

An LLVM Backend of the COGENT Compiler

by
Zhanlin Shang

Supervisor: Christine Rizkallah

University of New South Wales
May, 2020

ABSTRACT

COGENT is a purely functional programming language for implementing easily verified system code with a certifying compiler developed in Haskell. The current COGENT compiler generates C code and hence many optimizations rely on the C compiler. The task of this project is to implement a compiler backend of COGENT that produces LLVM Intermediate Representation (LLVM IR) instead of C in order to enable low level optimizations on the generated code, as well as to explore the possibility of using Vellvm, a Coq formalization of a core subset of the LLVM IR, to avoid the limitations in the C-based proof tools.

TABLE OF CONTENTS

ABSTRACT	ii
1 INTRODUCTION	1
2 BACKGROUND	1
2.1 COGENT	1
2.2 LLVM	3
3 IMPLEMENTATION	7
3.1 Mapping COGENT Core Language to LLVM	8
3.1.1 COGENT Types to LLVM Types	8
3.1.2 COGENT Expressions to LLVM IR	9
3.2 Interfacing with C	13
4 FUTURE WORK	13
5 RELATED WORK	14
6 CONCLUSIONS	15
A Calling Cogent from C	16
B Comparing Outputs from Different Backends	18
BIBLIOGRAPHY	21

1 Introduction

COGENT [9] is a purely functional programming language that is being compiled to C, a higher order logic embedding and a proof relating the two in the Isabelle/HOL theorem prover. COGENT can assist programmers in writing system programs that are easy to verify[1].

Since COGENT is a system programming language, it is necessary for it to be as efficient as possible. The current approach involves some C verification tools which are limited to a subset of C. This limits the efficiency of the C code that can be produced by the COGENT compiler.

LLVM [6] is a high performance modern (compared to GCC) compiler framework that has been developed over 17 years. Many low level optimizations can be done by LLVM, such as loop unrolling, partial redundancy elimination and tail call optimisation. As LLVM was designed to be a compiler framework, it is easy to add or change the passes of the compilation process.

The goal of this project is to implement a compiler backend of COGENT that targets the LLVM Intermediate Representation (LLVM IR) to avoid the limitations caused by the C verification tools.

To preserve the verifiability of COGENT, we plan to use a verification framework that targets LLVM called Vellvm [15] in the future. Since Vellvm has only formalized a core subset of the LLVM IR, our backend only uses the subset of LLVM IR that is formalized by Vellvm.

This report will introduce COGENT and LLVM first and then demonstrate the implementation of an LLVM backend of a core subset of COGENT.

2 Background

2.1 Cogent

COGENT [9] is designed to be a high level, purely functional language for implementing systems software. The COGENT compiler guarantees the memory safety and totality of COGENT programs.

There are several phases in the compilation of a COGENT program: parsing, type checking, desugaring, normalization, monomorphisation, and then code generation to the target language (C code in the existing compiler and LLVM IR in this project).

At the code generation phase, there is a small set of primitive types. The primitive types that are covered by this project are:

Booleans	:	Bool
Unsigned integers	:	U8, U16, U32, U64
Strings	:	String
Pairs	:	(A, B)
Records	:	{x :A , y : b}

primops	$o \in$	{+, *, /, >, <, >=, <=, ==, !=}
literals	$l \in$	{123, True, 'a', ...}
expressions	$e ::=$	$l \mid x \mid o(\bar{e})$ $\mid \mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2$ $\mid \mathbf{let!} \ (\bar{y}) \ x = e_1 \ \mathbf{in} \ e_2$ $\mid \mathbf{if} \ e_1 \ \mathbf{then} \ e_2 \ \mathbf{else} \ e_3$ $\mid \mathbf{promote} \ \langle \bar{C} \ \bar{\tau} \rangle e$ $\mid \{f = e\} \mid e.f$ $\mid \mathbf{take} \ x \ \{f = y\} = e_1 \ \mathbf{in} \ e_2$ $\mid \mathbf{put} \ e_1.f := e_2$
function def.	$d ::=$	$\langle f :: \pi, f \ x = e \rangle \mid \langle f :: \pi, \blacksquare \rangle$
programs	$P ::=$	\bar{d}
function names	\ni	f, g
variable names	\ni	x, y
constructors	\ni	A, B, C
record fields	\ni	f, g

Figure 1: The Syntax of COGENT Core Language Subset Covered

primops	$o \in$	{ , <<, ...}
expressions	$e ::=$	$() \mid f[\bar{\tau}]$ $\mid \mathbf{cast} \ t \ e$ $\mid \mathbf{case} \ e_1 \ \mathbf{of} \ Cx \rightarrow e_2 \ \mathbf{else} \ y \rightarrow e_3$ $\mid \mathbf{esac} \ e \mid C \ e$

Figure 2: The Syntax of COGENT Core Language Subset Not Covered

The syntax of COGENT [9] covered by this project is shown in Figure 1:

For primitive operations, they are defined on ints.

For literal values, types `int`, `boolean` and `string` are defined.

For expressions, literal value, variable, binary operations, let binding (and let bang), if branching, integer type promoting, record definition and record field accessing are implemented.

For functions, both function definitions and abstract function declarations are supported. Abstract function declaration will be treated as a `extern` declaration.

Some parts of the core syntax is left for future work due to the time limitations, as shown in Figure 2.

2.2 LLVM

LLVM [6] is a modern compiler framework that provides a language independent code representation (LLVM IR) which can be used for multiple purposes, including analysis and transformation. There are many analyses and transformation passes included by the LLVM framework, such as printing function call graphs, printing control flow graphs, dead code elimination, function inlining and loop unrolling.

In addition to assembly, LLVM includes a simple type system which is primitive in mainstream languages, such as integers of arbitrary length, floating point numbers, structures, arrays, functions and pointers, and operations on these types. LLVM also hides machine-specific details such as physical registers and calling conventions [6].

LLVM provides an infinite set of virtual registers in the Static Single Assignment (SSA) form [6], i.e. each virtual register can be written by only one instruction and an assignment of a new value to a variable will be represented as initialization of a new variable based on the old value, the instructions used in this report are explained in Figure 3. For example, the following C code:

```
1 int x = 1;
2 x = 2;
3 int y = x;
```

can be represented in LLVM IR as:

```
1 %x.1 = alloca i32 ; allocate a 32 bits int and assign the address to %x.1
2 %x.2 = alloca i32 ; allocate a 32 bits int and assign the address to %x.2
```

```

3 %y = alloca i32 ; allocate a 32 bits int and assign the address to %y
4 store i32 1, i32* %x.1 ; store a 32 bits int value 1 to address %x.1
5 store i32 2, i32* %x.2 ; store a 32 bits int value 2 to address %x.2
6 %tmp = load i32, i32* %x.2 ; load the value from address %x.2 into %tmp
7 store i32 %tmp, i32* %y ; store the value of %tmp to address %y

```

- *alloca type*:
Allocate memory for *type*
- *store type value, type * place*:
Put *value* into the address pointed by *place*
- *load type, type * place*:
Load value of type *type* from the address pointed by *place*
- *ret type operand*:
Return the value of *operand* from the current block, operand can be either literal value or register
- *getelementptr type, type * place, i32 idx {, i32 idx}**:
Get from the aggregate type pointed by *place* the pointer to the element determined by *idxs*, for example, given `int * a`, `a[1]` will be `getelementptr i32, i32 * %a, i32 1`
- *insertvalue agg-type agg, elm-type elm, idx*:
Insert *elm* into the *idx* place of the aggregate type value *agg*
- *call ret-type name (arg-type arg, ...)*:
Call the function *name* with argument *arg* and get the returned value as type *ret-type*

Figure 3: Subset of LLVM Instructions covered in this report

SSA can be seen as equivalent to a well-behaved subset of Continuation-Passing Style (CPS) [5]. In CPS [12], a procedure does not return, instead it is passed an additional “call back” argument, which is applied to the procedure’s return value. There are many studies on optimising CPS and the optimisations for CPS can all be applied for SSA.

In LLVM IR, the minimal unit is an **Instruction**. LLVM uses RISC-like instructions [6], plus higher level information such as types. An instruction has the form:

```

1 instruction-name arg_1, arg_2, ... , arg_n

```

where `args` can be one of the followings:

- operand: can be literal or register
e.g. in `add i32 1 %x`
both `1` and `%x` are operands
- type: can be type or type name
e.g. in `load i32, i32* %x`
`i32` and `i32*` are types
- attribute: different instructions have different attributes
e.g. in `alloca i32, align 4`
and `add nuw nsw i32 1, 2`
`align 4`, `nuw` and `nsw` are all attributes

The LLVM instruction set consists of several kinds of instructions: terminator instructions, binary instructions, bitwise binary instructions, memory instructions and other instructions [7]. Most of LLVM instructions are in three-address form: they take one or two operands and produce a single result [6].

Instructions are contained by **Blocks**, which form the body of a program, e.g. function body, conditional branches, etc. A block contains one and only one terminator. Here is a sample program in LLVM IR:

```
1 define i8 @foo({ i8, i8 }* %0) {
2 entry: ; this is a label, LLVM has first class labels that can be used for jumping
3   %1 = getelementptr inbounds { i8, i8 }, { i8, i8 }* %0, i32 0, i32 0
4   ; get pointer to the first field of %0
5   %2 = load i8, i8* %1 , align 1
6   ; load the value from %1 to %2
7   %3 = getelementptr inbounds { i8, i8 }, { i8, i8 }* %0, i32 0, i32 1
8   ; get pointer to the second field of %0
9   %4 = load i8, i8* %3 , align 1
10  ; load the value from %3 to %4
11  %5 = add i8 %2, %4
```



```

12   ; add %2 and %4, store result in %5
13   ret i8 %5
14   ; return %5
15 }
16
17 define i8 @main() {
18   %1 = alloca { i8, i8 }, align 1
19   ; allocate memory for a struct { i8, i8 }
20   %2 = insertvalue { i8, i8 } undef, i8 1, 0
21   ; insert the value 1 to the first field of { i8, i8 } (starts from undef)
22   ; and store the struct in %2
23   ; %2 is { 1, undef }
24   %3 = insertvalue { i8, i8 } %2, i8 2, 1
25   ; insert the value 2 to the second field of %2
26   ; %3 is { 1, 2 }
27   store { i8, i8 } %3, { i8, i8 }* %1
28   ; store the value of %3 to address %1
29   %4 = call i8 @foo({ i8, i8 }* %1)
30   ; call function @foo with %1 as argument and store the returned value in %4
31   ret i8 %4
32   ; return %4
33 }

```

This sample program defines two functions, one named `foo` and the other named `main`. Similar to C, an LLVM program starts with the function `main`.

In this `main` function, firstly we allocate memory for a structure that contains two 8-bit integers, then we assign the two fields in the allocated structure to 1 and 2 respectively, next we call the function `foo` with the allocated structure as argument and return the result from `main`.

In this `foo` function, firstly we take values from both fields in the structure given as argument, then we compute the sum of them and return the result from `foo`.

3 Implementation

The LLVM backend of COGENT compiler is a drop-in replacement of the existing C code generation pass. The following two components are of the most importance in the code generation process:

- `data Type t`
- `data Expr t v a e`

They correspond to COGENT’s types and expressions respectively and the mapping between them to LLVM IR will be explained in the following sections.

In this project, we use a LLVM’s Haskell binding named “llvm-hs” [2], which is an one-to-one binding from the LLVM IR to Haskell, with abstraction for generating conditional branching.

In the LLVM backend, two data types are responsible for generating instructions and blocks:

- `data BlockState`
- `data CodegenState`

`BlockState` is used to track the states of a basic block, such as instructions in a block and the terminator of a block. `CodegenState` is used to track the general states of the compilation, such as blocks contained in this program, the block we are currently working with, variables indexing, virtual registers usages, etc.

The following utility functions for the code generation will be used in the later sections:

- `fresh :: Codegen Word`

This generates a new virtual register with a fresh name.

- `instr :: Type → Instruction → Codegen Operand`

This function takes an instruction and stores the result into a newly assigned register by `fresh`, then returns the register.

- `addBlock :: ShortByteString → Codegen Name`

This function takes a name and creates a new block in the current `CodegenState`.

- `setBlock :: Name → Codegen Name`

This function takes a name and sets the block bound to the name as the working block. This is essential when generating codes for branching.

3.1 Mapping Cogent Core Language to LLVM

In the typechecked core language, expressions are of type `TypedExpr t v a` which states the type of the expression. Therefore, mapping the COGENT types to LLVM types should occur before compiling the COGENT expressions.

3.1.1 Cogent Types to LLVM Types

The function `toLLVMType` maps the following COGENT types to LLVM types:

- **Unit:**

Since the type `Unit` is similar to `void` in C, we generate `VoidType` of LLVM for the unit type in COGENT.

- **PrimInt:**

In COGENT, there are 5 primitive unsigned integer types: `Bool`, `U8`, `U16`, `U32`, `U64`. Since LLVM doesn't differentiate signed and unsigned value at the type level and arbitrary length integer value are allowed, denoted in LLVM as `in` where n can range from 1 to $2^{23} - 1$, we can generate integer types with corresponding length in LLVM.

In `llvm-hs`, the constructor `IntegerType` takes an integer n and produces the corresponding LLVM integer type `in`, we have `IntegerType 1` to `IntegerType 64` for the COGENT primitive integer types respectively.

- **Record and Pair:**

In LLVM, a structure type is defined as:

```
1 type { type1, type2, ..., typen }
```

In `llvm-hs` the constructor `StructureType` takes a list of LLVM types and produces the required type.

Since a COGENT record is a structure and a product is essentially a record without names, therefore, for a record R with field types $f_1 \dots f_n$, we can have:

$$\text{toLLVMType } R = \text{StructureType}\{\text{toLLVMType } f_i \mid i \in [1, n]\}$$

- **String:**

Since strings are char pointers (i.e. `char *`), they are mapped to `PointerType{IntegerType 8}`.

3.1.2 Cogent Expressions to LLVM IR

COGENT expressions are compiled by the function:

$$\text{exprToLLVM} :: \text{TypedExpr } t \ v \ a \rightarrow \text{Codegen } (\text{Either Operand (Named Terminator)})$$

The `Either` in the type signature determines whether the returned value is a terminator or not. When a value needs to be returned (e.g. the last expression in a `let` block), a terminator `Ret` will be constructed from the value and then returned as `Right`, otherwise we just return the value as `Left`.

The following expressions are currently covered:

- `ILit` *val bits*:

Since arbitrary length integer values in the LLVM IR are primitive, for a COGENT integer of length *bit* and value *val*, we generate:

$$\text{ConstantOperand Int } \{bits, val\}$$

- `Op` *op [operand₁, operand₂]*:

First we get both operands by recursively generating instructions from *operand₁* and *operand₂* and we have to ensure that both operands produce terminator values, since it doesn't make sense to have code: `1 + (return 2)` in C.

Then we generate an instruction according to *op* and construct the instruction through `instr` to ensure that the returned value can be used as an operand by its upper level.

- Variable *idx*, *name*:

In core COGENT, de Bruijn indices are used to denote the variables used, but LLVM uses an incremental indexing that starts from 0 in each function. For example, the COGENT program:

```
1 foo : (U8, U8) -> U8
2 foo (a, b) = a + b
```

will be A-normalised into (with irrelevant parts removed):

```
1 foo = take (a, x_1) = x_0<_v0> { .0 }
2   in take (b, x_2) = x_1<_v1> { .1 }
3     in a<_v2> + b<_v0>
```

The equivalent LLVM IR is as follows:

```
1 define i8 @foo({ i8, i8 }* %0) {
2 entry:
3   %1 = getelementptr inbounds { i8, i8 }, { i8, i8 }* %0, i32 0, i32 0
4   %2 = load i8, i8* %1, align 1 ; a
5   %3 = getelementptr inbounds { i8, i8 }, { i8, i8 }* %0, i32 0, i32 1
6   %4 = load i8, i8* %3, align 1 ; b
7   %5 = add nuw i8 %2, %4
8   ret i8 %5
9 }
```

To resolve this problem, list named `indexing` in `CodegenState` is introduced. Whenever a new value is introduced (e.g. entering a new scope), the new value will be added to the front of `indexing`, then the order will correspond to de Bruijn index.

- Struct [*fields*]:

In generating instructions for a structure, memory is firstly allocated for the struct, then values of each field in the struct is set.

Therefore we have the following algorithm (pseudo code) to generate structures:

```
1 struct ← Alloca { type = recordType }
2 flds ← [ GetElementPtr from struct then Store exprToLLVM(f) | f ∈ fields]
```

```

3 return fold (>>) struct flds
4 -- >> ensures the sequencing of the generation of each fields in the struct

```

- **Let** *name val body* / **Let!** *var name val body*:

The intuitive translation from a **let** is (in de Bruijn's notation):

$$(\lambda.body)(val)$$

As literally translation of anonymous functions to LLVM is not efficient, we instead use the **indexing** method as introduced in the **Var** section above.

- **Take** (*name_a, name_b*) *record field body*:

This expression is similar to **let**. At the stage of code generation, the linearity is already checked in previous passes, we can take this equivalence as an intuition:

$$\mathbf{take} (a, b) r f \mathbf{body} \Leftrightarrow \mathbf{let} a = r.f \mathbf{in} \mathbf{let} b = r \mathbf{in} \mathbf{body}$$

Hence, following the same idea as generating **let**, we load the field value into a fresh register with the **load** instruction, and push the value and the record to the front of **indexing**. Then, with the extended **indexing**, we generate instructions for *body*.

- **Put** *record field val*:

As opposed to **take**, **put** stores the value into a structure at a certain field. We need to generate the pointer to the field and construct the **store** instruction of LLVM to store the value to the address of the target field.

- **If** *cond then else*:

In SSA, conditional branching is achieved with the help of the Φ function [10]. The idea of the Φ function can be demonstrated with the following program:

```

1 x = 5
2 x = x - 3
3 if x > 3:
4   y = 1

```

```

5  else:
6    y = 2
7  w = y + x

```

A new definition of y will be generated and initialized to “choose” the correct branch:

```

1  x_1 = 5
2  x_2 = x_1 - 3
3  if x_2 > 3:
4    y_1 = 1
5  else:
6    y_2 = 2
7  y_3 =  $\Phi$ (y_1, y_2)
8  w = y_3 + x_2

```

As mentioned above, `llvm-hs` provides an abstraction for conditional branching, a special terminator named `CondBr`. Therefore we can write:

```

1  currentBlk <- gets currentBlock
2  blkTrue <- addBlock "brTrue"
3  blkFalse <- addBlock "brFalse"
4  setBlock blkTrue
5  exprToLLVM tb
6  setBlock blkFalse
7  exprToLLVM fb
8  setBlock currentBlk
9  CondBr { condition = cond
10         , trueDest = blkTrue
11         , falseDest = blkFalse
12         }

```

This produces the following LLVM program:

```

1  entry:
2  %1 = ;condition
3  br i1 %1, label %brTrue, label %brFalse
4  brTrue: ; preds = %entry
5  ...
6  brFalse: ; preds = %entry
7  ...

```

3.2 Interfacing with C

With the support of the LLVM infrastructure, inter-operation between COGENT and C can be simple:

- To use COGENT function in C, the COGENT function need to be marked as `extern`
- To use C function in COGENT, an abstract function declaration will be compiled into external linkage by default.

Then by compiling and assembling each components to LLVM bitcodes and linking them using ‘llvm-link’, we will have an executable as desired. An example is included in Appendix A.

4 Future Work

There is still a considerable number of features not implemented yet in this backend and the most important features still lacking are:

- Sum type (i.e. tagged union)
- Case / Esac (`case` and `esac` takes an value of sum type and branches for each case)

Due to the time limitations, these have been left for future development.

Our proposed approach on them is: we can append a “tag” to the value of sum type, e.g. for a sum type $R = \langle E \mid S \rangle$, we generate:

$$\%R = \{\text{i8}, \max(\text{sizeof}(E), \text{sizeof}(S))\}$$

In which the `i8` will be the tag and this should be used in C / C++ in the same way as:

```
1 enum tag : uint_8t { ... };
2 struct s {
3     tag t;
4     union {
5         E ...
```



```
6     s ...
7   } v;
8 };
```

Once we have the sum type implemented in this way, we can compile `case` and `esac` as `switch` statements in C.

5 Related Work

There are many functional programming languages that have LLVM backend, such as Haskell (GHC's LLVM backend) [14] and many Lisp dialects (e.g. Clasp [11], Open Dylan [4]). Modern languages with functional features such as Rust and Swift also target LLVM.

Because of historical reasons, the LLVM backend of GHC compiles the old intermediate language C minus minus (Cmm) to LLVM. Therefore, our work with the LLVM backend of COGENT compiler is more related to the Cmm generation.

Clasp [11] is a compiler for Common Lisp which aims to enable the interfacing between Common Lisp and C++. Their implementations could act as a reference when extending our LLVM backend to support interfacing with C, especially on the handling polymorphic types.

Swift and Rust are both languages with functional features such as first class closure without the use of runtimes or virtual machines. Swift has an intermediate language, Swift Intermediate Language (SIL) [3], which provides native support for closure capture. Rust also adopts a similar design by having an intermediate language named *mid-level* IR (MIR) [8]. The closure expansion of Rust is done on the MIR level [13]. As we are planning to add the support of closures in COGENT, the implementations of SIL and MIR could be good references for the future development.

A common feature among these works are, they all have intermediate languages between their core languages and LLVM IR. We may also want to adopt this approach by having an IR before LLVM IR so that we can bridge up the abstraction gap between core COGENT and LLVM, for example, when we are adding closure supports to COGENT in the future, having a mid level IR could help to resolve the captures.

6 Conclusions

We have presented a working prototype of LLVM backend of the COGENT compiler. As there is no longer a reliance on the C compilers, we can produce outputs that better suit our needs from the same COGENT input, such as real 1 bit `boolean`, and produce smaller binary (see Appendix B). This even allows us to combine COGENT program with not only C, but any language that is capable of being compiled into LLVM bitcode.

A Calling Cogent from C

A simple example for calling COGENT from C:

```
1 #include <stdlib.h>
2
3 typedef struct {
4     char a;
5     char b;
6 } arg;
7
8 extern char foo( arg * );
9
10
11 int main(){
12     arg * a = malloc(sizeof(arg));
13     a->a = 1;
14     a->b = 2;
15     char res = foo(a);
16     free(a);
17     return res;
18 }
```

Figure 4: main.c

```
1 foo : (U8, U8) -> U8
2 foo (a,b) = a + b
```

Figure 5: adder.cogent

```
1 CLANG=clang
2 AS=llvm-as
3 LINK=llvm-link
4 COGENT=cogent
5
6 all: main.bc adder.bc
7     $(LINK) main.bc adder.bc -o main
8     chmod +x main
9
10 main.bc: main.ll
11     $(AS) main.ll -o main.bc
12
13 main.ll: main.c
14     $(CLANG) -S -emit-llvm main.c -o main.ll
15
16 adder.bc: adder.ll
17     $(AS) adder.ll -o adder.bc
18
19 adder.ll: adder.cogent
20     $(COGENT) --llvm adder.cogent
```

Figure 6: Makefile

B Comparing Outputs from Different Backends

For the input file:

```
1 foo : (U8, U8) -> U8
2 foo (a,b) = a + b
```

Figure 7: simple-arith.cogent

different backends generate different outputs (some metadata were removed):

```

1 %struct.t1 = type { i8, i8 }
2
3 ; Function Attrs: noinline nounwind optnone uwtable
4 define dso_local signext i8 @foo(i16 %0) #0 {
5     %2 = alloca %struct.t1, align 1
6     %3 = alloca i8, align 1
7     %4 = alloca i8, align 1
8     %5 = alloca i8, align 1
9     %6 = bitcast %struct.t1* %2 to i16*
10    store i16 %0, i16* %6, align 1
11    %7 = getelementptr inbounds %struct.t1, %struct.t1* %2, i32 0, i32 0
12    %8 = load i8, i8* %7, align 1
13    store i8 %8, i8* %3, align 1
14    %9 = getelementptr inbounds %struct.t1, %struct.t1* %2, i32 0, i32 1
15    %10 = load i8, i8* %9, align 1
16    store i8 %10, i8* %4, align 1
17    %11 = load i8, i8* %3, align 1
18    %12 = sext i8 %11 to i32
19    %13 = load i8, i8* %4, align 1
20    %14 = sext i8 %13 to i32
21    %15 = add nsw i32 %12, %14
22    %16 = trunc i32 %15 to i8
23    store i8 %16, i8* %5, align 1
24    %17 = load i8, i8* %5, align 1
25    ret i8 %17
26 }

```

Figure 8: simple-arith.ll by compiling the generated C code

```
1 define i8 @foo({ i8, i8 }* %0) {
2 entry:
3   %1 = getelementptr inbounds { i8, i8 }, { i8, i8 }* %0, i32 0, i32 0
4   %2 = load i8, i8* %1, align 1
5   %3 = getelementptr inbounds { i8, i8 }, { i8, i8 }* %0, i32 0, i32 1
6   %4 = load i8, i8* %3, align 1
7   %5 = add nuw i8 %2, %4
8   ret i8 %5
9 }
```

Figure 9: simple-arith.ll by using the LLVM backend

Bibliography

- [1] AMANI, S., HIXON, A., CHEN, Z., RIZKALLAH, C., CHUBB, P., O’CONNOR, L., BEEREN, J., NAGASHIMA, Y., LIM, J., SEWELL, T., TUONG, J., KELLER, G., MURRAY, T., KLEIN, G., AND HEISER, G. Cogent: Verifying high-assurance file system implementations. In *International Conference on Architectural Support for Programming Languages and Operating Systems* (Atlanta, GA, USA, Apr. 2016), pp. 175–188.
- [2] COWLEY, A., DIEHL, S., KIEFER, M., AND SCARLET, B. S. *llvm-hs: General purpose LLVM bindings*, 2020 (accessed Apr. 2020). <https://hackage.haskell.org/package/llvm-hs>.
- [3] GROFF, J., AND LATTNER, C. *Swift Intermediate Language: A high level IR to complement LLVM*, 2015 (accessed Apr. 2020). <https://llvm.org/devmtg/2015-10/slides/GroffLattner-SILHighLevelIR.pdf>.
- [4] HOUSEL, P. S. LLVM Code Generation for Open Dylan. In *Proceedings of the 13th European Lisp Symposium on European Lisp Symposium* (2020), ELS2020, European Lisp Scientific Activities Association.
- [5] KELSEY, R. A. A correspondence between continuation passing style and static single assignment form. In *Papers from the 1995 ACM SIGPLAN Workshop on Intermediate Representations* (New York, NY, USA, 1995), IR ’95, Association for Computing Machinery, p. 13–22.
- [6] LATTNER, C., AND ADVE, V. LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In *Proceedings of the 2004 International Symposium on Code Generation and Optimization (CGO’04)* (Palo Alto, California, Mar 2004).
- [7] LLVM DEVELOPERS. *LLVM Language Reference Manual*, 2020 (accessed Apr. 2020). <https://llvm.org/docs/LangRef.html>.
- [8] MATSAKIS, N. *Introducing MIR*, 2016 (accessed Apr. 2020). <https://blog.rust-lang.org/2016/04/19/MIR.html>.
- [9] O’CONNOR, L., CHEN, Z., RIZKALLAH, C., AMANI, S., LIM, J., MURRAY, T., NAGASHIMA, Y., SEWELL, T., AND KLEIN, G. Refinement through restraint: Bringing down the cost of verification. In *International Conference on Functional Programming* (Nara, Japan, Sept. 2016).
- [10] ROSEN, B. K., WEGMAN, M. N., AND ZADECK, F. K. Global value numbers and redundant computations. In *Proceedings of the 15th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (New York, NY, USA, 1988), POPL ’88, Association for Computing Machinery, p. 12–27.
- [11] SCHAFMEISTER, C. A., AND WOOD, A. Clasp common lisp implementation and optimization. In *Proceedings of the 11th European Lisp Symposium on European Lisp Symposium* (2018), ELS2018, European Lisp Scientific Activities Association.
- [12] SUSSAM, G. J., AND STEELE JR., G. L. SCHEME: An Interpreter For Extended Lambda Calculus. *AI Memo*, 349 (1975).

- [13] TEAM, R. C. *Closure Expansion in Rustc*, (accessed May. 2020). <https://rustc-dev-guide.rust-lang.org/closure.html>.
- [14] TEREI, D. A., AND CHAKRAVARTY, M. M. An llvm backend for ghc. In *Proceedings of the Third ACM Haskell Symposium on Haskell* (New York, NY, USA, 2010), Haskell '10, Association for Computing Machinery, p. 109–120.
- [15] ZHAO, J., NAGARAKATTE, S., MARTIN, M. M., AND ZDANCEWIC, S. Formalizing the llvm intermediate representation for verified program transformations. *SIGPLAN Not.* 47, 1 (Jan. 2012), 427–440.