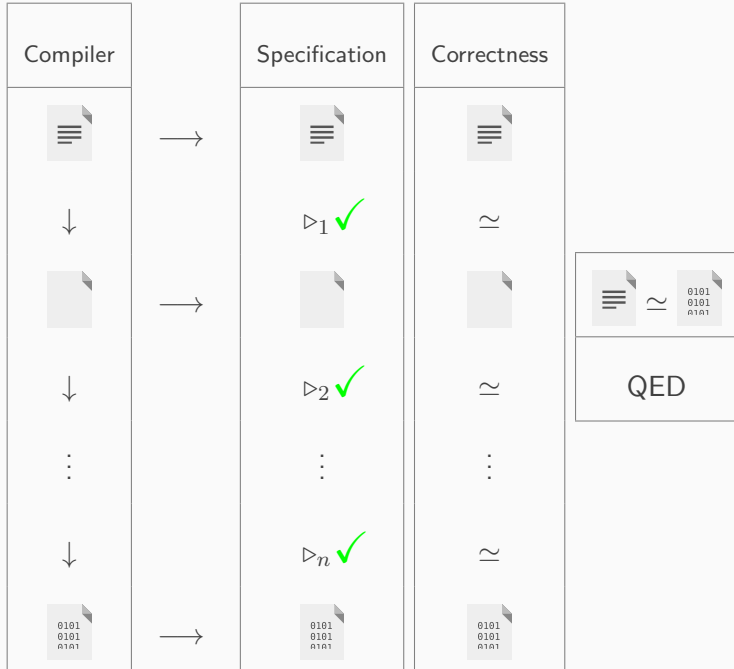
Compiler






















- Trusted computing base



- Trusted computing base
 - Definition of \simeq



- Trusted computing base
 - Definition of \simeq
 - Coq Kernel



- Trusted computing base
 - Definition of \approx
 - Coq Kernel
- Coq kernel: relatively small and well-studied code base

Additionally

- Formal pedantic mode

Additionally

- Formal pedantic mode
 - Integrate decision procedures directly in the compiler

Additionally

- Formal pedantic mode
 - Integrate decision procedures directly in the compiler
- Improved compiler testing

Additionally

- Formal pedantic mode
 - Integrate decision procedures directly in the compiler
- Improved compiler testing
 - Plutus: existing test suite (example input/outputs)

Additionally

- Formal pedantic mode
 - Integrate decision procedures directly in the compiler
- Improved compiler testing
 - Plutus: existing test suite (example input/outputs)
 - Property-based testing

Additionally

- Formal pedantic mode
 - Integrate decision procedures directly in the compiler
- Improved compiler testing
 - Plutus: existing test suite (example input/outputs)
 - Property-based testing
- Generate LaTeX from specification (Coq code)

- End-to-end verification
 - Plutus: Haskell \rightarrow PIR \rightarrow PLC
 - Alternative front-end: Gallina + CertiCoq
- Proof transportation?
 - Use equivalence proof
- Certifier implementation: Coq vs Agda?
- Hints for constructing typing derivation, instead of re-implementing type-checker

Conclusion

- Gradual form of verification: each layer improves the trustworthiness of the system.

Conclusion

- Gradual form of verification: each layer improves the trustworthiness of the system.
- We don't have to reason about an implementation, only a high-level specification

Conclusion

- Gradual form of verification: each layer improves the trustworthiness of the system.
- We don't have to reason about an implementation, only a high-level specification
- Robustness: certain changes in the compiler do not affect the verification at all.

Conclusion

- Gradual form of verification: each layer improves the trustworthiness of the system.
- We don't have to reason about an implementation, only a high-level specification
- Robustness: certain changes in the compiler do not affect the verification at all.
- Separation of concerns: compiler writers do not have to be proof engineers