

1 Bridging Specification and Implementation in 2 Smart Contract Languages

3 Cas van der Rest  

4 IOHK, The Netherlands <http://www.casvanderrest.nl/>

5 — Abstract —

6 Before entrusting a smart contract with our funds or data, it is essential to fully understand its terms.
7 A formal specification of the smart contract language’s semantics is crucial for expressing these terms
8 unambiguously. However, before a specification can serve as a reliable tool for assessing a contract’s
9 behavior, we must first establish a strong connection between the specification and implementation, a
10 task that is further complicated by the rapid evolution of smart contract languages and the common
11 prioritization of implementation over formal specification during (early) development.

12 In this paper, we propose a strategy to prevent divergence between implementation and specific-
13 ation by transpiling Nanopass intermediate representations (IRs) into mutually defined families of
14 inductive data types in Agda. We enforce completeness of typing relations using Agda’s dependent
15 type system and meta-programming capabilities. A key outcome of our approach is that any meaning-
16 ful syntactic change—such as the addition or removal of production rules or non-terminals—results
17 in a compile-time error in the specification. Although our approach was developed in the context of
18 Compact—the smart contract language of the Midnight blockchain—we believe that it may serve
19 as a general template for synchronization between Agda-based specifications and smart contract
20 language implementations.

21 **2012 ACM Subject Classification** [Replace ccsdesc macro with valid one](#)

22 **Keywords and phrases** Dummy keyword

23 **Digital Object Identifier** 10.4230/OASlcs.CVIT.2016.23

24 **Acknowledgements** I want to thank ...

25 **1** Introduction

26 Smart contracts potentially handle large amounts of funds, and as such it is crucial precisely
27 understand its terms, as well as trust that a system will faithfully execute these terms.
28 Failing to understand the terms of a smart contract may lead to substantial financial losses,
29 as exemplified by the infamous exploit of the TheDAO contract [8] in 2016. Subtleties around
30 the expression of control flow in Solidity smart contracts allowed hackers to launch a reentry
31 attack and capture the equivalent of \$50 in Ether, resulting in hard fork of the Ethereum
32 blockchain to recover the stolen funds.


33 This exploit perfectly illustrates the need to be able to unambiguously express the terms
34 of a smart contract. A formal mathematical *specification* of the smart contract language vital
35 for expressing a contract’s intended semantics. However, the utility of formal specifications
36 in assessing a contract’s validity hinges on whether it faithfully reflects how smart contracts
37 are executed by the *implementation* of a system. This problem of how to bridge the gap
38 between specification and implementation is an age-old question in the field of programming
39 languages, resulting in many different approaches to tackle the problem, ranging from fully-
40 verified compilers [4, 5] to more light-weight approaches such as certifying compilation [6] or
41 conformance testing.

42 While previous work offers us a plethora of techniques for connecting specification and
43 implementation, additional challenges arise when they are applied outside the context of
44 academic research. Typically, smart contract languages are developed in an environment



© Cas van der Rest;
licensed under Creative Commons License CC-BY 4.0
42nd Conference on Very Important Topics (CVIT 2016).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:7



OpenAccess Series in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

45 that is characterized by rapid evolution of the language’s design as well as prioritization of
 46 implementation over the development of (formal) specifications. This presents a dilemma: do
 47 we develop the specification *with* the language, embracing the additional maintenance costs
 48 incurred by an evolving design, or do we wait for the language’s design and implementation
 49 to stabilize. Although waiting appears tempting, it is important to highlight that there
 50 are downsides too. That is, we lose potential some synergy between specification and
 51 implementation, where by forcing ourselves to express mathematically what (we think) we
 52 are doing, we may uncover flaws in our thinking and design. This allows a specification to
 53 inform the design of the very language it is specifying. It becomes much harder to backtrack
 54 on mistakes when strictly sequencing the development of a language’s implementation and
 55 specification.

56 Clearly, the ideal scenario would be to develop a specification in conjunction with an
 57 implementation, while maintaining a strong connection between the two. However, if we
 58 require implementation and specification to be connected through formal proofs, e.g., by
 59 having a verified compiler, this becomes entirely infeasible for almost all projects due to the
 60 tremendous amount of resources required. This begs the question of whether we can settle
 61 for a more light-weight approach where we develop the implementation and specification
 62 simultaneously but separately, connecting them in a way that prevents divergence in the
 63 presence of design changes but does not impede the development process by requiring proofs
 64 to be written before code reaches production.

65 In this paper, we present such a light-weight strategy for bridging the specification and
 66 implementation of smart contract languages. In short, our approach works by transpiling
 67 the syntax definition of Nanopass [3] intermediate representations (IRs) to mutually-defined
 68 families of inductive data types in Agda [7], leveraging its dependent type system and meta-
 69 programming capabilities to enforce completeness of typing rules defined over the transpiled
 70 syntax. This approach has two important benefits. First, it guards against divergence of the
 71 specification and implementation, by using the compiler’s internal syntax definition as the
 72 source of truth for the languages abstract syntax. This ensures that any meaningful syntactic
 73 change to the languages abstract syntax—in the form of adding, removing or changing
 74 production rules or non-terminals—manifests as a compile time error when type checking
 75 the specification. As a result, we can fully automatically check whether the specification and
 76 implementation are synchronized. The second important benefit of our approach is that it
 77 does not impact the compiler development cycle in any way. By repurposing the definition
 78 of Nanopass IRs as a DSL for specifying the language’s abstract syntax, we integrate the
 79 specification and implementation in a way that neither induces any additional overhead in
 80 the compiler nor does it induce additional maintenance work for compiler developers.

81 More concretely, this paper is structured as follows:

- 82 ■ In Section 2, we illustrate, by example, how one would define and formally specify a
 83 smart contract language using our approach.
- 84 ■ In Section 3, we present an approach for leveraging Agda’s dependent type system and
 85 meta-programming capabilities to statically enforce completeness of typing relations
 86 defined over the transpiled syntax.

87 Finally, we conclude and discuss future work in Section 4.

88 1.1 Industrial Application

89 The techniques in this paper have been developed and applied in the context *Compact* [1],
 90 the smart contract language of the Midnight blockchain [2]. In that context, we successfully

```

(define-language Lstlc      -- Declares the name of the language
  (terminals              -- Declares the terminals of the language
    (string (name))
    (number (nat)))
  (Type (type)           -- Declares "Type" and its meta-variable
    (tbool)
    (tnat)
    (tfun type1 type2))
  (Expr (expr)          -- Declares "Expr" and its meta-variable
    (evar name)
    (etrue)
    (efalse)
    (elit nat)
    (eif expr1 expr2 expr3)
    (elam name expr1)
    (eapp expr1 expr2)
    (eadd expr1 expr2)))

```

■ **Figure 1** Nanopass IR definition for a Simply-Typed λ -calculus

91 applied the techniques described in this paper to develop an Agda specification of Compact’s
 92 static semantics next to its Nanopass implementation. While Compact is not a particularly
 93 large language, its abstract syntax still contains in the order of 100 production rules spanning
 94 several tens of non-terminals. By extension, the formal specification of its type system
 95 contains an equal number of typing rules respectively judgments. Using the techniques
 96 described in this paper, we managed to keep this formal specification from diverging from
 97 its implementation, despite the rapid pace at which the language is currently evolving. We
 98 choose to present them here in a more general setting to illustrate their potential application
 99 outside the context in which they were initially developed.

100 2 Integrating Specification and Implementation, by Example

101 In this section we demonstrate the process for defining, transpiling, and specifying a language
 102 definition. For this purpose, we show, as a three-step process, how to formally specify the
 103 static semantics of a simply-typed λ -calculus.

104 2.1 Step 1: Declare Abstract Syntax of the Target Language as a 105 Nanopass IR

106 We start by declaring the abstract syntax of our language as a Nanopass IR. This definition of
 107 the abstract syntax is regarded as the source of truth, and is used both by the implementation
 108 and specification. Figure 1 shows the definition of a Nanopass IR for a simply typed λ -
 109 calculus. Along with some metadata defining the name of the IR (`Lstlc`), it declares two
 110 non-terminals: `Type`, defining the syntax of types, and `Expr`, defining the abstract syntax
 111 of terms. For each non-terminal, we must also declare a *meta-variable* (respectively `type`
 112 and `expr` for types and terms of the language), which are used to refer to their associated
 113 non-terminal. Meta-variables also play a crucial role in transpilation; when converting the
 114 untyped Nanopass IR to a typed Agda definition, the required type information is recovered
 115 by resolving these meta-variables.

```

mutual
  data Expr : Set where
    evar : String → Expr
    etrue : Expr
    efalse : Expr
    enum : ℕ → Expr
    eif : Expr → Expr → Expr → Expr
    elam : String → Expr → Expr
    eapp : Expr → Expr → Expr
    eadd : Expr → Expr → Expr

  variable
    expr expr1 expr2 expr3 expr' : Expr

  data Type : Set where
    tbool : Type
    tnat : Type
    tfun : Type → Type → Type

  variable
    type type1 type2 type3 type' : Type

```

■ **Figure 2** Transpiled Agda definition of the Nanopass IR definition shown in Figure 1.

116 2.2 Step 2: Transpilation of Abstract Syntax to Agda

117 The next step is to transpile the Nanopass IR defined in Figure 1 to Agda. This generates
 118 a mutually-defined family of inductive data types. Figure 2 shows the generated Agda
 119 definition. Each non-terminal in the abstract syntax corresponds to a data type, and each
 120 production rule to a constructor. To make sure that the generated type signatures are accepted
 121 by Agda, we must resolve meta-variables referring to both terminals and non-terminals. For
 122 example, we transpile the production rule for λ -abstraction as follows:

123 $(\text{elam name expr1}) \mapsto \text{elam} : \text{String} \rightarrow \text{Expr} \rightarrow \text{Expr}$

124 Here, to ensure that the generated type signature is accepted by Agda, we must resolve the
 125 meta-variable `expr1` to the `Expr` data type, and the meta-variable `name` to the `String` type.

126 In addition to a mutually-defined family of inductive types, transpilation also generates
 127 several *generalized variables* for each data type, indicated by `variable`. These serve a similar
 128 purpose to meta-variables in the Nanopass IR, in that they are used to refer to universally-
 129 quantified values of their corresponding types.

130 2.3 Step 3: Define Typing Rules as an Inductive Relation over Abstract 131 Syntax

132 Now that the abstract syntax of the target language is available in Agda, we can define the
 133 type system. Typically, one does this in Agda by declaring an inductive relation over terms
 134 (or, term-indexed data type), whose constructors the typing rules, i.e., the different ways in

$$\begin{array}{l}
 \text{⊢-evar} : \text{ name } \mapsto \text{ type } \in \Gamma \\
 \hline
 \Gamma \vdash \text{evar name} : \text{ type} \\
 \\
 \text{⊢-eif} : \Gamma \vdash \text{expr}_1 : \text{tbool} \\
 \rightarrow \Gamma \vdash \text{expr}_2 : \text{type} \\
 \rightarrow \Gamma \vdash \text{expr}_3 : \text{type} \\
 \hline
 \Gamma \vdash \text{eif expr}_1 \text{ expr}_2 \text{ expr}_3 : \text{type} \\
 \\
 \text{⊢-eapp} : \Gamma \vdash \text{expr}_1 : \text{tfun type}_1 \text{ type}_2 \\
 \rightarrow \Gamma \vdash \text{expr}_2 : \text{type}_1 \\
 \hline
 \Gamma \vdash \text{eapp expr}_1 \text{ expr}_2 : \text{type}_2
 \end{array}$$

■ **Figure 3** Excerpt of the typing rules for the simply typed λ -calculus defined in Figures 1 and 2, defined as constructors of an inductive relation over context, term, and type.

135 which a proof of well-typedness can be constructed. Generally, this relation has additional
 136 positions for tracking type information and contextual information. In this case, we use a
 137 three-place relation over context, terms, and types:

```
data ⊢_ : _ (Γ : Context) : Expr → Type → Set where
```

138 Figure 3 shows an excerpt of the definition of the static semantics for the target lan-
 139 guage, by declaring constructors corresponding to the typing rules for variables, if-then-else
 140 expressions, and function application.

141 3 Ensuring Completeness of Typing Relations

142 Right now, a compile-time error will be triggered in the following cases:

- 143 ■ a non-terminal is removed,
- 144 ■ a production rule is removed, or
- 145 ■ a production rule is changed.

146 In summary, by defining our formal specification on top of the transpiled syntax, we are
 147 guaranteed that our specification can only refer to syntactic elements that are part of the
 148 compiler's internal definition. What is explicitly not guaranteed is that the specification is a
 149 *complete* one. That is, the following changes to the language will not trigger a compile-time
 150 error:

- 151 ■ a non-terminal is added, or
- 152 ■ a production rule is added.

153 More likely than not, such syntactic changes signify the addition of a new feature to
 154 the language, for which we should also extend the formal specification. To ensure that
 155 the addition of new syntax to the language triggers a compile-time error, we must perform

156 additional checks on the Agda side. Specifically, we must (1) declare which data types make
 157 up the language, (2) check that we have a corresponding typing relation for each data type
 158 in the language's syntax, and (3) check that for every constructor there is a corresponding
 159 typing rule in the associated typing relation.

160 For (1), the transpilation tool generates an additional definition that collects the non-
 161 terminals of the language together with their meta-variables:

```
Lstlc : List (Set × String)
Lstlc = (Expr , "expr") :: (Type , "type") :: []
```

162 To ensure that each non-terminal has an associated typing relation, we declare an instance
 163 of the `HasTyping` for the `Lstlc` language. This instance contains a proof witnessing that every
 164 non-terminal of the language has an associated typing relation in the form of a three-place
 165 relation over context, terms, and types.

```
Typing S = ∃₂ λ (CTX : Set) (I : CTX → Set) → (ctx : CTX) → S → I ctx → Set
record HasTyping (syn : List (Set × String)) : Set₁ where
  field rels : All (Typing ∘ proj₁) syn
```

166 As a result, by defining an instance of the form `HasTyping Lstlc`, we are forced to declare a
 167 typing relation for each non-terminal. Furthermore, whenever a new non-terminal is added in
 168 the Nanopass IR, after transpilation of the updated syntax, this instance becomes ill-typed,
 169 and we are must declare a new typing relation for the newly added non-terminal.

170 Finally, we use Agda's meta-programming capabilities to check that the declared typing
 171 relations cover every production rule of the syntax. In short, this check is performed by
 172 asserting that for every constructor in the untyped syntax, there is a constructor of the
 173 corresponding typing relation that has that constructor in the term position (i.e., S in the
 174 definition of `Typing`). We invoke this check for `Lstlc` as follows:

```
175 unquoteDecl = checkRels (getTyping Lstlc) []
```

176 If the typing rules shown in Figure 2 were the only rules we defined, invoking the meta-
 177 program above would result in the following type error when checking the specification,
 178 indicating which relations are missing which typing rules.

```
179 Discovered missing rule(s) while checking coverage of relation _⊢_: _
180 ---> No typing rule found for constructor etrue
181 ---> No typing rule found for constructor efalse
182 ---> No typing rule found for constructor enum
183 ---> No typing rule found for constructor elam
184 ---> No typing rule found for constructor eadd
```

185 4 Conclusion and Future Work

186 In this paper, we presented an approach for connecting the specification and implementation
 187 of smart contract languages through the transpilation of Nanopass IRs. While our approach
 188 was developed and applied in the context of the Compact language, its principles and
 189 methodology are generalizable to other (smart contract) languages.

190 For future work, we aim to further enhance the accessibility and applicability of our
 191 approach by making our tools publicly available as open-source projects, and leveraging our
 192 extraction mechanism to enable further compiler verification. For example, we could explore
 193 certification [6] techniques to formally verify compiler correctness.

194 — **References** —

- 195 **1** URL: <https://docs.midnight.network/develop/reference/compact/>.
- 196 **2** URL: <https://midnight.network/>.
- 197 **3** Andrew W. Keep and R. Kent Dybvig. A nanopass framework for commercial compiler
198 development. In Greg Morrisett and Tarmo Uustalu, editors, *ACM SIGPLAN International*
199 *Conference on Functional Programming, ICFP'13, Boston, MA, USA - September 25 - 27,*
200 *2013*, pages 343–350. ACM, 2013. doi:10.1145/2500365.2500618.
- 201 **4** Ramana Kumar, Magnus O. Myreen, Michael Norrish, and Scott Owens. Cakeml: a verified
202 implementation of ML. In Suresh Jagannathan and Peter Sewell, editors, *The 41st Annual*
203 *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14,*
204 *San Diego, CA, USA, January 20-21, 2014*, pages 179–192. ACM, 2014. doi:10.1145/2535838.
205 2535841.
- 206 **5** Xavier Leroy, Sandrine Blazy, Daniel Kästner, Bernhard Schommer, Markus Pister, and
207 Christian Ferdinand. Compcert-a formally verified optimizing compiler. In *ERTS 2016:*
208 *Embedded Real Time Software and Systems, 8th European Congress*, 2016.
- 209 **6** George C. Necula and Peter Lee. The design and implementation of a certifying compiler.
210 In Jack W. Davidson, Keith D. Cooper, and A. Michael Berman, editors, *Proceedings of*
211 *the ACM SIGPLAN '98 Conference on Programming Language Design and Implementation*
212 *(PLDI), Montreal, Canada, June 17-19, 1998*, pages 333–344. ACM, 1998. doi:10.1145/
213 277650.277752.
- 214 **7** Ulf Norell. Dependently typed programming in agda. In Pieter W. M. Koopman,
215 Rinus Plasmeijer, and S. Doaitse Swierstra, editors, *Advanced Functional Programming,*
216 *6th International School, AFP 2008, Heijen, The Netherlands, May 2008, Revised Lec-*
217 *tures*, volume 5832 of *Lecture Notes in Computer Science*, pages 230–266. Springer, 2008.
218 doi:10.1007/978-3-642-04652-0_5.
- 219 **8** Cryptopedia Staff. The dao: What was the dao hack?, Oct 2023. URL: [https://www.gemini.](https://www.gemini.com/cryptopedia/the-dao-hack-makerdao)
220 [com/cryptopedia/the-dao-hack-makerdao](https://www.gemini.com/cryptopedia/the-dao-hack-makerdao).