



# Lambda the Ultimate SSA: Optimizing Functional Programs in SSA

Siddharth Bhat  
CSTAR  
IIIT Hyderabad  
Hyderabad, India  
siddharth.bhat@research.iiit.ac.in

Tobias Grosser  
School of Informatics  
University of Edinburgh  
Edinburgh, United Kingdom  
tobias.grosser@ed.ac.uk

**Abstract**—Static Single Assignment (SSA) is the workhorse of modern optimizing compilers for imperative programming languages. However, functional languages have been slow to adopt SSA and prefer to use intermediate representations based on minimal lambda calculi due to SSA’s inability to express higher-order constructs. We exploit a new SSA construct — regions — in order to express functional optimizations via classical SSA-based reasoning. Region optimization currently relies on ad-hoc analyses and transformations on imperative programs. These ad-hoc transformations are sufficient for imperative languages as regions are used in a limited fashion. In contrast, we use regions pervasively to model sub-expressions in our functional IR. This motivates us to systematize region optimizations. We extend classical SSA reasoning to regions for functional-style analyses and transformations. We implement a new SSA+regions based backend for LEAN4, a theorem prover that implements a purely functional, dependently typed programming language. Our backend is feature-complete and handles all constructs of LEAN4’s functional intermediate representation  $\lambda\text{rc}$  within the SSA framework. We evaluate our proposed region optimizations by optimizing  $\lambda\text{rc}$  within an SSA+regions based framework implemented in MLIR and demonstrating performance parity with the current LEAN4 backend. We believe our work will pave the way for a unified optimization framework capable of representing, analyzing, and optimizing both functional and imperative languages.

**Index Terms**—Optimizing compilers, Functional programming

## I. INTRODUCTION

Many optimizing compilers for imperative programming languages use Static Single Assignment (SSA) as their intermediate representation (IR) [1][2]. Such IRs impose structural and semantic rules on values to aid the compiler’s reasoning. Traditionally, such IRs express control flow using a Control Flow Graph (CFG) – a flat collection of sequences of instructions (basic blocks) that transfer control amongst one another. On the other hand, functional intermediate representations use an *expression-based* IR, where control flow is represented via particular expressions such as case statements. An IR for functional constructs necessarily needs the ability to manipulate sub-expressions.

The famous slogan “SSA is continuations” [3] is not entirely true as stated; the correspondence between SSA and continuation-passing-style (CPS) is established between SSA and a *syntactically restricted* fragment of CPS, which is *further annotated* with information about which continuations

### A Dead Expression Elimination

```

out =
let x = e in y...
↓
out = y

```

```

%x = rgn.val { e }
%y = ...
return %y
↓
return %y

```

### B Case Elimination

```

out = case True of
True -> e
False -> f
↓
out = e

```

```

%ve = rgn.val { e }
%vf = rgn.val { f }
%r = select true, %ve, %vf
rgn.run %r
↓
rgn.run %ve

```

### C Common Branch Elimination

```

out = case x of
True -> e
False -> e
↓
out = e

```

```

%ve = rgn.val { e }
%vf = rgn.val { e }
%r = select %x, %ve, %vf
rgn.run %r
↓
%ve = rgn.val { e }
rgn.run %ve

```

Fig. 1: We express regions as SSA values to adapt transformations in functional compilers in an SSA setting. We see that eliminating dead let bindings, eliminating case of known values, and fusing case branches are all variants of classical SSA transformations operating on region values declared by `rgn.val` and executed by `rgn.run`.

represent intra-procedural control flow and which represent inter-procedural control flow. Thus, the above translation is not practically useful in order to design an SSA-based intermediate representation for functional programming languages, and most functional IRs continue to use expression-based intermediate representations. This has resulted in a schism within the compiler community, where the infrastructure built around optimizing compilers for imperative languages does not get reused for functional languages and vice versa.

In this work, we set out to heal this schism by providing *convenient, easily analyzable and optimizable* encodings of

core functional IR constructs within SSA. Our key innovation is to use regions to represent *functional sub-expressions as SSA values*. This allows us to extend SSA’s core strength, the ability to reason about values via sophisticated algorithms, to also cover reasoning about sub-expressions. This enables us to mostly reuse SSA algorithms, and in some cases extend SSA algorithms, to recover key functional language optimizations within SSA. We also gain the ability to use the tooling provided by MLIR’s compiler infrastructure for (a) testing, (b) having a stable textual and in-memory representation, (c) sophisticated infrastructure for parallel peephole rewriting, and (d) support for parallelization and vectorization.

We contribute this novel encoding of a functional intermediate representation, and an implementation of the encoding within the MLIR compiler framework. We also design and implement new algorithms and modifications to existing algorithms to reason about regions as sub-expressions. We test our design by building a fully-functional backend for the LEAN4 [4], functional programming language. Our new backend uses the MLIR compiler framework to implement SSA+regions. We embed LEAN4’s intermediate representation,  $\lambda_{rc}$ , within MLIR and show how to perform common functional programming language optimizations on  $\lambda_{rc}$  via SSA-based reasoning. Concretely, our contributions are:

- A minimal intermediate representation  $lp$  which allows functional constructs to be lowered, analyzed, and optimized within an SSA setting.
- A novel use of regions to encode functional constructs within an imperative SSA-style IR, called  $rgn$ .
- The design of functional-style optimizations for  $rgn$ , which are analogues of classical SSA optimizations, applied to regions (Figure 1).
- An evaluation of our feature-complete backend for the LEAN4 compiler and its  $\lambda_{rc}$  intermediate representation which demonstrates the soundness of our approach towards representing functional programs in SSA via  $lp$  and  $rgn$ .

## II. BACKGROUND

We provide background on SSA-based IRs, LEAN4, and functional programming constructs.

### A. Static Single Assignment (SSA) and Regions

An intermediate representation (IR) is in SSA form [5] if each variable is assigned exactly once and no variable remains undefined. SSA has gained popularity in imperative compilers such as GCC [6], LLVM [1], and many others, as data-flow information is explicitly expressed through dependencies from the definition of a value to its uses (def-use chains). SSA-based IRs typically use basic blocks that hold lists of sequentially executed operations, each taking a list of argument values and returning a tuple of return values. Terminator operations at the end of each basic block, which either branch to another basic block or return from a function, combine these basic blocks into a control flow graph (CFG). While IRs typically use a flat CFG, MLIR [7] recently introduced nested control flow as a first-class concept to support abstractions that require control over

scoping. Operations can now receive regions, nested single-entry sub-CFGs, as additional arguments. Regions make it easy to express concepts such as loop bodies, branches of an if-statement, and case statements in functional programs.

### B. LEAN, $\lambda_{pure}$ , and $\lambda_{rc}$

LEAN<sup>1</sup> [8] is an open source theorem prover based on a minimal dependently typed [9] kernel. After type checking the LEAN program, the compiler compiles the program to  $\lambda_{pure}$ , a minimal, pure, strict, higher order intermediate representation that is suitable for further lowering into assembly.  $\lambda_{pure}$  is lowered to  $\lambda_{rc}$ , an extension of  $\lambda_{pure}$  with reference counting. The current LEAN backend then lowers  $\lambda_{rc}$  to C. Runtime features such as primitives for I/O, numerics, reference counting, and task-based parallelism are provided by a custom runtime library, `libleanrt`.

### C. MLIR

MLIR [7] is a new compiler infrastructure that aims to simplify the development of domain-specific compilers. For this, MLIR provides a minimal SSA-based intermediate representation (IR), with the ability to easily instantiate extensions to the core IR (known as *dialects*) which follow SSA conventions. Having a customizable IR allows compiler developers to model domain-specific concerns by introducing custom types, operations, and attributes. Here, we briefly describe the relevant aspects of MLIR that are used in our compiler.

1) *Modules*: Each program in MLIR is called a *Module*. A Module consists of several global functions. Function names such as `@foo` are global and allow for linking function calls across modules. A function consists of basic-blocks. Each basic-block is a sequence of operations, ending with a terminator operation. SSA values are local to the scope of the function and have names of the form `%bar`.

2) *Operations*: SSA values are produced by *Operations*, such as `addi`, which stands for the integer addition operation. Each operation takes zero or more SSA operands that are defined before it and returns zero or more SSA values. Operations can also have compile time constants attached to them, such as `{phase = 90.0 : f64}`. These are called *attributes*.

3) *Dialects*: A *Dialect* in MLIR is a collection of operations and types. The type system in MLIR consists of either primitive types or user defined custom types to encode more complex type systems. In this paper, we use two existing dialects, `std` and `scf`. The `std` dialect contains all basic operations such as constant declarations, arithmetic operations and memory manipulation. The Structured Control Flow Dialect `scf` contains `if-else` and `for` loop constructs.

## III. EXPRESSING FUNCTIONAL PROGRAMS VIA SSA

In this section, we survey  $\lambda_{rc}$ , LEAN’s intermediate representation and discuss our embedding of  $\lambda_{rc}$  into MLIR’s SSA-based compiler IR.

<sup>1</sup>We refer to LEAN 4 when using the versionless expression LEAN.

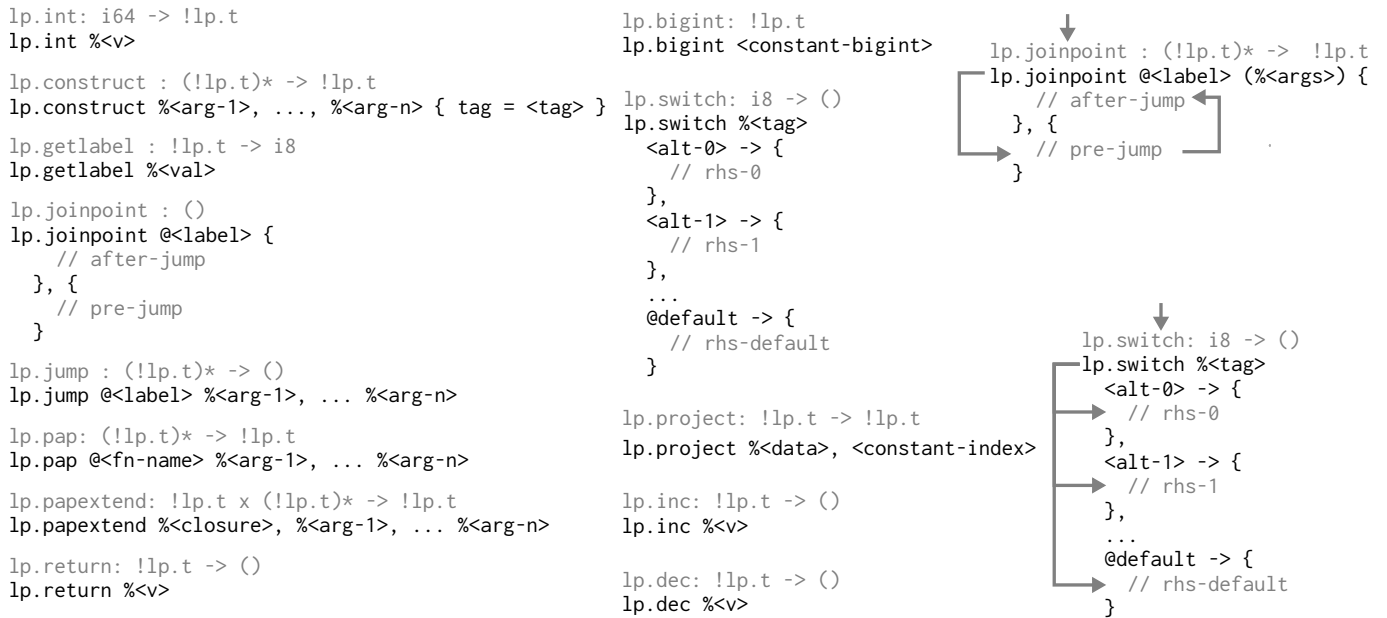


Fig. 2: `lp`'s operations, their syntax and their types. Variable length argument lists are denoted by  $(arg-ty)^*$ . `lp` expresses data constructors, pattern matching, closures, and join-points within an SSA framework. The control flow of `lp.joinpoint` and `lp.switch` are shown on the right. Control flow proceeds as per the arrows.




Fig. 3: LEAN4 compilation pipeline with our alternative `lp` backend, supported by `rgn` to encode functional constructs. We reuse the LEAN frontend, as well as the LEAN runtime to ensure compatibility with the runtime and FFI.

LEAN4's `λpure` IR is a minimal functional language with support for expressing data constructors, pattern matching, function applications, and closure creation. A lower level IR, known as `λrc`, extends `λpure` with operational concerns such as reference counting within the same IR. The current LEAN4 compiler compiles `λrc` down to C, followed by invoking a C compiler to link with the runtime and generate binaries. Our MLIR backend generates an MLIR dialect (`lp`) from `λrc`, and then continues compilation using the MLIR compilation pipeline (Figure 3).

We introduce the `lp` dialect (which stands for `λpure`, though we also support `λrc` reference counting instructions) within MLIR, with the following set of high-level operations (Figure 2):

- `lp.int`, `lp.bigint` to create machine integers and GMP-based big-integers.
- `lp.switch` to pattern-match on integers.
- `lp.construct` to create algebraic data types.
- `lp.getlabel` to extract the tag of an algebraic data type.
- `lp.project` to extract out fields from an algebraic data type.

- `lp.pap`, `lp.papextend`, to express partial function applications (closure creation).
- `lp.joinpoint`, `lp.jump`, for representing structured control flow using join points.
- `lp.inc`, `lp.dec` for reference counting.
- `lp.return` to return values from `lp` control flow.

We note that the dialect is *feature complete* to represent a functional programming language. It supports all the core constructs necessary, including closures and partial applications, data constructors and pattern matching on data constructors, as well as support for reference counting. This completeness is a strength of `lp`, since we are able to compile a realistic functional programming language (LEAN4) with our compiler pipeline. Like `λrc`, the `lp` dialect uses a single type, denoted by `!lp.t`, to represent values that live on the heap, i.e., boxed values. We also use standard integer types such as `i32` and `i8` as necessary to express interactions with machine integers. `lp` is *type erased*, as we have erased most typing information from the LEAN source program. We are left with just enough of the type information to generate code that is aware of types such as `int`, `float`, pointers, as well as integer bit widths.

### A. Integers & Switch Cases

LEAN supports arbitrarily sized natural numbers and integers with the `Nat` and `Int` types. LEAN also guarantees that small integers are represented by a machine word. To support this, `lp` introduces the `lp.bigint` operation to construct arbitrary sized integers and the `lp.int` operation to construct machine word sized integers. The LEAN4 compiler lowers `lp.bigint`

```

def intUsage : Nat -> Nat
| 42 => 43
| _ => 99999999

```

↓

```

func @intUsage(%arg: !lp.t) -> !lp.t {
  %c42 = lp.int 42 : !lp.t
  %eq42 = call @lean_nat_dec_eq(%arg, %c42) : i8
  lp.switch %eq42
  1 -> {
    %out = lp.int 43 : !lp.t
    lp.return %out : !lp.t
  },
  @default -> {
    %out = lp.bigint "99999999" : !lp.t
    lp.return %out : !lp.t
  }
}

```

Fig. 4: Program demonstrating integer manipulation with `lp.int`, `lp.bigint`, and `lp.switch`. We use `lp.switch` to pattern match on raw integer values. LEAN integers are type-erased to create a uniform representation for small and large integers. Integers are compared using the runtime call `@lean_nat_dec_eq`.

to runtime calls to create big-integers, while `lp.ints` are lowered to machine integers.

Note that the runtime does not distinguish between naturals and integers — they lower to the same signed machine level representation, and are only distinguished at the LEAN level by the dependently typed theory.

As LEAN integers can be both machine integers and big-integers, pattern matching on a LEAN integer must anticipate the scrutinee as being either a machine integer or a big-integer. `λrc` hides this complexity by staging the pattern match: (1) First, we check whether one LEAN integer is equal to another using a runtime call to `@lean_nat_dec_eq`. This function handles equality checking between machine-machine, machine-bigint, and bigint-bigint integers uniformly. (2) If the integers are equal, we execute the pattern match.

In more detail, the `@lean_nat_dec_eq` function receives two integers of type `!lp.t` as input, and returns either a 0 or a 1 of type `i8` to indicate whether the inputs are equal. We switch-case on the return value with the `lp.switch` operation which dispatches control to a matching switch arm. If no switch arm matches, control is dispatched to the `@default` arm (Figure 2).

For concreteness, consider the example program in Figure 4. To compile the first arm of the pattern match `42 =>`, we call `lean_nat_eq_dec` to check if the argument `%arg` equals 42. If this is true, we proceed to execute the right hand side of the pattern match and return 43. The other arm `_ =>` of the switch becomes the `@default` branch of the `lp.switch` which returns 99999999.

### B. Join Points

The naive lowering of case statements to switch operations can result in code duplication across the right hand sides of nested cases. We wish to avoid such duplication to ensure

### A Pseudocode example

```

def eval: Int -> Int -> Int -> Int
| 0, 2, _ -> 40 ①
| 0, _, 2 -> 50 ②
| _, _, _ -> 60 ③

```

### B Code Duplication

```

def eval(%x: i32, %y: i32, %z: i32) -> i32 {
  lp.switch %x
  0 -> {
    lp.switch %y
    2 -> {
      lp.return 40 ①
    },
    @default -> { // 0 2 _ -> 40
      lp.switch %z
      2 -> { // 0 _ 2 -> 50
        lp.return 50 ②
      },
      @default -> { // 0 _ _ -> 60
        lp.return 60 ③
      } // end switch %z
    } // end switch %y
  },
  @default -> { // _ _ _ -> 60
    lp.return 60 ③
  } // end switch %x
}

```

### C Deduplication via join-points

```

def eval(%x: i32, %y: i32, %z: i32) -> i32 {
  lp.joinpoint @jpc { ④ ←
  return 60
}, {
  lp.switch %x
  0 -> {
    lp.switch %y
    2 -> {
      lp.return 40
    },
    @default -> {
      lp.switch %z
      2 -> {
        lp.return 50
      },
      @default -> {
        lp.jump @jpc
      } // end switch %z
    } // end switch %y
  },
  @default -> {
    lp.jump @jpc
  } // end switch %x
} // end join point
}

```

Fig. 5: Complex pattern matching introduces code duplication across branches of control flow. This pattern matching is lowered using `lp.jump` to optimize the right hand side. The pattern matcher uses `lp.jump` to jump to the deduplicated right hand side. The arrows on the left indicate code duplication; the arrows on the right indicate deduplication via `lp.jump`.

efficient compilation of pattern matching. Consider a case statement that matches on two arguments (Figure 5). For the first two cases ① and ②, we expect that the first argument is 0, and the second or the third argument is 2. If this is not the case, we execute the default case ③.

As `lp.switch` can only match on a single integer at a time, we introduce two nested `lp.switch` statements when lowering this pattern match. The outer switch analyzes the

first argument and the inner one analyzes the second argument. If either case ① or case ② fails, we must execute the code for the default case ③. When directly lowering to nested switch operations, the expression that computes the result of the default case is duplicated into multiple branches of the switch operations (duplication indicated by arrows in Figure 5 A). In case of complex pattern matches with many variables, this duplication can be very costly.

To remedy such code duplication, as well as to increase the expressiveness of the IR to encode complex control flow, LEAN uses *join-points* [10], [11]. These join-points introduce the ability to create *labels* within the IR that control flow can jump to. Intuitively, a label is a *point* in the program, and jumping to a label *joins* control flow across the program at that point. The code duplication that we observed in Figure 5 A is eliminated via join-points (Figure 5 B). We create a common region ④ that performs the default action of returning 60. Code duplication is avoided by introducing two `lp.jumps` as indicated by the arrows. Conceptually, the combination of `lp.joinpoint` and `lp.jump` allows us to represent local, named closures within the IR, since the join-point region can refer to values that were defined before the join-point, as well as the arguments that are passed by `lp.jump`. The difference between the join-point closures and regular closures is that join-point closures are known to not escape, and all of their call sites are within the function body. This allows join-points to be efficiently lowered to jumps, unlike regular closures which require us to build heap objects that represent the closure. Together, `lp.joinpoint` and `lp.jump` provide a powerful tool to decrease code duplication incurred by complex pattern matching.

### C. Data Constructors & Pattern Matching

In this section, we describe `lp`'s support for constructing data types via data constructors, and destructuring data types via pattern matching. A data constructor is conceptually a tagged union. Thus, we construct data using `lp.construct`, which receives a tag and a sequence of arguments. The tag denotes which variant of a given data type is being built, and the argument list denotes the fields of the data constructor.

For example, if we wish to build linked lists (Figure 6), we would need two tags, one for the empty list (`Nil` / tag 0) and one for a cons-cell (`Cons` / tag 1). The function `@singleton` constructs a list with a single cons-cell by constructing a nil value and linking a cons-cell to the nil value.

To compute the length of the linked list, we pattern match on the tag of the data constructor. If the tag is 0 (the object is `Nil`), we return zero. Otherwise, the tag must be 1 (the object is `Cons`) and so we recursively call `length` on the rest of the list and add one. We extract the tag via `lp.getlabel`, pattern match on the tag using `lp.switch`, and extract the list pointer from the cons-cell using `lp.project`.

### D. Closures

LEAN supports two types of function application: (1) regular (saturated) function applications where the caller provides as

```

inductive List
| Nil | Cons (i: int) (l: List)

def singleton (n : Nat) : List := Cons n Nil

def length : List -> Nat
| Nil => 0 | Cons n l => 1 + length l

```

↓

```

func @singleton (%n: !lp.t) -> !lp.t {
  %nil = lp.construct { tag = 0 }
  %cons = lp.construct %n, %nil { tag = 1 }
  lp.return %cons
}

func @length(%arg0: !lp.t) -> !lp.t {
  %lbl = lp.getlabel %arg0
  lp.switch %lbl
  0 -> { // Nil
    %c0 = lp.int 0 : !lp.t
    lp.return %c0 : !lp.t
  },
  1 -> { // Cons
    %l = lp.project %arg0, 1
    %len = call @length(%l)
    %c1 = lp.int 1 : !lp.t
    %out = call @lean_nat_add(%c1, %len)
    lp.return %out : !lp.t
  }
}

```

Fig. 6: We encode a data constructor with `lp.construct`, which receives the tag and the values held by the constructor as arguments. We encode a case expression as a combination of `lp.getlabel`, `lp.project`, and `lp.switch`. `lp.getlabel` extracts the label of the data constructor that is switched on by `lp.switch`. The fields of the value that is being case'd upon are extracted using `lp.project`.

many arguments as the arity of the function, (2) partial function applications, where the caller provides fewer parameters than the function's arity. In the case of a regular function application, we call the function eagerly, and reuse MLIR's function call infrastructure to lower these. For partial function applications, we build a *closure* that holds onto the arguments that are supplied, and waits for further arguments to be provided before the call can be made. Once a closure receives all its arguments, the function held by the closure is invoked with all the arguments that are held by the closure.

Partial applications require some care. First, to build a closure, we introduce the `lp.pap` (partial application) operation, which partially applies a function to some arguments and builds a closure. To *extend* a closure with more arguments, we introduce the `lp.papextend` (partial application extend) operation, which takes as inputs a closure and further arguments to extend the closure. If these extra arguments *saturate* the closure (i.e., the closure now has all arguments), then the function stored in the closure is invoked with arguments stored in the closure. If the closure is not yet saturated the closure is *extended* by storing the newly passed arguments in the closure.

```

def k (x: Nat) (y: Nat): Nat := x
def k10: Nat -> Nat := k 10
def ap42 (f: Nat -> Nat -> Nat): Nat -> Nat := f 42 →
def k42: Nat -> Nat := ap42 k
    ↓
global @kslot : !lp.t
def @init() {
  %k = lp.pap @k
  store @kslot, %k // initialize @kslot
}
def @entrypoint(%argc: i64, %argv: i8*) {
  ...
}
def @main(%argc: i64, %argv: i8*) {
  call @init() // setup top-level closures
  call @entrypoint(%argc, %argv) // run program
}

func @k (%x: !lp.t, %y: !lp.t) -> !lp.t {
  return %y: !lp.t
}
func @k10(%x: !lp.t) -> !lp.t {
  %c10 = lp.int 10 : !lp.t
  %out = lp.pap @k, %c10
  return %out : !lp.t
}
func @ap42(%f: !lp.t) -> !lp.t {
  %c42 = lp.int 42 : !lp.t
  %out = lp.papextend %f, %c42
  return %out : !lp.t
}
func @k42(%x: !lp.t) -> !lp.t {
  %k = load @kslot : !lp.t
  %out = call @ap42(%k): !lp.t
  return %out : !lp.t
}

```

Fig. 7: Closures are built using `lp.pap`, and closures are extended using `lp.papextend`. We see that these are ubiquitous operations which almost always occur in the presence of higher order functions. A program written in the functional style will typically have many closures, so a compact and optimizable representation is important. `λrc` performs lambda lifting which creates top-level closures such as `@kslot` which must be initialized before the program is run. We first run `init()` which initializes closures, followed by `entrypoint` which is the entry point to the LEAN program.

Consider the example in Figure 7. The function `@k10` partially applies `k` to the value 10. This is done by calling `lp.pap` with arguments `@k` and 10. This creates a closure which waits for the other argument `y` to invoke `k(10, y)`.

The function `@ap42` takes an arbitrary closure and applies it to the value 42. This is achieved by using `lp.papextend` to extend the closure of `f` with the argument 42.

Finally, we draw attention to subtlety when invoking `ap42` with the argument `k`. `ap42`, which obeys LEAN’s calling convention, expects a *closure* as an argument. But what is the closure associated to `k`? We don’t have one, as `k` is a top-level, raw function, not a LEAN closure. This mismatch requires us to be able to create *empty closures* for top-level functions. This is performed by the runtime, where we have an initialization phase which creates top-level closures. In this case, the closure associated to `k` is initialized in the global variable `@kslot` by the function `@init`.

As for lowering, `lp.pap` and `lp.papextend` are lowered into runtime calls which manipulate closures. A closure is represented in-memory by a function pointer and a list of pointers to the arguments which are held by the closure.

### E. Tail Calls

`λrc` explicitly keeps track of calls which must be tail calls for memory consumption guarantees. To respect this, we generate an LLVM call annotated with a `musttail` attribute, which enforces that these calls must be tail calls. If a call that is annotated as a `musttail` cannot be guaranteed to be a tail call, then the LLVM optimizer provides an error. This is critical as tail-call-optimization is guaranteed by the LEAN language semantics, but is not guaranteed by the lowering to C as the C standard does not require guaranteed tail call elimination.

### F. Reference Counting

Since LEAN is a functional language, one does not explicitly refer to memory allocation or de-allocation. Hence, LEAN implements automatic memory management via reference counting. To incorporate reference counting for objects that are created on the heap, we expose the operations `lp.inc` and `lp.dec` to increment and decrement reference counts. These are lowered to corresponding LEAN runtime calls.

### G. Linking against the Lean Runtime

We link against the LEAN runtime which is written in C, by compiling to LLVM, and then linking using `llvm-link`. We perform this step as many performance critical runtime routines, including `refcounting`, are written in C, and are inlined by the `leanc` backend. Thus, for performance parity, we compile the C code down to LLVM and link the runtime, thereby providing LLVM visibility to LEAN’s runtime symbols.

## IV. THE RGN DIALECT

In this section, we introduce the `rgn` dialect to represent and optimize control flow within `λrc`. We describe the semantics of the `rgn` dialect, the lowering from `lp`’s control-flow constructs (`lp.switch`, `lp.joinpoint`, `lp.jump`) to `rgn` constructs, and the adaptation of classical SSA optimizations to `rgn`. The `rgn` dialect has two instructions:

- The `rgn.val` instruction creates an SSA value which names a region definition. These named values declare regions prior to being called. This is conceptually a *continuation*, where the region represents a computation that is to be performed when invoked.
- The `rgn.run` terminator instruction transfers control flow to a region with the supplied arguments. This is conceptually *invoking a continuation*. We branch to the

region that is to be executed and continue execution from this region.

We allow `rgn.val` values to be passed as operands to MLIR’s `select` and `switch` instructions. These instructions allow us to pick a value based on a boolean or an integer respectively. We use this to express choosing which region is to be executed for control flow. We do not allow `rgn.val` operations to interact with other operations; in particular, they may not be passed to other functions as arguments and may not be returned. This ensures that all uses of `rgn.val` can be statically analyzed by our compiler, while still allowing us to leverage the analyses and optimizations provided for `select` and `switch`.

### A. Lowering `lp` to `rgn`

#### A 2-Way Switch

```
lp.switch %<tag>
<alt-0> -> {
  // rhs-0
},
<alt-1> -> {
  // rhs-1
}
```



```
%r0 = rgn.val {
  // rhs-0
}
%r1 = rgn.val {
  // rhs-1
}
%r = select %<tag>,
           %r0, %r1
rgn.run %r
```

#### B N-Way Switch

```
lp.switch %<tag>
<alt-0> -> {
  // rhs-0
},
<alt-1> -> {
  // rhs-1
},
...
@default -> {
  // rhs-default
}
```



```
%r0 = rgn.val {
  // rhs-0
}
%r1 = rgn.val {
  // rhs-1
}
...
%rd = rgn.val {
  // rhs-default
}
%r = switch %tag
         <alt-0> -> %r0,
         <alt-1> -> %r1,
         ...
         default -> %rd
rgn.run %r
```

#### C Join Points

```
lp.joinpoint @<lbl> {
  // after-jump
}, {
  // pre-jump
  ...
  lp.jump @<lbl> (%arg)*
}
```



```
%lbl = rgn.val {
  // after-jump
}
// pre-jump
...
rgn.run %lbl (%arg)*
```

Fig. 8: Lowering the control flow of `lp` to `rgn`. `lp.switch` with two cases is lowered via `select` (A). `lp.switch` with many cases is lowered via `switch` (B). `lp.joinpoint` is lowered to a combination of `rgn.val` and `rgn.run` (C).

The lowering from `lp` to `rgn` is straightforward (Figure 8). `lp` has two control flow constructs, `lp.switch` & `lp.joinpoint`. `lp.switch` is lowered by converting every right hand side of a pattern match to a `rgn.val`, then selecting the correct right hand side that needs to be executed, and finally executing the selected right hand side via `rgn.run`. We lower the selection to either a `select` or a `switch` depending on the number of cases. We lower `lp.joinpoint` by converting the jump target to a `rgn.val`, the `lp.jump`

to a `rgn.run`, replacing the joinpoint by the region that is to be executed before the jump.

### B. Optimization Passes for `rgn`

In this section, we adapt classical optimization passes to `rgn`’s region-values. We explore optimizations that are made possible by these adapted transformations.

1) *Dead Region Elimination*: Dead code elimination requires no changes to work with region values. If a region value is never referenced, then it is never executed. It is thus dead and can safely be removed.

<p><b>A</b></p> <pre>out = case True of   True -&gt; 3   False -&gt; 5</pre> <p><b>C</b></p> <pre>%t = constant true %x = rgn.val { return 3 } %y = rgn.val { return 5 } rgn.run %x</pre>	<p><b>B</b></p> <pre>%x = rgn.val { return 3 } %y = rgn.val { return 5 } %z = select true, %x, %y rgn.run %z</pre> <p><b>D</b></p> <pre>return 3</pre>
---	--

Figure A above is the original program. We show the naive translation to `rgn` in B. The optimizations for `select` on a constant value `true` kick in and replace `%z` by `%x` in C. Finally, the running of a known region is replaced by the constant value that corresponds to the region in D.

2) *Global Region Numbering*: We introduce global value numbering for regions.<sup>2</sup> For *straight line regions*, the value number of the region is defined as a rolling hash of the value numbers of all instructions within the region. Two regions are defined to have the same value number if and only if the sequence of instructions in the two regions have the same value numbers in identical order. The restriction that a region has a single basic block is not as restrictive as it first appears, since languages with high level control flow can always be lowered to regions with a single basic block, with control flow expressed via `rgn` and `select`. Global value numbering of regions identifies redundant computations across branches of control flow. This allows us to fold away equivalent computations.

<p><b>A</b></p> <pre>out = case b of   True -&gt; 7   False -&gt; 7</pre> <p><b>C</b></p> <pre>%b = &lt;external&gt; %w = rgn.val { return 7 } %z = select %b, %w, %w rgn.run %z</pre>	<p><b>B</b></p> <pre>%b = &lt;external&gt; %x = rgn.val { return 7 } %y = rgn.val { return 7 } %z = select %b, %x, %y rgn.run %z</pre> <p><b>D</b></p> <pre>%b = &lt;external&gt; %w = rgn.val { return 7 } rgn.run %w</pre>
--	--

Figure A above is the original program. We show the naive translation to `rgn` in B, where the value `%b` is external to the current scope being analyzed. The Region-CSE algorithm

<sup>2</sup>MLIR does not perform global value numbering as it is unclear how to define value numbers for instructions with regions, as in general arbitrary regions do not yet have a prescribed semantics.

fuses  $\%x$  and  $\%y$  into one region, relabelled  $\%w$  in **C**. Finally, the select of two equal values ( $\%w$ ,  $\%w$ ) on both branches is folded away in **D**.

### C. Lowering *rgn* to *std*

In this section, we describe how to lower *rgn* to a traditional SSA-based IR without regions. For concreteness, we describe a lowering to MLIR’s *std* dialect. We anticipate no changes necessary to lower to another purely SSA based IR such as LLVM.

Since the semantics of *rgn* is given entirely by adding extra structure to flat CFGs, *rgn* can be lowered by forgetting this extra structure. The lowering is driven entirely by *rgn.run*. First, we lower *rgn.runs* by matching on the argument — (1) A *rgn.run* of a known *rgn.val* is compiled to a branch of the region that is run, (2) A *rgn.run* of a switch (or select) is compiled to a jump-table. Finally, dead *rgn.val* instructions are entirely dropped.

## V. EVALUATION

We evaluate our work by analyzing the correctness, performance, and memory usage of the proposed SSA-based LEAN backend. Our evaluation was run on an Intel Xeon CPU @ 2.20GHz with 126GB memory. We first discuss the correctness of our compiler. We then consider the end-to-end performance, by comparing runtimes of programs compiled by our compiler versus the baseline Lean compiler on LEAN’s benchmark suite, and then perform an evaluation of the *rgn* optimizations on the benchmark suite. Finally, we have a holistic discussion about the costs and benefits of our new backend in comparison to the current C backend of LEAN4.

### A. Correctness

We test our compiler for correctness against the LEAN test suite, which consists of 648 test cases, out of which we pass 648 (100%) of tests. This shows that our compiler correctly implements the semantics of  $\lambda_{rc}$  and interfaces properly with the LEAN runtime system. This ensures that our evaluation is representative of functional programming workloads.

### B. Performance Analysis

We characterize the performance of our backend in comparison to LEAN’s default backend (commit hash `be4cf60`) using LEAN’s benchmark suite. The programs in the LEAN benchmark suite represent workloads commonly encountered by functional programming languages:

- `binarytrees` and `binarytrees-int` implement a purely functional binary tree lookup, insert, and delete benchmark.
- `const_fold` implements constant folding on an expression based language.
- `deriv` benchmarks derivative computations on expression trees.
- `filter` implements filtering values from a linked list based on a predicate.
- `qsort` implements real in-place quicksort using LEAN’s arrays.

- `rbmap_checkpoint` implements red-black tree insertion and lookup.
- `unionfind` implements Tarjan’s union-find algorithm.




Fig. 9: Speedup of our runtimes in comparison to LEAN4’s existing C backend. The geomean speedup over the baseline LEAN4 compiler across all benchmarks is  $1.09\times$ . Thus, we achieve performance parity with the LEAN4 backend.

The reported timing results (Figure 9) are gathered as the geometric mean of runtimes over ten runs. We report a geomean speedup of  $1.09\times$  with the baseline LEAN compiler, which validates our claim that we achieve performance parity with the LEAN4 compiler.




Fig. 10: Speedup of *rgn* dialect optimizations over  $\lambda_{rc}$  (red) and speedup of no optimizations over  $\lambda_{rc}$  (gray). The numbers over the bar are the speedups of *rgn* over  $\lambda_{rc}$ . The geomean speedup of *rgn* over  $\lambda_{rc}$  across all benchmarks is  $1.0\times$ . Thus, our simplification pipeline achieves performance parity with  $\lambda_{rc}$ .

Finally, we compare the effectiveness of the *rgn* dialect to that of the  $\lambda_{rc}$  simplifier (Figure 10) by running three versions of the pipeline: (a) a baseline of our MLIR pipeline which receives optimized code from the  $\lambda_{rc}$  simplifier, which we compare against (b) unoptimized  $\lambda_{rc}$  code which is then optimized by *rgn* (we disable LEAN’s `simplcase` pass which performs *rgn* style switch simplification), as well as (c) unoptimized  $\lambda_{rc}$  code which is left unoptimized by *rgn* before passing to LLVM. We find that the performance is almost identical across the three variants. We conjecture that LLVM’s sophisticated control flow analysis is capable of optimizing away the unoptimized IR we pass it.

In summary, we show that our SSA-based optimization pipeline matches the hand-written LEAN optimization pipeline, while reasoning about functional constructs purely via SSA. The *rgn* dialect allows us to cleanly express  $\lambda_{rc}$ ’s semantics with SSA, thereby making SSA+regions an attractive choice to represent functional programming languages in traditional imperative compiler tooling.



### C. Qualitative Analysis

In this section, we compare and contrast our backend to LEAN’s backend from the perspective of a *compiler engineer*. We first survey a list of salient differences between LEAN’s  $\lambda_{rc}$  backend and our MLIR-based backend (Figure 11).

Feature	$\lambda_{rc} + C$	lp + rgn
Backend	C	MLIR
Vectorization	No	affine, linalg FileCheck,
Testing harness	makefile	llvm-lit
Constant folding	Hand-written	MLIR rewriter
CSE	Hand-written	MLIR builtin
DCE	Hand-written	MLIR builtin
Inliner	Hand-written	MLIR builtin
Test minimization	None	mlir-reduce
Debug information	None	Possible
IDE support	None	LSP
Tail call optimization	Heuristic	Guaranteed

Fig. 11: Ecosystem differences between our MLIR-based backend and the current LEAN4 tooling. Note that MLIR provides us access to a rich ecosystem of tooling for compiler development.

Notice the availability of rich tooling due to our use of the MLIR compiler framework — we are able to reuse existing optimization infrastructure for constant folding, common subexpression elimination (extended by *rgn*), dead code elimination, and inlining. Similarly, we use LLVM’s test harness infrastructure for parallel test running and failure reporting. The  $\lambda_{rc}$  pipeline generates C code from which crashes and miscompiles are difficult to debug as there is no association between the generated C program and the source LEAN program. One can potentially teach the LEAN frontend to preserve debug information to be passed to MLIR. The MLIR compiler framework makes retaining debug information a top priority that enables the compiler engineer to correlate miscompiles to LEAN program source locations accurately. Finally, as indicated before, LEAN requires guaranteed tail call elimination which the current C based backend cannot. Our MLIR-based compiler guarantees this by using LLVM’s *musttail* annotations.

In summary, our compiler lays the foundation for LEAN to access an extensive suite of analyses and optimizations. It also provides the compiler engineer tooling for debugging and unit testing. Lastly, it allows us to provide deterministic tail call elimination, which is necessary for the semantics of  $\lambda_{rc}$ .

## VI. RELATED WORK

In this section, we survey related work on functional intermediate representations and uses of SSA in compilers for functional languages.

The Glasgow Haskell Compiler (GHC) [12] is an optimizing compiler for Haskell, a lazy language with a focus on having an expressive type system and the ability to write optimized

programs. GHC uses an intermediate representation known as Core, which is a strictly typed encoding of  $F_{\omega}$  [13]. Most optimizations to do with exposing computation happen at this layer. After this stage, the encoding is lowered to STG [14], a lower-level intermediate representation in administrative normal form [15]. Next, STG is lowered to C- [16], which is a target-independent assembly language that supports garbage collection. From this stage, GHC can either emit assembly or generate LLVM. The impedance mismatch is synthesized during the translation from STG to C-. C- is a traditional assembly IR that is well-suited to borrow traditional optimizations from imperative compilers. Unfortunately, it is precisely at this step that GHC chooses to simultaneously lower the encodings of laziness and algebraic data types, making the assembly hard to optimize.

The Intel Haskell research compiler [17] is a whole-program optimizing compiler for Haskell which focuses on vectorization and other program transformations for performance. This compiler starts from GHC Core, and then translates to a lazy ANF (Administrative Normal Form) [15]. The compiler performs demand analysis and abstract simplification on ANF. The demand analysis is performed using traditional abstract interpretation techniques. Next, this demand information is used to interpret the program and perform abstract simplification. It then compiles to an intermediate representation called MIL, which is a loosely typed CFG based, SSA-like intermediate representation. They represent laziness as heap values, and manipulate the heap. Thus, their representation and analysis of lazy values uses memory semantics instead of value semantics. MIL also does not have a notion of nested regions. Therefore, MIL extends the traditional control flow controls with finer-grained information, called as *cut* and *interproc*. Finally, MIL is vectorized, and then lowered to Intel’s low-level IR, Pillar, which mirrors LLVM [1]. Finally, assembly is generated from Pillar.

GRIN [18] is an alternative monadic intermediate representation for lazy and strict functional programming languages which explicitly represents heap manipulation and case analysis. Due to the monadic encoding, it has pointer semantics, not value semantics. GRIN therefore chooses to run a sophisticated whole-program points-to analysis to resolve these pointers for optimization. GRIN is well-suited for the whole-program paradigm, while we focus on optimizing using local, per-module information without incurring the penalty of a costly global analysis.

Thorin [19] is a higher-order, functional IR based on continuation-passing style. Thorin chooses to not use explicit nesting, and uses a dependency graph instead. This has the advantage of providing greater flexibility during compilation. In contrast, we use explicit nesting based on regions, which is easier to analyze and adapt into an SSA based framework.

The MLTon compiler [20] is a whole program optimizing compiler for Standard ML. MLTon uses a variant of SSA to encode algebraic data types and case analysis. They perform whole program compilation, and use aggressive whole program analyses reminiscent of GRIN to analyze and eliminate

overhead. Furthermore, their adaptations to SSA do not involve regions. They choose to encode special terminator operations that represent control flow by case analysis.

## VII. CONCLUSION

In this work, we have described `lp`, which implements a new backend for the LEAN4 compiler within MLIR. To optimize a functional language within an SSA compiler framework, we introduce `rgn`, an intermediate representation that is designed for region analysis and optimization. We adapt classical functional language transformations to the SSA setting by using and extending SSA algorithms to operate on `rgn`. Finally, we implement our suggestions and demonstrate that our new LEAN4 backend based on `lp` and `rgn` passes all regression tests and achieves performance parity with the existing LEAN4 compiler. This foundation is a force multiplier, as we can now express the LEAN4 semantics in a compiler framework that has been designed to be analyzed, optimized, and offers mature support for performance and regression testing. We envision `lp` and `rgn` together acting as the bedrock for SSA-based optimizing compilers for many more functional programming languages.

## APPENDIX

### A. Abstract

The artifact’s goal is to show how regions+SSA allows us to create an MLIR-based backend for LEAN which achieves performance parity with the LEAN4 compiler. The artifact consists of a docker container with accompanying scripts to replicate figure 9, 10. The docker container is the only piece needed to run all the experiments. Scripts to generate the figures and the table come with the docker container.

### B. Artifact Check-List (Meta-Information)

- **Program:** A custom LEAN4 backend based on the MLIR compiler toolchain, along with LEAN4’s test suite for testing, and LEAN4’s benchmark suite for performance analysis.
- **Compilation:** A C++ compatible compiler to bootstrap LLVM/MLIR as well as LEAN4.
- **Run-time environment:** Any operating system that supports Docker.
- **Hardware:** Any x86 machine.
- **Output:** PDF files replicating Figure 9 and Figure 10, and a successful run of the entire LEAN4 test suite.
- **How much disk space required (approximately)?:** 10GB.
- **How much time is needed to prepare workflow (approximately)?:** 2 hours.
- **How much time is needed to complete experiments (approximately)?:** 1 hour.
- **Publicly available?:** Yes.

### C. Description

1) *How Delivered:* The artifact is delivered as a Docker container and is available at <http://doi.org/10.5281/zenodo.5786074>.

2) *Hardware Dependencies:* None.

3) *Software dependencies:* The docker image has dependences needed to compile MLIR, Lean4, and our compiler tooling to run the test suite.

### D. Experiment Workflow

Access the docker image `cgo22.docker` from (<http://doi.org/10.5281/zenodo.5786074>), then run:

```
$ docker load -i cgo22.docker
$ docker run -it siddudruid/cgo21-v4
$ su nonroot # switch to non-root user
$ export PATH=/code/llvm-project/build/bin:$PATH
$ export PATH=/code/lean4/build/release/stage1/bin:$PATH
$ export PATH=/code/lz/build/bin:$PATH
$ export LEANLIB=/code/lean4/build/release/stage1/lib
$ export LD_LIBRARY_PATH=$LEANLIB:$LD_LIBRARY_PATH
$ cd /code/lean4/build/release && \
    make -j4 test
$ cd /code/lz/test/lambdapure/compile/bench && \
    ./speedup-time.py --data --plot --nrns 10
$ cd /code/lz/test/lambdapure/compile/bench && \
    ./speedup-rgn-time.py --data --plot --nrns 10
```

Upon running `make -j4 test`, the test output is printed to stdout. The scripts `speedup-time.py` and `speedup-rgn-time.py`, produce PDFs `speedup-time.pdf` and `speedup-rgn-time.pdf` in the directory `/code/lz/test/lambdapure/compile/bench/`:

```
/code/lz/test/lambdapure/compile/bench/speedup-time.pdf
/code/lz/test/lambdapure/compile/bench/speedup-rgn-time.pdf
```

To open the pdf file, keep the container running, and in another shell instance, use the `docker cp` command to copy files from within the container out to the host:

```
$ docker container ls # find ID
$ docker cp <CONTAINERID>:<PATH/INSIDE/CONTAINER> \
    <PATH/OUTSIDE/CONTAINER>
```

For more about `docker cp`, please see: (<https://docs.docker.com/engine/reference/commandline/cp/>)

### E. Evaluation and Expected Result

On running the test suite with:

```
$ cd /code/lean4/build/release && make -j test
```

We find that the output is:

```
100% tests passed, 0 tests failed out of 648
```

See that we pass all tests. The test script that is run can be found at `/code/lean4/test/common.sh` which can be seen to invoke our compiler pipeline.

To generate performance plots, we run:

```
$ cd /code/lz/test/lambdapure/compile/bench && \
    ./speedup-time.py --data --plot --nrns 10
$ cd /code/lz/test/lambdapure/compile/bench && \
    ./speedup-rgn-time.py --data --plot --nrns 10
```

For `speedup-time.pdf`, we expect a geometric speedup of 1.0x with the baseline `lean4`, and for `speedup-rgn-time.pdf`, we expect a geometric speedup of 1.0x between region optimizations and no region optimizations.

### F. Full Workflow Example

```
# Grab docker image from http://doi.org/10.5281/zenodo.5786074
$ curl https://zenodo.org/record/5786074/files/cgo22.docker?download=1 > cgo22.docker
$ docker load -i cgo22.docker
$ docker run -it siddudruid/cgo21-v4
$ su nonroot # switch to non-root user
$ export PATH=/code/llvm-project/build/bin:$PATH
$ export PATH=/code/lean4/build/release/stage1/bin:$PATH
$ export PATH=/code/lz/build/bin:$PATH
$ export LEANLIB=/code/lean4/build/release/stage1/lib
$ export LD_LIBRARY_PATH=$LEANLIB:$LD_LIBRARY_PATH
$ cd /code/lean4/build/release/ && make -j test # NO test failures
100% tests passed, 0 tests failed out of 648

Total Test time (real) = 399.28 sec
[100%] Built target test

$ cd /code/lz/test/lambdapure/compile/bench && \
    ./speedup-time.py --data --plot --nrns 10

$ cd /code/lz/test/lambdapure/compile/bench && \
    ./speedup-rgn-time.py --data --plot --nrns 10
```

## REFERENCES

- [1] C. Lattner and V. Adve, “Llvm: A compilation framework for lifelong program analysis & transformation,” in *International Symposium on Code Generation and Optimization, 2004. CGO 2004*. IEEE, 2004, pp. 75–86.
- [2] D. Novillo, “Design and implementation of tree ssa,” 2004.
- [3] R. A. Kelsey, “A correspondence between continuation passing style and static single assignment form,” *ACM SIGPLAN Notices*, vol. 30, no. 3, pp. 13–22, 1995.
- [4] L. de Moura and S. Ullrich, “The lean 4 theorem prover and programming language,” in *International Conference on Automated Deduction*. Springer, 2021, pp. 625–635.
- [5] F. Rastello, *SSA-based Compiler Design*. Springer, 2010.
- [6] R. M. Stallman, “Gnu compiler collection internals,” *Free Software Foundation*, 2002.
- [7] C. Lattner, J. Pienaar, M. Amini, U. Bondhugula, R. Riddle, A. Cohen, T. Shpeisman, A. Davis, N. Vasilache, and O. Zinenko, “Mlir: A compiler infrastructure for the end of moore’s law,” *arXiv preprint arXiv:2002.11054*, 2020.
- [8] L. de Moura, S. Kong, J. Avigad, F. Van Doorn, and J. von Raumer, “The lean theorem prover (system description),” in *International Conference on Automated Deduction*. Springer, 2015, pp. 378–388.
- [9] B. C. Pierce, *Advanced topics in types and programming languages*. MIT press, 2005.
- [10] P. Downen, L. Maurer, Z. M. Ariola, and S. Peyton Jones, “Sequent calculus as a compiler intermediate language,” in *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming*, 2016, pp. 74–88.
- [11] L. Maurer, P. Downen, Z. M. Ariola, and S. Peyton Jones, “Compiling without continuations,” in *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2017, pp. 482–494.
- [12] S. P. Jones, C. Hall, K. Hammond, W. Partain, and P. Wadler, “The glasgow haskell compiler: a technical overview,” in *Proc. UK Joint Framework for Information Technology (JFIT) Technical Conference*, vol. 93, 1993.
- [13] B. C. Pierce and C. Benjamin, *Types and programming languages*. MIT press, 2002.
- [14] S. L. P. Jones, “Implementing lazy functional languages on stock hardware: the spineless tagless g-machine version 2.5,” *Journal of functional programming*, vol. 2, no. 2, pp. 127–202, 1992.
- [15] C. Flanagan, A. Sabry, B. F. Duba, and M. Felleisen, “The essence of compiling with continuations,” in *Proceedings of the ACM SIGPLAN 1993 conference on Programming language design and implementation*, 1993, pp. 237–247.
- [16] S. P. Jones, N. Ramsey, and F. Reig, “C—: A portable assembly language that supports garbage collection,” in *International Conference on Principles and Practice of Declarative Programming*. Springer, 1999, pp. 1–28.
- [17] H. Liu, N. Glew, L. Petersen, and T. A. Anderson, “The intel labs haskell research compiler,” in *Proceedings of the 2013 ACM SIGPLAN symposium on Haskell*, 2013, pp. 105–116.
- [18] U. Boquist and T. Johnsson, “The grin project: A highly optimising back end for lazy functional languages,” in *Symposium on Implementation and Application of Functional Languages*. Springer, 1996, pp. 58–84.
- [19] R. Leiba, M. Köster, and S. Hack, “A graph-based higher-order intermediate representation,” *2015 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, pp. 202–212, 2015.
- [20] S. Weeks, “Whole-program compilation in mlton,” *ML*, vol. 6, pp. 1–1, 2006.