

Translation Certification for Smart Contracts

Anonymous

No Institute Given

Abstract. Compiler correctness is an old problem, but with the emergence of *smart contracts* on blockchains that problem presents itself in a new light. Smart contracts are self-contained pieces of software that control (valuable) assets in an adversarial environment; once committed to the blockchain, these smart contracts cannot be modified. Smart contracts are typically developed in a high-level contract language and compiled to low-level virtual machine code before being committed to the blockchain. For a smart contract user to trust a given piece of low-level code on the blockchain, they must convince themselves that (a) they are in possession of the matching source code and (b) that the compiler has correctly translated the source code to the given low-level code.

Classic approaches to compiler correctness tackle the second point. We argue that *translation certification* also squarely addresses the first. We describe the proof architecture of a novel translation certification framework, implemented in Coq, for a functional smart contract language. We demonstrate that we can model the compilation pipeline as a sequence of translation relations that facilitate a modular verification methodology and are robust in the face of an evolving compiler implementation.

1 Introduction

Compiler correctness is an old problem that has received renewed interest in the context of *smart contracts* — that is, compiled code on public blockchains, such as Ethereum or Cardano. This code often controls a significant amount of financial assets, must operate under adversarial conditions, and can no longer be updated once it has been committed to the blockchain. Bugs in smart contracts are a significant problem in practice [5]. Recent work has also established that smart contract language compilers can exacerbate this problem [23, Section 3] (in this case, the Vyper compiler). More specifically, the authors report (a) that they did find bugs in the Vyper compiler that compromised smart contract security and (b) that they performed verification on generated low-level code, because they were wary of compiler bugs.

Hence, to support reasoning about smart contract source code, we need to get a handle on the correctness of smart contract compilers. On top of that, we do also need a *verifiable link* between the source code and its compiled code to prevent *code substitution attacks*, where an adversary presents the user with source code that doesn't match the low-level code committed on-chain.

In this paper, we are reporting on our ongoing effort to develop a certification engine for the open-source on-chain code compiler of the Plutus smart contract

system¹ for the Cardano blockchain.² Specifically, we make the following contributions:

- We describe a novel architecture for a translation certifier based on *translation relations*, which enables us to generate *translation certificates*—proof objects that relate the source code to the resulting compiled code and establish the correctness of the translation (Section 2).
- We provide formal definitions for the transformation passes that step-by-step translate PIR (Plutus Intermediate Representation) to PLC (Plutus Core) and briefly discuss the challenges associated with the certification of each of these passes (Section 3).
- We present a taxonomy of existing approaches to compiler correctness and discuss the importance of generating translation certificates in the domain of smart contracts (Section 4).

We also evaluate how our approach to incremental certification copes with changes to the compiler, which is being developed in an independent open source project. Finally, we discuss related work in Section 5 and future work in Section 6.

2 The Architecture of the Certifier

On-chain code in the Plutus smart contract system is written in a subset of Haskell called *Plutus Tx*. The Plutus Tx compiler is implemented as a plugin for the widely-used, industrial-strength GHC Haskell compiler, combining large parts of the GHC’s compilation pipeline with custom translation steps to generate Plutus Core. In this context, it seems infeasible to apply full-scale compiler verification à la CompCert [18]. We will therefore outline the design of a certification engine that, using the Coq proof assistant, generates a proof object, a *translation certificate*, asserting the validity of a Plutus Core program with respect to a given Plutus Tx source contract. In addition to asserting the correct translation of *this one program*, the translation certificate serves as a verifiable link between source and generated code.

We model the compiler as a composition of pure functions that transform one abstract syntax tree into another. Figure 1 illustrates the architecture for a single transformation, where the gray area marks the compiler implementation as a function $f_i : \text{AST}_i \rightarrow \text{AST}_{i+1}$. We use a family of types AST_i to illustrate that the representation of the abstract syntax might change after each transformation.

To support certification, the compiler outputs each intermediate tree t_i , so that we can parse these in our Coq implementation of the certifier. Within Coq, we define a high-level specification of each pass. We call this specification a *translation relation*: a binary relation on abstract syntax trees that specifies the intended behaviour of the compiler pass. The orange area in Figure 1 displays

¹ <https://developers.cardano.org/docs/smart-contracts/plutus/>

² <http://cardano.org> is, at the time of writing, the 5th largest public blockchain by market capitalisation.

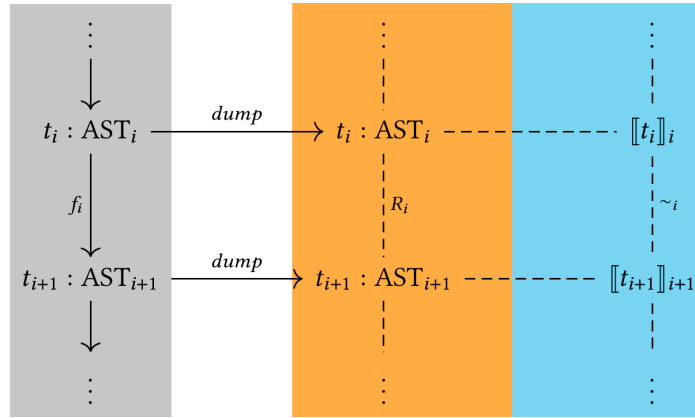


Fig. 1. Architecture for a single compiler pass. The grey area represents the compiler, orange and blue represent the certification component in Coq.

the translation relation R_i of pass i , where the vertical dashed line indicates that $R_i(t_i, t_{i+1})$ holds. To establish this, we define a search procedure that, given two subsequent trees produced by the compiler, can construct a derivation relating the two.

The translation relation is purely syntactic—it does not assert anything about the correctness of the compiler—but rather *specifies* the behaviour of a particular compiler pass. To verify that the compilation preserves language semantics requires an additional proof, the blue area in Figure 1, that establishes that any two terms related by R_i have the same semantics.

We have implemented this approach for a range of concrete passes of the Plutus Tx compiler. To illustrate our approach in this section, we will use an untyped lambda calculus, extended with non-recursive let-bindings.

$$t ::= x \mid \lambda x. t \mid t t \mid \text{let } x = t \text{ in } t$$

In the following section, we will extend this to a lambda calculus that is closer to the intermediate language used by the Plutus Tx compiler.

2.1 Characterising a transformation

To assert the correctness of a single compiler stage f_i , we begin by defining a translation relation R_i on a pair of source and target terms t_i and t_{i+1} , respectively. This relation characterises the admissible translations of that compiler stage. That is, for all t_i, t_{i+1} , we have $f_i(t_i) = t_{i+1}$ implies $R_i(t_i, t_{i+1})$.

As a concrete example, consider an inlining pass. We have characterised this as an inductively defined relation in Figure 2. Here, $\Gamma \vdash s \triangleright t$ asserts that program s can be translated into t given an environment Γ of let-bound variables, paired with their definition. According to Rule [Inline-Var] the variable x may

$$\begin{array}{c}
\frac{\Gamma(x) = t' \quad \Gamma \vdash t' \triangleright t}{\Gamma \vdash x \triangleright t} \quad [\text{Inline-Var}] \\
\frac{\Gamma \vdash t_1 \triangleright t'_1 \quad (x, t_1), \Gamma \vdash t_2 \triangleright t'_2}{\Gamma \vdash \mathbf{let} \ x = t_1 \ \mathbf{in} \ t_2 \triangleright \mathbf{let} \ x = t'_1 \ \mathbf{in} \ t'_2} \quad [\text{Inline-Let}] \\
\frac{\Gamma \vdash t_1 \triangleright t'_1 \quad \Gamma \vdash t_2 \triangleright t'_2}{\Gamma \vdash t_1 \ t_2 \triangleright t'_1 \ t'_2} \quad [\text{Cong-App}] \\
\frac{}{\Gamma \vdash x \triangleright x} \quad [\text{Cong-Var}] \\
\frac{\Gamma \vdash t_1 \triangleright t'_1}{\Gamma \vdash \lambda x. t_1 \triangleright \lambda x. t'_1} \quad [\text{Cong-Abs}]
\end{array}$$

Fig. 2. Characterisation of an inliner

be replaced by t when the pair (x, t') can be looked up in Γ and t' can be translated to t , accounting for repeated inlining. The remaining rules are congruence rules, where Rule [Inline-Let] also extends the environment Γ . We omitted details about handling variable capture to keep the presentation simple: hence, we assume that variable names are globally unique.

Crucially, these rules do *not* prescribe which variable occurrences should be inlined, since the [Inline-Var] and [Cong-Var] rules overlap. The choice in the implementation of the pass may rely on a complex set of heuristics internal to the compiler. Instead, we merely define a relation capturing the *possible* ways in which the compiler *may* behave. This allows for a certification engine that is robust with respect to changes in the compiler, such as the particular heuristics used to decide when to replace a variable with its definition or not.

We can now encode the relation as an indexed algebraic datatype in Coq, as shown in Figure 3. Here, we use a de Bruijn representation for variables. The environment is extended at both `Inl_Lam` and `Inl_Let`, shifting any previous references in the old environment.

2.2 Proof search

After defining a translation relation R_i characterising one compiler stage, we now define a search procedure to construct a proof that for two particular terms t_i and t_{i+1} , produced by a run of the compiler, the relation $R_i(t_i, t_{i+1})$ holds. To find and implement such a search procedure, we generally follow these steps:

1. We write proofs for specific compilations by hand using Coq's *tactics*. For simple relations, like the inline example sketched above, a proof can often be found with a handful of tactics such as `auto` or `constructor`. This is particularly useful for debugging the design of our relations describing compiler passes. The drawback of this approach is, however, that it is difficult to reason when such proof search may fail. Furthermore, proofs written using such tactics quickly become slow for large terms.

```

Inductive Inline : Env -> Term -> Term -> Type :=
| Inl_Lam : forall {env s t},
  Inline (LambdaBound :: shiftEnv env) s t ->
  Inline env (Lam s) (Lam t)

| Inl_App : forall {env s sx t tx},
  Inline env s t -> Inline env sx tx ->
  Inline env (App s sx) (App t tx)

| Inl_Let : forall {env s t s' t'},
  Inline env s s' ->
  Inline (LetBound s :: shiftEnv env) t t' ->
  Inline env (Let s t) (Let s' t')

| Inl_Var : forall {env n},
  Inline env (Var n) (Var n)

| Inl_Inl : forall {env n t},
  nth_error env n = Some (LetBound t) ->
  Inline env (Var n) t

```

Fig. 3. Characterisation of an inliner in Coq

2. Once we are sufficiently confident that a relation accurately captures admissible compiler behaviour, we write a decision procedure of the form `forall (t1 t2 : Term), option (R t1 t2)`. These procedures can still produce large proof terms and may not always successfully construct a proof, but they form a useful intermediate step towards full-on proof by reflection.
3. Finally, we write a boolean decision procedure in the style of `ssreflect` [15] of type `Term -> Term -> Bool`, together with a soundness proof stating that it will only return `true` when two terms are related through R_i . Verifying such boolean functions for complex compilation passes is non-trivial; hence, we only invest the effort once we have a reasonable degree of confidence that the relation we have defined accurately describes a given compiler pass.

2.3 Semantics preservation

Given the relational specification of each individual compiler pass, we can now establish the correctness properties for each pass. In the simplest case, this could be asserting the preservation of a program's static semantics, i.e., a proof of type preservation. On the other end of the spectrum, we can demonstrate that the translated term is semantically equivalent to the original program. Proving such properties for Plutus, however, requires advanced techniques such as step-indexed logical relations [2], which go beyond the scope of the current paper.

In Figure 1, we denote R_i 's correctness properties in the blue area by means of an abstract binary relation \sim_i on the semantic objects $[[t_i]]_i$ of ASTs t_i . In the

case of static semantics, we can choose typing derivations as semantic objects, and (for most passes) relate these by simply comparing types syntactically.

We can construct these proofs independently and incrementally for each step in the translation. In fact, even without any formal proof about the semantics, manual inspection of a translation relation may already provide some degree of confidence that the translation step is correct. After all, the translation relation asserts the specification of this compiler pass’ admissible behaviour.

2.4 Certificate generation

A complete translation certificate includes the entire set of ASTs t_1, \dots, t_n together with the proof objects witnessing the translation relations R_i for these ASTs. In addition, it includes the instantiated proofs of the refinement relation to produce a single proof object.

This certificate together with the source and compiled code can be independently checked by a trusted proof checker, such as the Coq kernel [8]. The proof itself can be inspected to confirm it proves the right theorem. One can then be confident that the compiled program is a faithful translation of the source code.

3 Translation Relations of the Plutus Tx Compiler

The Plutus Tx compiler translates a subset of Haskell to Plutus Core, a variant of System F_{ω}^{μ} [12]. The Plutus Core code is committed to the Cardano blockchain, constituting the definitive reference to any deployed smart contract.

Plutus Core programs are pure, self-contained functions (i.e., they do not link to other code) and are passed a representation of the transaction whose validation they contribute to. The programs are run by an interpreter during the transaction validation phase of the blockchain.

The Plutus Tx compiler itself is implemented as a core-to-core pass plugin [13] in the GHC compiler pipeline. On a high level, the compiler comprises three steps:

1. The parsing, type-checking and desugaring phase of GHC are reused to obtain a GHC Core program.
2. An large subset of GHC Core is transformed into an intermediate language named Plutus Intermediate Representation (PIR). All referred definitions are included so that the program is self-contained.
3. The PIR program is then transformed and compiled down into Plutus Core

The certification effort reported here focuses on Step 3, which consists of several optimisation passes and translation steps. PIR is a superset of the Plutus Core language: it adds several conveniences, such as user-defined datatypes, strict and non-strict let-bindings that may be (mutually) recursive. The compilation steps translate these constructs into simpler language constructs.

In Figure 4 we present a simplified version of the PIR syntax, where we omit some constructs for the sake of presentation. The full PIR language specification

$t ::= x \mid \lambda(x : \tau). t \mid t t$	variable, lambda, function application
$\mid \Lambda(\alpha : \kappa). t \mid t \{\tau\}$	type abstraction, type application
$\mid \text{let}_r^s x = t \text{ in } t$	term bindings
$\mid \text{data } T \bar{\alpha} = \overline{C}_i \overline{\tau}_i \text{ with } x \text{ in } t$	datatype binding
$r ::= \text{rec} \mid \text{nonrec}$	recursion type of binding
$s ::= \text{strict} \mid \text{nonstrict}$	strictness of binding
$\tau ::= \dots$	types

Fig. 4. Simplified PIR

has been formalised elsewhere [12, 16]. In particular, we ignore the fact that in PIR, let-bindings may contain a group of (mutually recursive) bindings. Similarly, we do not include mutually-recursive datatypes. Furthermore, we omit the syntax of types, and the term-level witnesses of iso-recursive types. We occasionally omit type annotations, when they are not relevant.

We introduce the individual compiler passes that the Plutus compiler performs using the following Haskell program to illustrate their behaviour:

```
-- | Either a specific end date, or "never".
data EndDate = Fixed Integer | Never

pastEnd :: EndDate -> Integer -> Bool
pastEnd end current =
  let inlineMe = False
  in case end of
    Fixed n -> (let floatMe = if current `greaterThanEqInteger` 0
                       then n else 0 in floatMe) `lessThanEqInteger` current
    Never -> inlineMe
```

This program is a basic implementation of a *timelock*, a contract that states that funds may be moved after a certain date, or not at all. It contains a few contrived bindings (`inlineMe` and `floatMe`) that will be useful to illustrate some transformations. After the program is desugared to GHC Core, it is converted to a term in PIR that corresponds to the following Simplified PIR term:

```
data Bool = True | False with Bool_match in
data Unit = Unit with Unit_match in
let nonrec strict lessThanEqInteger = ... in
  data EndDate = Fixed Integer | Never with EndDate_match in
    \ (end : EndDate) .
    \ (current : Integer) .
      let nonrec nonstrict inlineMe = False in
```

```

EndDate_match end
(\unit n -> lessThanEqInteger
 (let nonrect nonstrict floatMe =
   Bool_match (greaterThanEqInteger current 0)
   (\unit -> n) (\unit -> 0)
   Unit
   in floatMe)
 current)
(\unit -> inlineMe)
Unit

```

Note that case distinction of a type T is encoded as the application of a pattern match function T_match , which is introduced as part of a data definition. Furthermore, branches of a case distinction are delayed by abstracting over a unit value, since PIR is a strict language.

Next we will discuss the compiler passes, we have included each intermediate form of the above program with some commentary in Appendix A.

3.1 Variable Renaming

The renaming pass transforms a program into an α -equivalent program, with globally unique names. The correctness of some subsequent transformations depends on this property. We can express this pass as a translation relation $\Delta \vdash t \triangleright_{\alpha} t'$, stating that under the renaming environment Δ (consisting of pairs of variables), t is renamed to t' . The environment Δ records all variables that are in scope in t , paired with their corresponding name in t' .

The case for lambda abstractions is defined as follows:

$$\frac{(x, y), \Delta \vdash t \triangleright_{\alpha} t' \quad \{z \mid (z, y) \in \Delta\} \cap FV(t) = \emptyset}{\Delta \vdash \lambda x.t \triangleright_{\alpha} \lambda y.t'} \quad [\text{Rename-Abs}]$$

The [Rename-Abs] rule states that a lambda-bound variable may be renamed at its binding-site, when t and t' are related under the extended environment. Additionally, the new name y should not capture any free variable that is also renamed to y . Very similar rules can be stated for other binding constructs such as `let`. Note that the variables x and y may be equal, in which case no renaming was performed.

The variable case simply follows from the environment Δ :

$$\frac{(x, y) \in \Delta}{\Delta \vdash x \triangleright_{\alpha} y} \quad [\text{Rename-Var}]$$

3.2 Inlining

The rules of the translation relation for inlining in PIR are similar to those in Section 2.1. However, the Plutus Tx compiler does more than just inlining let-bound definitions. It also performs dead-code elimination (removing those let-bindings

that have been inlined exhaustively) and it renames variables to ensure the global uniqueness of bound variables, also known as the *Barendregt-convention*. This introduces a problem for our certification approach, as we cannot observe and dump the intermediate ASTs, since the transformations are fused into a single pass in the compiler.

We solve this by modeling the individual transformations, composing them using *relational composition*, $\exists t_2. R_1(t_1, t_2) \wedge R_2(t_2, t_3)$. To construct a proof relating two terms, then amounts to also finding the *intermediate term*, t_2 witnessing the composite transformation. To simplify the search of this intermediate AST, we adjust the compiler to emit supporting information about the performed pass; in this case, a list of the eliminated variables. If the compiler emits incorrect information, we may fail to construct a certificate, but we will never produce an incorrect certificate.

3.3 Let-floating

During let-floating, let-bindings can be moved upwards in the program. This may save unnecessarily repeated computation and makes the generated code more readable. The Plutus compiler constructs a dependency graph to maintain a correct ordering when multiple definitions are floated. For the translation relation, we first consider the interaction of a `let` expression with its parent node in the AST. For example, consider the case of a lambda with a non-strict `let` directly under it:

$$\frac{x \notin FV(t_1) \quad x \neq y \quad t_1 \triangleright_{let} t'_1 \quad t_2 \triangleright_{let} t'_2}{\lambda x. \text{let}_r^{\text{nonstrict}} y = t_1 \text{ in } t_2} \text{ [Float-Let-Lam]}$$

$$\text{let}_r^{\text{nonstrict}} y = t'_1 \text{ in } \lambda x. t'_2$$

This rule states that a non-strict let-binding may float up past a lambda, if the bound term does not reference the lambda-bound variable. Furthermore, we require $x \neq y$, to avoid variable capture in t_2 . This rule does not apply to `strict` let-bindings, as floating them outside a lambda might change termination behaviour of the program. Similar rules then express when a let may float upwards past the other language constructs. Most of these are much simpler, only binding constructs pose additional constraints on scoping and strictness. Since the compiler pass may float `lets` more than just one step up, we define the translation relation as the transitive closure of \triangleright_{let} . Note that we do not need to maintain a dependency graph in the certifier, but only need to assert that transformations do not break dependencies.

3.4 Dead-code elimination

By means of a live variable analysis, the compiler determines which let-bound definitions are unused. This is mainly useful for definitions that are introduced by other compiler passes. Since PIR is a strict language, however, the compiler

can only eliminate those bindings for which it can determine they have no side-effects. For example, a let-bound expression that is unused but diverges cannot be removed, as that could change the termination behaviour of the program.

The analysis in the compiler is not as straightforward as counting occurrences. Even a let-bound variable that does occur in the code, may be dead-code, if it is only used in other dead bindings. This is also known as strongly live variable analysis [14]. We define a translation relation $t \triangleright_{dce} t'$ that captures dead code elimination. The crucial rule is for let-bindings.

$$\frac{t_2 \triangleright_{dce} t'_2 \quad x \notin FV(t'_2)}{\mathbf{let}_{\substack{\text{nonstrict} \\ r}} x = t_1 \mathbf{in} t_2 \triangleright_{dce} t'_2} \quad [\text{DCE-Let-nonstrict}]$$

Note that the condition $x \notin FV(t'_2)$ mentions the *resulting* body of the let t'_2 . This is justified since the rules of \triangleright_{dce} can remove bindings only, but cannot change any other language constructs. This illustrates how succinct we can describe the specification of a complex compiler pass.

In practice, the Plutus compiler also eliminates some strict bindings that obviously do not diverge, such as values.

3.5 Encoding of non-strict bindings

The PIR language allows both for strict and non-strict let-bindings, but Plutus Core does not. The *thunking transformation* is used to obtain semantic equivalent definitions which use a strict let-binding. We define the rules as a relation $\Gamma \vdash t \triangleright_{\circlearrowleft} t'$, where Γ records for every bound variable whether it was bound strictly or non-strictly. The rule for a non-strict binding site is:

$$\frac{\Gamma \vdash t_1 \triangleright_{\circlearrowleft} t'_1 \quad (x, \mathbf{nonstrict}), \Gamma \vdash t_2 \triangleright_{\circlearrowleft} t'_2 \quad y \notin FV(t_1)}{\mathbf{let}_{\substack{\text{nonstrict} \\ \text{nonrec}}} x = t_1 \mathbf{in} t_2 \triangleright_{\circlearrowleft}} \quad [\text{Thunk-Let-nonstrict}]$$

$$\mathbf{let}_{\substack{\text{strict} \\ \text{nonrec}}} x = \lambda \mathit{unit}. t'_1 \mathbf{in} t'_2$$

This rule states that a right hand side is thunked by introducing a lambda abstraction that expects a trivial unit value as its argument.

The rules for other variable binders extend Γ . The rule for a recursive let-binding also extends the environment under which t_1 is transformed. Finally, we also replace the occurrences of nonstrict variables, adding an application to the unit value, thereby forcing evaluation.

$$\frac{(x, \mathbf{nonstrict}) \in \Gamma}{\Delta \vdash x \triangleright_{\circlearrowleft} x \ ()} \quad [\text{Thunk-Var}]$$

3.6 Encoding of recursive bindings

The Plutus compiler translates (mutually) recursive let-bindings in non-recursive ones using fixpoint combinators. Here we only consider the rule for individual recursive lets in simplified PIR:

$$\frac{t_1 \triangleright_\mu t'_1 \quad t_2 \triangleright_\mu t'_2 \quad y \notin FV(t_1)}{\text{let}_{\text{rec}}^s x = t_1 \text{ in } t_2} \text{ [EncRec-Let]}$$

$$\text{let}_{\text{nonrec}}^{\text{strict}} fix = \dots \text{ in let}_{\text{nonrec}}^s x = fix (\lambda x. t'_1) \text{ in } t'_2$$

This rule relates recursive bindings to non-recursive ones, and expects an explicit definition of the fixpoint operator as well. Since PIR has no primitive construct for term-level fix-points, the compiler generates a definition fix . Note that fix is defined in a non-recursive let, its construction relies on recursive types [16].

The actual transformation for PIR is much more involved, since mutually recursive binding groups require a more involved fixpoint combinator of which the definition depends on the size of the group.

3.7 Encoding of datatypes

Datatype definitions are encoded using lambda and type abstractions according to the Scott encoding [1]. To show the idea of the rather general \triangleright_{data} translation relation, we show a rule specialised to the *Maybe* datatype.

$$\frac{t \triangleright_{data} t'}{\text{data } Maybe \alpha = Just \alpha \mid Nothing \text{ with } maybe \text{ in } t} \text{ [Scott-Maybe]}$$

$$(\wedge Maybe. \lambda Just. \lambda Nothing. \lambda maybe. t') \tau_{Maybe} t_{Just} t_{Nothing} t_{maybe}$$

The [Scott-Maybe] rule relates the datatype definition to a term that abstracts over the type *Maybe*, its constructors *Just* and *Nothing* and the matching function *maybe*, which are each lambda encoded. For the exact definitions of τ_{Maybe} , t_{Just} , $t_{Nothing}$ and t_{maybe} we refer to the general formalisation of PIR [16].

3.8 Encoding of non-recursive bindings

A non-recursive let-binding is simply compiled into a β redex:

$$\frac{t_1 \triangleright_\beta t'_1 \quad t_2 \triangleright_\beta t'_2}{\text{let}_{\text{nonrec}}^{\text{strict}} x = t_1 \text{ in } t_2 \triangleright_\beta (\lambda x. t'_2) t'_1} \text{ [Redex-Let]}$$

Note that at this point in the compiler pipeline, $\text{let}_{\text{nonrec}}^{\text{strict}}$ is the only type of let-binding that can still occur.

4 Evaluation

In this section, we evaluate our approach to proof engineering for an independently developed, constantly evolving compiler under the application constraints imposed by smart contracts.

4.1 Compilers and correctness

The standard approach to compiler correctness is *full compiler verification*: a proof that asserts that the compiler is correct as it demonstrates that, for any valid source program, the translation produces a semantically equivalent target program. Examples of this approach include the CompCert [18] and CakeML [17] projects, showing that (with significant effort) it is possible to verify a compiler end-to-end. To do so, the compiler is typically implemented in a language suitable for verification, such as the Coq proof assistant or the HOL theorem prover.

In contrast, the technique that we propose for the Plutus compiler is based on *translation validation* [24]. Instead of asserting an entire compiler correct, translation validation establishes the correctness of individual compiler runs.

A statement of full compiler correctness is, of course, the stronger of the two statements. Translation validation may fail to assert the correctness of some compiler runs; either because the compiler did not produce correct code or because the translation certifier is incomplete. In exchange for being the weaker property, translation validation is potentially (1) less costly to realise, (2) easier to retrofit to an existing compiler, and (3) more robust in the face of changes to the compiler.

The idea of *proof-carrying code* [20] is closely related to translation validation, shifting the focus to compiled programs, rather than the compiler itself. A program is distributed together with a proof of a property such as memory or type safety. Such a proof excludes certain classes of bugs and gives direct evidence to the users of such a program, who may independently check the proof before running it. Our certification effort, while related, differs in that we keep proof and program separate and in that we are interested in full semantic correctness and not just certain properties like memory and type safety.

4.2 Certificates and smart contracts

Smart contracts often manage significant amounts of financial and other assets. Before a user engages with such a contract, which has been committed to the blockchain as compiled code, they may want to inspect the source code to assert that it behaves as they expect. In order to be able to rely on that inspection, they need to know without doubt that (1) they are looking at the correct source code and (2) that the source code has been compiled correctly.

While a verified smart contract compiler addresses the second point, it doesn't help with the first without an infrastructure of reproducibility. In contrast, a *certifying compiler* [21] that generates an independently verifiable certificate of correct translation, squarely addresses both points. By verifying a smart contract's translation certificate, a smart contract user can convince themselves that they are in possession of the matching source code and that this was correctly compiled to the code committed to the blockchain.

4.3 Engineering considerations

Incremental verification. The certifier architecture outlined in this paper allows for an incremental approach to verification: during the development of the certification engine, each individual step in the process increases our overall confidence in the compiler’s correctness, even if we have not yet completed the end-to-end verification of the compiler pipeline.

By defining only the translation relations, we have an independent formal specification of the compiler’s behaviour. This makes it easier to reason informally and to spot potential mistakes or problems with the implementation.

Implementing the decision procedures for translation relations ties the implementation to the specification: we can show on a per-compilation basis that a pass is sound with respect to its specification as a translation relation. Furthermore, we can test and debug translation relations by automatically constructing evidence for various input programs.

Finally, by proving semantics preservation of a translation relation, we gain full confidence in the corresponding pass for compiler runs that abide by that translation relation.

Agility. The Plutus Tx compiler is developed independently of our certification effort. Moreover, it depends on large parts of a large code base — namely, that of the Glasgow Haskell Compiler (GHC). In addition, both GHC and the Plutus Tx-specific parts evolve on a constant basis; for example, to improve code optimisation or to fix bugs.

In that context, full verification appears an insurmountable task and a proof on the basis of the compiler source code would constantly have to adapt to the evolving compiler source. Hence, the architecture of our certification engine is based on a *grey box approach*, where the certifier matches the general outline (such as the phases of the compiler pipeline), but not all of the implementation details of the compiler. For example, our translation relation for the inliner admits any valid inlining. Improvements of the compiler heuristics to produce more efficient programs by being selective about what precisely to inline don’t affect the inliner’s translation relation, and hence, don’t affect the certifier.

Trusted Computing Base (TCB). The fact that the Plutus Tx compiler is not implemented in a proof assistant, but in Haskell complicates direct compiler verification. It might be possible to use a tool like `hs-to-coq` [26], which translates a subset of Haskell into Coq’s Gallina and has been used for proving various properties about Haskell code [10]. However, given that those tools often only cover language subsets, it is not clear that they are applicable. More importantly, such an approach would increase the size of the trusted computing base (TCB), as the translation from Haskell into Coq’s Gallina is not verified. Similarly, extraction-based approaches suffer from the same problem if the extraction itself is not verified, although there are projects like `CertiCoq` [3] that try to address that issue.

In any case, our architecture has a small TCB. We directly relate the source and target programs, taking the compiler out of the equation. Trusting a translation certificate comes down to trusting the Coq kernel that checks the proof, the theorem with its supporting definitions and soundness of the interpreter with respect to the formalised semantics. Of course, these components are part of the TCB of a verified compiler too. This aspect also motivated our choice of Coq over other languages such as Agda, due to its relatively small and mature kernel.

5 Related Work

Ethereum was the first blockchain to popularise use of smart contracts, written in the Solidity programming language. Solidity is an imperative programming language that is compiled to EVM bytecode, which runs on a stack machine operating on persistent mutable state. The DAO vulnerability [11] has underlined the importance of formal verification of smart contracts. Notably, a verification framework has been presented [9] for reasoning about embedded Solidity programs in F^* . The work includes a decompiler to convert EVM bytecode, generated by a compiler, into Solidity programs in F^* . The authors propose that correctness of compilation can be shown by proving equivalence of the embedded source and (decompiled) target program using relational reasoning [6]. However, this would involve a manual proof effort on a per-program basis, and relies on the F^* semantics since the embeddings are shallow. Furthermore, components such as the decompiler are not formally verified, adding to the size of the TCB.

The translation validation technique has been used for the verification of a particular critical Ethereum smart contract [23] using the K framework. The work demonstrates how translation validation can successfully be applied to construct proofs about the low-level EVM bytecode by mostly reasoning on the (much more understandable) source code. The actual refinement proof is still constructed manually, however.

The Tezos blockchain also uses a stack-like language, called Michelson. The Mi-Cho-Coq framework [7] formalises the language and supports reasoning with a weakest precondition logic. There is ongoing work for developing a certified compiler in Coq for the Albert intermediate language, intended as a target language for certified compilers of higher-level languages. This differs from our approach as it requires the compiler to be implemented in the proof assistant.

ConCert is a smart contract verification framework in Coq [4]. It enables formal reasoning about the source code of a smart contracts, defined in a different (functional) language. The programs are translated and shallowly embedded in Coq’s Gallina. Interestingly, the translation is proven sound, in contrast with approaches such as *hs-to-coq* [26], since it is implemented using Coq’s metaprogramming and reasoning facility *MetaCoq* [25].

The Cogent certifying compiler [22] has shown that it is possible to use translation validation for lowering the cost of functional verification of low-level code: a program can be written and reasoned about in a high-level functional language, which is compiled down to C. The generated certificate then proves a

refinement relation, capable of transporting the verification results to the corresponding C code. The situation is different from ours: the Cogent compiler goes through a range of languages with different semantic models and uses the forward-simulation technique as a consequence. In contrast, we are working with variations of lambda calculi that have similar semantics, allowing us to use logical relations and translation relations.

In their Coq framework [19], Li and Appel use a similar technique for specifying compiler passes as inductive relations in Coq. Their tool reduces the effort of implementing program transformations and corresponding correctness proofs. The tool is able to generate large parts of an implementation together with a partial soundness proof with respect to those relations. The approach is used to implement parts of the CertiCoq backend.

6 Further work

The Plutus Tx compiler translates a Haskell subset into Plutus Core. The compiler consists of two main parts: the first one reuses various stages of GHC to compile the Haskell subset to GHC Core — GHC’s principal intermediate language. The second part compiles GHC Core to Plutus Core. As Plutus Core is strict and doesn’t directly support datatypes, both parts are quite complex. Moreover, both consist of a significant number of successive transformation steps.

In this paper, we focused on the certification effort covering the second part of that pipeline; specifically, the translation from PIR to Plutus Core. We developed translation relations for all passes described in Section 3, such that we can, for example, relate previously described timelock example in PIR to its final form in Plutus Core. For some of these passes, such as inlining, we have implemented a verified decision procedure, but most of the evidence is generated semi-automatically by using Coq tactics. We have not yet covered all transformations in their full generality; for example, we do not cover (mutually) recursive datatypes yet. We have also started the semantic verification of key passes of the translation and are investigating different ways to improve the efficiency of proof search for larger programs.

Our next steps comprise the following: (1) filling in the remaining gaps in translation relations (such as covering mutually recursive datatypes); (2) complete all decision procedures; (3) drive the semantic verification forward; and (4) develop techniques to further automate our approach and improve the efficiency of the certifier.

The first three steps pose a significant amount of work, but we do not expect major new conceptual questions or obstacles. This is different for Step (4), where we anticipate the need for further research work. This includes more compositional definitions of the translation relations, such that we can generate at least part of the decision procedures (semi-)automatically. Moreover, we already perceive efficiency to be a bottleneck and we plan to work on optimising the proof search. Finally, we plan to apply our approach to the first part of the Plutus Tx compiler (Haskell subset to GHC Core).

References

1. Abadi, M., Cardelli, L., Plotkin, G.: Types for the scott numerals (1993)
2. Ahmed, A.: Step-indexed syntactic logical relations for recursive and quantified types. In: European Symposium on Programming. pp. 69–83. Springer (2006)
3. Anand, A., Appel, A., Morrisett, G., Paraskevopoulou, Z., Pollack, R., Belanger, O.S., Sozeau, M., Weaver, M.: Certicoq: A verified compiler for coq. In: The third international workshop on Coq for programming languages (CoqPL) (2017)
4. Annenkov, D., Nielsen, J.B., Spitters, B.: Concert: a smart contract certification framework in coq. In: Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs. pp. 215–228 (2020)
5. Atzei, N., Bartoletti, M., Cimoli, T.: A survey of attacks on Ethereum smart contracts (SoK). In: Principles of Security and Trust (POST 2017). LNCS, vol. 10204 (2017)
6. Barthe, G., Fournet, C., Grégoire, B., Strub, P.Y., Swamy, N., Zanella-Béguelin, S.: Probabilistic relational verification for cryptographic implementations. ACM SIGPLAN Notices **49**(1), 193–205 (2014)
7. Bernardo, B., Cauderlier, R., Hu, Z., Pesin, B., Tesson, J.: Mi-cho-coq, a framework for certifying tezos smart contracts. In: International Symposium on Formal Methods. pp. 368–379. Springer (2019)
8. Bertot, Y., Castéran, P.: Interactive theorem proving and program development: Coq’Art: the calculus of inductive constructions. Springer Science & Business Media (2013)
9. Bhargavan, K., Delignat-Lavaud, A., Fournet, C., Gollamudi, A., Gonthier, G., Kobeissi, N., Kulatova, N., Rastogi, A., Sibut-Pinote, T., Swamy, N., et al.: Formal verification of smart contracts: Short paper. In: Proceedings of the 2016 ACM workshop on programming languages and analysis for security. pp. 91–96 (2016)
10. Breitner, J., Spector-Zabusky, A., Li, Y., Rizkallah, C., Wiegley, J., Weirich, S.: Ready, set, verify! applying hs-to-coq to real-world haskell code (experience report). Proceedings of the ACM on Programming Languages **2**(ICFP), 1–16 (2018)
11. Buterin, V.: CRITICAL UPDATE Re: DAO Vulnerability. <https://blog.ethereum.org/2016/06/17/critical-update-re-dao-vulnerability/> (2016), retrieved December 10, 2021
12. Chapman, J., Kireev, R., Nester, C., Wadler, P.: System F in Agda, for Fun and Profit. In: Mathematics of Program Construction (MPC 2019). LNCS, vol. 11825 (2019)
13. GHC Team: GHC 9.0 User Manual. https://downloads.haskell.org/~ghc/9.0.1/docs/html/users_guide/extending_ghc.html
14. Giegerich, R., Möncke, U.: Invariance of approximative semantics with respect to program transformations. In: GI—11. Jahrestagung, pp. 1–10. Springer (1981)
15. Gonthier, G., Le, R.S.: An Ssreflect Tutorial. Ph.D. thesis, INRIA (2009)
16. Jones, M.P., Gkoumas, V., Kireev, R., MacKenzie, K., Nester, C., Wadler, P.: Unraveling recursion: compiling an ir with recursion to system f. In: International Conference on Mathematics of Program Construction. pp. 414–443. Springer (2019)
17. Kumar, R., Myreen, M.O., Norrish, M., Owens, S.: Cakeml: a verified implementation of ml. ACM SIGPLAN Notices **49**(1), 179–191 (2014)
18. Leroy, X., Blazy, S., Kästner, D., Schommer, B., Pister, M., Ferdinand, C.: Compcert—a formally verified optimizing compiler. In: ERTS 2016: Embedded Real Time Software and Systems, 8th European Congress (2016)

19. Li, J.M., Appel, A.W.: Deriving efficient program transformations from rewrite rules. *Proceedings of the ACM on Programming Languages* **5**(ICFP), 1–29 (2021)
20. Necula, G.C.: Proof-carrying code. In: *Proceedings of the 24th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. pp. 106–119 (1997)
21. Necula, G.C., Lee, P.: The design and implementation of a certifying compiler. *SIGPLAN Not.* **39**(4), 612–625 (Apr 2004)
22. O’Connor, L., Rizkallah, C., Chen, Z., Amani, S., Lim, J., Nagashima, Y., Sewell, T., Hixon, A., Keller, G., Murray, T., et al.: Cogent: certified compilation for a functional systems language. *arXiv preprint arXiv:1601.05520* (2016)
23. Park, D., Zhang, Y., Rosu, G.: End-to-end formal verification of Ethereum 2.0 deposit smart contract. In: *Computer Aided Verification (CAV 2020)*. LNCS, vol. 12224 (2020)
24. Pnueli, A., Siegel, M., Singerman, E.: Translation validation. In: *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. pp. 151–166. Springer (1998)
25. Sozeau, M., Anand, A., Boulier, S., Cohen, C., Forster, Y., Kunze, F., Malecha, G., Tabareau, N., Winterhalter, T.: The metacoq project. *Journal of Automated Reasoning* (2020)
26. Spector-Zabusky, A., Breitner, J., Rizkallah, C., Weirich, S.: Total haskell is reasonable coq. In: *Proceedings of the 7th ACM SIGPLAN International Conference on Certified Programs and Proofs*. pp. 14–27 (2018)

A Compiler dumps for the timelock program

In this appendix we show step-by-step how the timelock example in section 3 is transformed by the passes in the Plutus compiler. These programs were obtained by running the Plutus compiler on the Haskell source code program, after modifying the Plutus pretty-printer to output a bit more compact presentation. We occasionally omit some sub-terms to improve readability (indicated as ...).

A.1 Original PIR Term

The Plutus compiler converts the GHC Core program into the following PIR program. Note that variables in PIR are represented as pairs of names and unique integers. The name is only maintained for readability, whereas the integers are used for actual program transformations. We pretty-print the integer in subscript after the name.

The conversion includes definitions for all the built-in types and functions that may be used in PIR program, since the program has to be self-contained. Starting from line 34 we can recognise the timelock example. Note that Haskell’s lazy `case` expression has been translated to a call to `EndDate_match`, where the case branches have been “thunked” by abstracting over a unit value. This thunking prevents the (strict) function application of `EndDate_match` from evaluating all the branches.

```

1  let nonrec type ByteString0 = ... in
2  let nonrec data Bool11 = True13 : ... | False14 : ... with Bool_match12 in
3  let nonrec strict verifySignature57 = ... in
4  let nonrec type String2 = ... in
5  let nonrec data Unit60 = Unit62 : ... with Unit_match61 in
6  let nonrec strict trace70 = ... in
7  let nonrec type Integer1 = ... in
8  let nonrec strict takeByteString5 = ... in
9  let nonrec strict subtractInteger27 = ... in
10 let nonrec strict sha38 = ... in
11 let nonrec strict sha27 = ... in
12 let nonrec strict remainderInteger32 = ... in
13 let nonrec strict quotientInteger31 = ... in
14 let nonrec strict multiplyInteger28 = ... in
15 let nonrec strict modInteger30 = ... in
16 let nonrec strict lessThanInteger44 = ... in
17 let nonrec strict lessThanEqInteger48 = ... in
18 let nonrec strict lessThanByteString20 = ... in
19 let nonrec strict greaterThanInteger36 = ... in
20 let nonrec strict greaterThanEqInteger40 = ... in
21 let nonrec strict greaterThanByteString24 = ... in
22 let nonrec strict error64 = ... in
23 let nonrec strict equalsInteger52 = ... in
24 let nonrec strict equalsByteString16 = ... in
25 let nonrec strict emptyString66 = ... in

```

```

26 let nonrec strict emptyByteString25 = ... in
27 let nonrec strict dropByteString6 = ... in
28 let nonrec strict divideInteger29 = ... in
29 let nonrec strict concatenate4 = ... in
30 let nonrec type Char3 = ... in
31 let nonrec strict charToString67 = ... in
32 let nonrec strict appendString65 = ... in
33 let nonrec strict addInteger26 = ... in
34 let nonrec data EndDate71 = Fixed73 : ... | Never74 : ... with EndDate_match72 in
35 λds75 : EndDate71 .
36 λds76 : Integer .
37 let nonrec nonstrict inlineMe77 = False in
38 let nonrec nonstrict wild78 = ... in
39 (((EndDate_match72 ds75 { Unit60 -> Bool11 }) (λn79 : Integer .
40 λthunk84 : Unit60 .
41 (lessThanEqInteger48 (let nonrec nonstrict floatMe83 = ... in
42 floatMe83)) ds76)) (λthunk85 : Unit60 .
43 inlineMe77)) Unit62

```

A.2 Renaming

The first pass does a global renaming to ensure each variable is in fact globally unique. Note that the integers of all bound variables are indeed renamed compared to the previous version of the program.

```

1 let nonrec type ByteString86 = ... in
2 let nonrec data Bool87 = True88 : ... | False89 : ... with Bool_match90 in
3 let nonrec strict verifySignature91 = ... in
4 let nonrec type String96 = ... in
5 let nonrec data Unit97 = Unit98 : ... with Unit_match99 in
6 let nonrec strict trace100 = ... in
7 let nonrec type Integer103 = ... in
8 let nonrec strict takeByteString104 = ... in
9 let nonrec strict subtractInteger105 = ... in
10 let nonrec strict sha3106 = ... in
11 let nonrec strict sha2107 = ... in
12 let nonrec strict remainderInteger108 = ... in
13 let nonrec strict quotientInteger109 = ... in
14 let nonrec strict multiplyInteger110 = ... in
15 let nonrec strict modInteger111 = ... in
16 let nonrec strict lessThanInteger112 = ... in
17 let nonrec strict lessThanEqInteger116 = ... in
18 let nonrec strict lessThanByteString120 = ... in
19 let nonrec strict greaterThanInteger124 = ... in
20 let nonrec strict greaterThanEqInteger128 = ... in
21 let nonrec strict greaterThanByteString132 = ... in

```

```

22 let nonrec strict error136 = ... in
23 let nonrec strict equalsInteger140 = ... in
24 let nonrec strict equalsByteString144 = ... in
25 let nonrec strict emptyString148 = ... in
26 let nonrec strict emptyByteString149 = ... in
27 let nonrec strict dropByteString150 = ... in
28 let nonrec strict divideInteger151 = ... in
29 let nonrec strict concatenate152 = ... in
30 let nonrec type Char153 = ... in
31 let nonrec strict charToString154 = ... in
32 let nonrec strict appendString155 = ... in
33 let nonrec strict addInteger156 = ... in
34 let nonrec data EndDate157 = Fixed158 : ... | Never159 : ... with EndDate_match160 in
35 λds161 : EndDate157 .
36 λds162 : Integer .
37 let nonrec nonstrict inlineMe163 = False in
38 let nonrec nonstrict wild164 = ... in
39 (((EndDate_match160 ds161 { Unit97 -> Bool87 }) (λn165 : Integer .
40 λthunk166 : Unit97 .
41 (lessThanEqInteger116 (let nonrec nonstrict floatMe167 = ... in
42 floatMe167)) ds162)) (λthunk171 : Unit97 .
43 inlineMe163)) Unit98

```

A.3 Dead Code Elimination

In this pass, the compiler cleans up the unused definitions that were present after the GHC core translation.

```

1 let nonrec data Bool1 = True2 : ... | False3 : ... with Bool_match4 in
2 let nonrec data Unit11 = Unit12 : ... with Unit_match13 in
3 let nonrec strict lessThanEqInteger30 = ... in
4 let nonrec strict greaterThanEqInteger42 = ... in
5 let nonrec data EndDate71 = Fixed72 : ... | Never73 : ... with EndDate_match74 in
6 λds75 : EndDate71 .
7 λds76 : Integer .
8 let nonrec nonstrict inlineMe77 = False in
9 (((EndDate_match74 ds75 { Unit11 -> Bool1 }) (λn79 : Integer .
10 λthunk80 : Unit11 .
11 (lessThanEqInteger30 (let nonrec nonstrict floatMe81 = ... in
12 floatMe81)) ds76)) (λthunk85 : Unit11 .
13 inlineMe77)) Unit12

```

A.4 Inlining

The compiler performs an inlining pass and decides to inline the let-bound definition `inlineMe` on line 8 in Section A.3. This results in the following program, where the let-binding has been eliminated and the inlined definition (`False`) can be seen on line 12.

```

1 let nonrec data Bool1 = True2 : ... | False3 : ... with Bool_match4 in
2 let nonrec data Unit11 = Unit12 : ... with Unit_match13 in
3 let nonrec strict lessThanEqInteger30 = ... in
4 let nonrec strict greaterThanEqInteger42 = ... in
5 let nonrec data EndDate71 = Fixed72 : ... | Never73 : ... with EndDate_match74 in
6 λds75 : EndDate71 .
7 λds76 : Integer .
8 (((EndDate_match74 ds75 { Unit11 -> Bool1 }) (λn79 : Integer .
9 λthunk80 : Unit11 .
10 (lessThanEqInteger30 (let nonrec nonstrict floatMe81 = ... in
11 floatMe81)) ds76)) (λthunk85 : Unit11 .
12 False3)) Unit12

```

A.5 Thinking recursive definitions

The next pass thinks recursive term bindings (similar to the encoding of non-strict let bindings in Section A.7), to make sure that they are of a function type and work well with the fixpoint combinator that is introduced in a later pass (Section A.9). Since this program does not include any recursive term bindings, the result is unchanged.

A.6 Let-floating

Next, the Plutus compiler decides to float a let-bound definition. In this run, the `floatMe` definition is moved outside of the first argument of `lessThanEqInteger`, as can be seen on line 10. Additionally, this pass performs *merging* of adjacent let definitions into a single let with a group of bindings, printed on line 1. We did not mention this transformation in Section 3.3, since simplified PIR has no binding groups. The order of these definitions has also changed, but this is fine as long as no dependencies are broken. We use a translation relation that is reminiscent of the one in `Sectionsub:let-float`, but for bindings only.

```

1 let nonrec data Bool1 = True2 : ... | False3 : ... with Bool_match4;
2   strict greaterThanEqInteger42 = ...;
3   data Unit11 = Unit12 : ... with Unit_match13;
4   data EndDate71 = Fixed72 : ... | Never73 : ... with EndDate_match74;
5   strict lessThanEqInteger30 = ... in
6 λds75 : EndDate71 .
7 λds76 : Integer .
8 (((EndDate_match74 ds75 { Unit11 -> Bool1 }) (λn79 : Integer .
9 λthunk80 : Unit11 .
10 let nonrec nonstrict floatMe81 = ... in
11 (lessThanEqInteger30 floatMe81) ds76)) (λthunk85 : Unit11 .
12 False3)) Unit12

```

A.7 Encoding of non-strict let bindings

The non-strict binding on line 10 is transformed in a strict binding by thinking. From the type we can see that the Plutus compiler actually abstracts over the Scott-encoded version of a unit value. The occurrence is applied to a unit value on line 11.

```

1 let nonrec data Bool1 = True2 : ... | False3 : ... with Bool_match4;
2   strict greaterThanEqInteger42 = ...;
3   data Unit11 = Unit12 : ... with Unit_match13;
4   data EndDate71 = Fixed72 : ... | Never73 : ... with EndDate_match74;
5   strict lessThanEqInteger30 = ... in
6   λds75 : EndDate71 .
7   λds76 : Integer .
8   (((EndDate_match74 ds75 { Unit11 -> Bool1 }) (λn79 : Integer .
9   λthunk80 : Unit11 .
10  let nonrec strict floatMe81 = λarg207 : ∀a0 : *.a0 -> a0 . ... in
11  (lessThanEqInteger30 (floatMe81 (λa0 : *.
12  λx1 : a0 .
13  x1))) ds76)) (λthunk85 : Unit11 .
14  False3)) Unit12

```

A.8 Encoding of datatypes

Next, the three datatype definitions are Scott encoded. For example, the `Bool` datatype with its constructors and matching function are bound on line 1 to 4, and the corresponding definitions are found as arguments on line 40 for `Bool`, line 41-44 for `True`, line 45-48 for `False` and line 49-50 for `Bool_match`.

```

1 (((λABool1 : *.
2   λTrue2 : Bool1 .
3   λFalse3 : Bool1 .
4   λBool_match4 : Bool1 -> (∀a221 : *.a221 -> (a221 -> a221)) .
5   let nonrec strict greaterThanEqInteger42 = ... in
6   ((λAUnit11 : *.
7     λUnit12 : Unit11 .
8     λUnit_match13 : Unit11 -> (∀a217 : *.a217 -> a217) .
9     (((λAEndDate71 : *.
10      λFixed72 : Integer -> EndDate71 .
11      λNever73 : EndDate71 .
12      λEndDate_match74 : EndDate71 -> (∀a208 : *. (Integer -> a208) -> (a208 -> a208)) .
13      let nonrec strict lessThanEqInteger30 = ... in
14      λds75 : EndDate71 .
15      λds76 : Integer .
16      (((EndDate_match74 ds75 { Unit11 -> Bool1 }) (λn79 : Integer .
17      λthunk80 : Unit11 .
18      let nonrec strict floatMe81 = ... in

```

```

19 (lessThanEqInteger30 (floatMe81 ( $\Lambda$ a0 : *.
20  $\lambda$ x1 : a0 .
21 x1))) ds76)) ( $\lambda$ thunk85 : Unit11 .
22 False3)) Unit12 {  $\forall$ a208 : *. (Integer  $\rightarrow$  a208)  $\rightarrow$  (a208  $\rightarrow$  a208) }
23 ( $\lambda$ arg_0212 : Integer .
24    $\Lambda$ a209 : *.
25    $\lambda$ case_Fixed210 : Integer  $\rightarrow$  a209 .
26    $\lambda$ case_Never211 : a209 .
27   case_Fixed210 arg_0212))
28 ( $\Lambda$ a213 : *.
29    $\lambda$ case_Fixed214 : Integer  $\rightarrow$  a213 .
30    $\lambda$ case_Never215 : a213 .
31   case_Never215))
32 ( $\lambda$ x216 :  $\forall$ a208 : *. (Integer  $\rightarrow$  a208)  $\rightarrow$  (a208  $\rightarrow$  a208) .
33   x216)
34 {  $\forall$ a217 : *. a217  $\rightarrow$  a217 }
35 ( $\Lambda$ a218 : *.
36    $\lambda$ case_Unit219 : a218 .
37   case_Unit219))
38 ( $\lambda$ x220 :  $\forall$ a217 : *. a217  $\rightarrow$  a217 .
39   x220)
40 {  $\forall$ a221 : *. a221  $\rightarrow$  (a221  $\rightarrow$  a221) }
41 ( $\Lambda$ a222 : *.
42    $\lambda$ case_True223 : a222 .
43    $\lambda$ case_False224 : a222 .
44   case_True223))
45 ( $\Lambda$ a225 : *.
46    $\lambda$ case_True226 : a225 .
47    $\lambda$ case_False227 : a225 .
48   case_False227))
49 ( $\lambda$ x228 :  $\forall$ a221 : *. a221  $\rightarrow$  (a221  $\rightarrow$  a221) .
50   x228)

```

A.9 Recursive let bindings, inlining and dead-code elimination

The next three passes encode recursive term-bindings, and perform another round of inlining and dead code elimination. In this example program however, they have no effect and the program does not change.

A.10 Non-recursive let bindings

The final pass encodes non-recursive let bindings as a beta-redex. The `floatMe` binding in Section A.8 line 18 can be recognised below on line 18, where it is now lambda-bound, and line 43 where the definition is provided as an argument.

```

1  (((λBool0 : *.
2  λTrue1 : Bool0 .
3  λFalse2 : Bool0 .
4  λBool_match3 : Bool0 -> (∀a4 : *.a4 -> (a4 -> a4)) .
5  (λgreaterThanEqInteger5 : Integer -> (Integer -> Bool0)) .
6  ((λUnit9 : *.
7  λUnit10 : Unit9 .
8  λUnit_match11 : Unit9 -> (∀a12 : *.a12 -> a12)) .
9  (((λEndDate13 : *.
10 λFixed14 : Integer -> EndDate13 .
11 λNever15 : EndDate13 .
12 λEndDate_match16 : EndDate13 -> (∀a17 : *. (Integer -> a17) -> (a17 -> a17)) .
13 (λlessThanEqInteger18 : Integer -> (Integer -> Bool0)) .
14 λds22 : EndDate13 .
15 λds23 : Integer .
16 (((λEndDate_match16 ds22 { Unit9 -> Bool0 }) (λn24 : Integer .
17 λthunk25 : Unit9 .
18 (λfloatMe26 : (∀a27 : *.a27 -> a27) -> Integer .
19 (lessThanEqInteger18 (floatMe26 (λa32 : *.
20 λx33 : a32 .
21 x33))) ds23) (λarg28 : ∀a29 : *.a29 -> a29 .
22 (((λBool_match3 ((greaterThanEqInteger5 ds23) 0) { Unit9 -> Integer }) (λthunk30 : Unit9 .
23 n24)) (λthunk31 : Unit9 .
24 0)) Unit10))) (λthunk34 : Unit9 .
25 False2)) Unit10) (λarg19 : Integer .
26 λarg20 : Integer .
27 (λb21 : Bool .
28 ((ifThenElse { Bool0 } b21) True1) False2) ((lessThanEqInteger arg19) arg20)) {...})
29 (λarg036 : Integer .
30 λa37 : *.
31 λcase_Fixed38 : Integer -> a37 .
32 λcase_Never39 : a37 .
33 case_Fixed38 arg036)) (λa40 : *.
34 λcase_Fixed41 : Integer -> a40 .
35 λcase_Never42 : a40 .
36 case_Never42)) (λx43 : ∀a44 : *. (Integer -> a44) -> (a44 -> a44)) .
37 x43) { ∀a45 : *.a45 -> a45 } (λa46 : *.
38 λcase_Unit47 : a46 .
39 case_Unit47)) (λx48 : ∀a49 : *.a49 -> a49 .
40 x48)) (λarg6 : Integer .
41 λarg7 : Integer .
42 (λb8 : Bool .
43 ((ifThenElse { Bool0 } b8) True1) False2) ((greaterThanEqInteger arg6) arg7)) {...})
44 (λa51 : *.
45 λcase_True52 : a51 .

```



```

46  $\lambda \text{case\_False}_{53} : a_{51} .$ 
47  $\text{case\_True}_{52})) (\lambda a_{54} : * .$ 
48  $\lambda \text{case\_True}_{55} : a_{54} .$ 
49  $\lambda \text{case\_False}_{56} : a_{54} .$ 
50  $\text{case\_False}_{56})) (\lambda x_{57} : \forall a_{58} : * . a_{58} \rightarrow (a_{58} \rightarrow a_{58}) .$ 
51  $x_{57})$ 

```