

Compiling a Higher-Order Smart Contract Language to LLVM

Vaivaswatha Nagaraj
Zilliqa Research
vaivaswatha@zilliqa.com

Jacob Johannsen
Zilliqa Research
jacob@zilliqa.com

Anton Trunov
Zilliqa Research
anton@zilliqa.com

George Pirlea
Zilliqa Research
george@zilliqa.com

Amrit Kumar
Zilliqa Research
amrit@zilliqa.com

Ilya Sergey
Yale-NUS College
National University of Singapore
ilya.sergey@yale-nus.edu.sg

Abstract

SCILLA is a higher-order polymorphic typed intermediate level language for implementing smart contracts. In this talk, we describe a SCILLA compiler targeting LLVM, with a focus on mapping SCILLA types, values, and its functional language constructs to LLVM-IR.

The compiled LLVM-IR, when executed with LLVM’s JIT framework, achieves a speedup of about 10x over the reference interpreter on a typical SCILLA contract. This reduced latency is crucial in the setting of blockchains, where smart contracts are executed as parts of transactions, to achieve peak transactions processed per second. Experiments on the Ackermann function achieved a speedup of more than 45x.

This talk is aimed at both programming language researchers looking to implement an LLVM based compiler for their functional language, as well as at LLVM practitioners.

1 Introduction

SCILLA (Smart Contract Intermediate Level Language) [16] is a functional programming language in the ML family, designed to enable writing safe smart contracts. SCILLA is the main smart contract language of the ZILLIQA blockchain [18]. Contracts written in SCILLA are deployed *as-is* (i.e., as source code) and are interpreted during transaction validation by the nodes participating in the blockchain consensus (*aka* miners). The project is under active development and is free and open-source [12–14].

The SCILLA Virtual Machine (SVM) (Fig. 1) aims to eventually replace the SCILLA interpreter with faster execution, thus achieving higher transaction throughput. This project comprises of the SCILLA compiler (SC) which translates SCILLA to LLVM-IR, the SCILLA runtime library (SRTL) and the JIT driver which JIT compiles the LLVM-IR and executes it. Execution of a SCILLA contract may result in accessing the blockchain state, with such interactions facilitated by the runtime library. The runtime library also implements SCILLA built-in operations.

In the course of this talk, we will briefly describe the phases of our compiler pipeline (Sec. 2), elaborate on mapping the SCILLA AST to LLVM-IR (Sec. 3), present characteristic benchmarks used to evaluate our compiler (Sec. 4), and share our

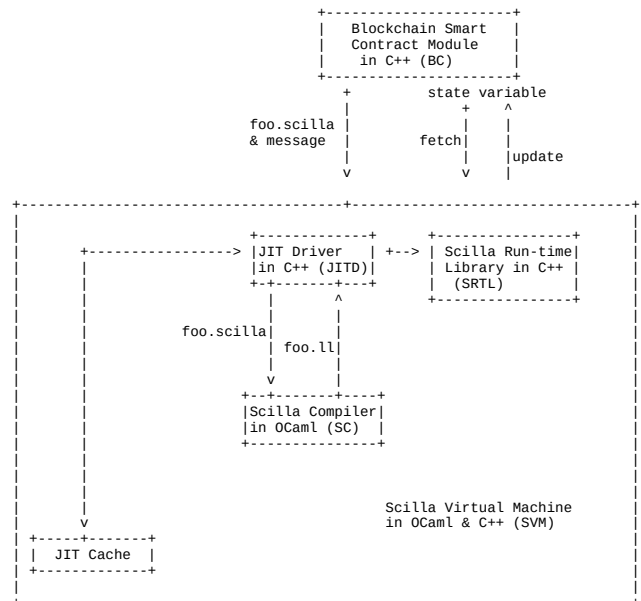


Fig. 1. Design of the SCILLA Virtual Machine

vision on some desired LLVM features which would facilitate projects similar to our own (Sec. 5).

2 The Compiler Pipeline

SCILLA’s reference interpreter and the accompanying type-checker are written in OCaml. To take advantage of the developed infrastructure, we decided to write the compiler also in OCaml, thus enabling re-using much of the frontend code and the type-checker. The compiler uses LLVM’s OCaml bindings to generate LLVM-IR.

Serving as a background for the next section, we now briefly discuss our compiler pipeline. While in this proposal we use textual description, our talk will make use of examples to illustrate the transformations performed in each pass.

2.1 Parsing and Type Checking

The parser and type checker reuse code from the reference interpreter and the result is an AST with each identifier

```

1  fun (p : List (Option Int32)) =>
2  match p with
3  | Nil => z
4  | Cons (Some x) xs => x
5  | Cons _ _ => z
6  end

```

Fig. 2. Nested pattern in a match-expression.

annotated with its type. The type annotations are necessary for code generation later on.

2.2 Dead Code Elimination

SCILLA follows a whole-program compilation model. This means that along with the program currently being compiled, we also compile imported libraries. While the DCE pass is useful in the traditional sense, our primary goal for it here is to simply eliminate unused imported library functions, to improve compile time.

2.3 Pattern Match Flattening

SCILLA allows for nested patterns in a match expression. For example, the expression shown in Fig. 2 requires checking for a `Some` constructor inside an object created with the `Cons` constructor. To efficiently translate pattern match constructs into an LLVM switch statement, match expressions with nested patterns are flattened into nested match expressions with flat patterns [10].

2.4 Uncurrying

As is common in functional programming languages, SCILLA functions and their applications (calls) follow curried semantics [2]. This means that a function taking n arguments is actually a function that takes in one argument and returns a function to process the remaining $n - 1$ arguments and produce the result. The uncurrying pass transforms the AST to have regular C-like (and LLVM-IR) function call semantics, *i.e.*, functions may take $n \geq 0$ arguments, and their calls will provide exactly n arguments. The change in semantics is accompanied by an optimization pass that analyzes all calls of a function, and when all of them provide the same n number of arguments, the function is translated to a function taking n arguments. If it cannot prove so, then the function is translated to take just one argument, and a nested function handling the rest. An example: If all calls of the function $f : T1 \rightarrow T2 \rightarrow T3$ provide two arguments (of type $T1$ and $T2$), then f is transformed to have the type $f : (T1, T2) \rightarrow T3$.

At the time of writing this proposal, our uncurrying pass changes the semantics of the AST to use uncurried semantics, but we do not have the optimization implemented yet.

2.5 Monomorphize

Parametric polymorphism (also called generic programming) is typically implemented using type-erasure (Java, ML, etc) or monomorphization (C++, Rust). The former erases all

type information in generic code after type-checking and generates a single copy of the code. This requires that runtime values be boxed (*i.e.*, values of all types represented uniformly [8]) so that they can all be handled uniformly in the generic code. On the other hand monomorphization specializes the code for each possible (at runtime) ground type (*i.e.*, types that don't have a type parameter) [17], creating copies of the code. While this results in code replication, each copy can be efficiently compiled because the runtime values no longer need to be boxed [3].

A third approach to parametric polymorphism is to wait until polymorphic code is used, and specialize it then (typically using JIT compilation). This approach is used in C# [5] and Ada [6].

SCILLA is based on System F [4, 15]— an expressive type system with *explicit* type bindings and instantiations. We implement parametric polymorphism by specializing each polymorphic SCILLA expression with all possible ground types. This uses a higher-order data-flow analysis to track flow of types into polymorphic expressions [9]. As type instantiation may happen in stages for SCILLA expressions with several type parameters, we choose to compile a polymorphic expression as a dynamic dispatch table, associating ground types with the corresponding monomorphic version of the expression. This provides for a simpler compilation that saves us from replicating code outside of the polymorphic expression, but at the cost of a runtime indirection indexing into the (possibly nested) dynamic dispatch tables. While this is a mix of existing strategies to implement parametric polymorphism, we are not aware of one that uses this exact combination.

2.6 Closure Conversion

As is typical of higher-order languages, SCILLA allows one to pass functions as arguments to other functions and return them, as well as to define functions nested inside other functions. This means a function may have not only its arguments available for use inside, but also a set of “free variables” that are defined outside its syntactic definition.

The closure conversion pass [1] computes the set of free variables for each function, constructs an aggregate type to hold all the free variables of the function (conventionally called the function’s environment) and lifts the function definition to the top / global level, as is required for LLVM-IR generation (where all functions are defined at the top level). The lifted function takes the environment as an argument.

When functions are used as values in the program, each such variable is now a pair of (1) pointer to the function definition and (2) the environment. Function applications (calls) are translated into calling the function pointer, with the paired environment being passed as an additional argument.

In our implementation, this pass also flattens the AST by transforming nested **let-in** expressions into a sequence of assignment statements.

3 Mapping SCILLA to LLVM-IR

The final stage of the compiler pipeline translates the closure converted AST to LLVM-IR.

3.1 Type Descriptors

Although SCILLA code is monomorphized by the time we generate LLVM-IR for it (which means we know the ground type of code and values that we’re generating code for), at times it is necessary for code in the runtime library to operate on SCILLA values. For example, many SCILLA contract executions end with sending out a message (which for the purposes of this discussion is just a composite SCILLA value). The message needs to be serialized into a JSON. It is much simpler to do the JSON serialization outside in the runtime library rather than generate code to do that.

Such a scenario gives rise to the requirement that we communicate the type of a SCILLA value to the runtime library. To this end we define type descriptor structs, defined to be identical in both the runtime library and in the generated code. Every type in the program being compiled will have a type descriptor built for it. A global type descriptor table is added into the compiled LLVM module, and interactions with the runtime library that require the type of a value will use the type’s index in this table.

3.2 Primitive Types

- **Integer types:** SCILLA has signed and unsigned integer types with widths 32, 64, 128 and 256. These types are mapped to LLVM integer types, but with a wrapper struct type. The wrapper struct enables having a 1-1 relation b/w SCILLA types and their translated LLVM types (because LLVM does not distinguish b/w signed and unsigned integers). For example, `Int32` translates to the LLVM type `%Int32 = type { i32 }`
- **String types:** `String` and `ByStr` (byte strings of arbitrary size) in SCILLA translate to LLVM structs containing a pointer to their contents and an integer field with the size.
- **ByStrX:** Fixed-sized byte strings (*i.e.*, X is known at compile time) translate to LLVM array type `[X × i8]`.

3.3 ADTs

Algebraic Data Types (also called variants or tagged unions) are composite types that have one or more “constructors” (sum type), each of which allow defining the value to be a tuple (product type). Figure 3 shows an integer list defined in SCILLA. It has two constructors `Nil` and `Cons`¹.

Because different constructors may be used to construct a value of a specific ADT, ADTs cannot easily be represented unboxed. We represent ADT values via a pointer to a packed struct containing the actual ADT value.

```

1  type MyList =
2      | Nil
3      | Cons of Int32 MyList

```

Fig. 3. Defining an integer list in SCILLA.

For an ADT `tname` with constructors `cname1`, `cname2`, ... `cnameN`, we define LLVM types `%tname = type { i8, %cname1*, %cname2*, ... %cnameN* }`, `%cname1 = type <{ i8, [types in cname1’s tuple]>`, `%cname2 = ...` etc. The LLVM type `%tname` acts as the placeholder for all values of this ADT, and a pointer to this type is used to represent the boxed ADT values. It is always only dereferenced for its first `i8` field, which is the tag representing the specific constructor used to build this value. The other fields are defined only for type completeness, with the idea that, in the future, we can type-check the LLVM-IR at the level of SCILLA types. The tag field is dereferenced at pattern matches, and based on which branch is taken, the pointer is cast to a `%cnameI` pointer and used.

We use packed LLVM structs to represent ADT values so that, when we build or deconstruct ADT values in the runtime library, we do not have to bother with architecture specific struct packing / padding.

3.4 Maps

Maps are boxed and hence handled using an opaque pointer. We rely on the runtime library to create a map and perform operations on it. In the runtime library, maps are defined as `unordered_map<string, any>`. This map’s value type is `std::any` because it needs to be able to represent any SCILLA type, including nested maps and program specific user-defined types.

3.5 Messages

Message types in SCILLA are used to create events, exceptions and outgoing messages. Message objects are created similar to tuples in other languages, but with each component being given a name. In other words, they do not have an explicit predefined type. To enable the runtime library to work with Message objects, we encode them as a sequence of triples, each consisting of a name, type descriptor, and the value itself. The sequence is headed by an integer indicating the number of fields.

3.6 Closures

All SCILLA functions are represented as closures, irrespective of whether they have free variables or not. In the generated LLVM-IR, a closure is represented by an anonymous struct type `{ fundef_sig*, void* }` where `fundef_sig` is the signature of the LLVM function definition. The `void *` represents the environment pointer. A stronger type isn’t used for the environment pointer because we want to represent different SCILLA functions with the same type, but whose

¹At the time of writing this, SCILLA does not support *user-defined* self-referencing ADTs. `List` is a builtin type.

environments may be different, in a uniform way. By convention, all generated LLVM functions will take an environment pointer as the first argument. If the function’s return type cannot be “by value”, then the second argument will be a stack pointer (“sret”) where the return value must be stored.

To avoid ABI complexities, generated LLVM functions and the hand written functions in the runtime library that they may call, all follow the simple rule that if the value size is larger than two eightbytes [7], we define that parameter (or return value) to be passed by reference (stack pointer).

4 Evaluation

The SCILLA compiler and the runtime library are still under active development and are not feature complete. We do however have interesting SCILLA programs and synthetic tests that we can compile and execute end-to-end.

1. **Simple-Map**: This is a simple map access contract that is representative of common constructs used in a typical SCILLA contract. It fetches a map entry, does a simple computation on it and stores back the result. While the interpreted execution takes about 5ms to execute this, the compiled code runs in about 0.5ms.
2. **Ackermann**: The Ackermann function, implemented in SCILLA takes about 2.2s to compute `ackermann(3, 7)`. The compiled code computes the same result in about 46ms.
3. **Church Encoding**: This code computes Church-encoded numerals with the base type equal to `UInt32` and two operations on them: addition and multiplication. We compute a term that evaluates to Church-encoded 131099 and then we convert it to the native `UInt32`. The interpreted execution takes about 0.7s while the compiled execution completes with the result in about 4ms.

These experiments do not include the time taken by the SCILLA LLVM-IR compiler or the LLVM JIT compiler, but just the time spent in the actual execution. (1) We haven’t prioritized compile time improvements yet, so we expect it to be high and not worthy of benchmarking. (2) In the production environment, we cache compiled code on disk, and the most frequently used ones in memory. So we expect the impact of compile times to be low.

5 LLVM Feature Requests

LLVM is a mature framework for building compilers. While it has largely simplified writing compiler frontends for new languages, in the course of writing the SCILLA compiler we found the following features, which, if implemented in LLVM, would further simplify and ease projects such as ours.

5.1 DebugIR

While attaching source debug information to LLVM-IR is useful for debugging, at times, for a compiler developer, being able to debug the LLVM-IR itself can save time. LLVM’s `-debug-ir` pass served exactly this purpose, but has long

been removed from the tree. Considering that a [previous attempt](#) at reviving this wasn’t successful, we decided to adapt the code into a standalone tool [11].

5.2 Shared Definitions with Pre-compiled Code

Many SCILLA operations are implemented in the SCILLA runtime library (SRTL). This requires invoking functions in SRTL from the dynamically (JIT) compiled SCILLA code and passing SCILLA values. We currently specify common type and function definitions both when generating LLVM-IR (in the compiler) as well as in `.h` files in SRTL. This requires both definitions to be kept in sync. While this can be automated by compiling the `.h` file to LLVM-IR and using that during compilation, a more systematic (LLVM provided) method for doing this may lead to a cleaner approach.

References

- [1] Andrew W. Appel. 1992. *Compiling with Continuations*. Cambridge University Press.
- [2] Richard S. Bird and Philip Wadler. 1988. Introduction to functional programming. In *Prentice Hall International series in computer science*.
- [3] Richard Eisenberg and Simon Peyton Jones. 2017. Levity polymorphism. (June 2017), 525–539. <https://www.microsoft.com/en-us/research/publication/levity-polymorphism/>
- [4] Jean-Yves Girard. 1972. *Interprétation fonctionnelle et élimination des coupures de l’arithmétique d’ordre supérieur*. Thèse d’État. Université de Paris VII, Paris, France.
- [5] Anders Hejlsberg. 2004. Generics in C#, Java, and C++. A Conversation with Anders Hejlsberg, Part VII by Bill Venners with Bruce Eckel. <https://www.artima.com/intv/generics.html#part2>.
- [6] Xavier Leroy. 1992. Unboxed Objects and Polymorphic Typing. In *POPL*. ACM Press, 177–188.
- [7] H.J. Lu, Michael Matz, Milind Girkar, Jan Hubička, Andreas Jaeger, and Mark Mitchell. 2018. System V Application Binary Interface. (2018). <https://github.com/hjl-tools/x86-psABI/wiki/x86-64-psABI-1.0.pdf>
- [8] Ronald Morrison, Alan Dearle, Richard C. H. Connor, and Alfred L. Brown. 1991. An Ad Hoc Approach to the Implementation of Polymorphism. *ACM Trans. Program. Lang. Syst.* 13, 3 (1991), 342–371.
- [9] Flemming Nielson, Hanne Riis Nielson, and Chris Hankin. 1999. *Principles of program analysis*. Springer.
- [10] Simon L. Peyton Jones. 1987. *The Implementation of Functional Programming Languages*. Prentice-Hall.
- [11] Zilliqa Research. 2020. DebugIR: Debugging LLVM-IR Files. <https://github.com/vaivaswatha/debugir>
- [12] Zilliqa Research. 2020. Scilla - A Smart Contract Intermediate Level Language. <https://github.com/Zilliqa/scilla>.
- [13] Zilliqa Research. 2020. Scilla Compiler. <https://github.com/Zilliqa/scilla-compiler>.
- [14] Zilliqa Research. 2020. Scilla Virtual Machine. <https://github.com/Zilliqa/scilla-vm>.
- [15] John C. Reynolds. 1974. Towards a theory of type structure. In *Programming Symposium (LNCS)*, Vol. 19. Springer, 408–423.
- [16] Ilya Sergey, Vaivaswatha Nagaraj, Jacob Johannsen, Amrit Kumar, Anton Trunov, and Ken Chan Guan Hao. 2019. Safer smart contract programming with Scilla. *PACMPL* 3, OOPSLA (2019), 185:1–185:30.
- [17] Andrew Tolmach and Dino P. Oliva. 1998. From ML to Ada: Strongly-Typed Language Interoperability via Source Translation. *J. Funct. Program.* 8, 4 (July 1998), 367–412.
- [18] Zilliqa Team. 2017. The Zilliqa Technical Whitepaper. <https://docs.zilliqa.com/whitepaper.pdf> Version 0.1.