

Verifying Peephole Rewriting In SSA Compiler IRs

Siddharth Bhat ✉ 

Cambridge University, United Kingdom

Alex Keizer ✉ 

Cambridge University, United Kingdom

Chris Hughes ✉

University of Edinburgh, United Kingdom

Andrés Goens ✉ 

University of Amsterdam, Netherlands

Tobias Grosser ✉ 

Cambridge University, United Kingdom

Abstract

There is an increasing need for domain-specific reasoning in modern compilers. This has fueled the use of tailored intermediate representations (IRs) based on static single assignment (SSA), like in the MLIR compiler framework. Interactive theorem provers (ITPs) provide strong guarantees for the end-to-end verification of compilers (e.g., CompCert). However, modern compilers and their IRs evolve at a rate that makes proof engineering alongside them prohibitively expensive. Nevertheless, well-scoped push-button automated verification tools such as the Alive peephole verifier for LLVM-IR gained recognition in domains where SMT solvers offer efficient (semi) decision procedures. In this paper, we aim to combine the convenience of automation with the versatility of ITPs for verifying peephole rewrites across domain-specific IRs. We formalize a core calculus for SSA-based IRs that is generic over the IR and covers so-called regions (nested scoping used by many domain-specific IRs in the MLIR ecosystem). Our mechanization in the Lean proof assistant provides a user-friendly frontend for translating MLIR syntax into our calculus. We provide scaffolding for defining and verifying peephole rewrites, offering tactics to eliminate the abstraction overhead of our SSA calculus. We prove correctness theorems about peephole rewriting, as well as two classical program transformations. To evaluate our framework, we consider three use cases from the MLIR ecosystem that cover different levels of abstractions: (1) bitvector rewrites from LLVM, (2) structured control flow, and (3) fully homomorphic encryption. We envision that our mechanization provides a foundation for formally verified rewrites on new domain-specific IRs.

2012 ACM Subject Classification Software and its engineering → Compilers; Software and its engineering → Semantics; Computing methodologies → Theorem proving algorithms; Theory of computation → Rewrite systems

Keywords and phrases compilers, semantics, mechanization, MLIR, SSA, regions, peephole rewrites

Digital Object Identifier 10.4230/LIPIcs.ITP.2024.23

Supplementary Material *Software*: <https://zenodo.org/records/11519739>

Funding This project has received funding from the European Union's Horizon EUROPE research and innovation program under grant agreement no. 101070374 (CONVOLVE).

Acknowledgements We thank Sébastien Michelland and Sebastian Ullrich for their early help in this project and feedback, as well as Anton Lorenzen for his helpful feedback.

1 Introduction

Static single assignment (SSA) [30] is the workhorse of modern compilers such as LLVM [16]. A key optimization that is enabled by SSA is to syntactically match a program pattern, and

44 replace the matched pattern with an optimized, semantically-equivalent program fragment.
 45 Despite their simplicity, these local “peephole optimizations” [22] remain important in
 46 compiler development. A quick overview of the program transformation libraries of
 47 LLVM shows that more than 10% of its code¹ is dedicated to LLVM’s peephole optimizer
 48 InstCombine, which is beyond the size of the LLVM loop optimizer. Despite the large size
 49 and scope of the LLVM project, Alive [20] is regularly referenced in LLVM commits. This
 50 is evidence that SMT-based, low-effort tooling for peephole rewrites can enable the use of
 51 verification in day-to-day compiler development.

52 Peephole rewriting has been formalized in its simpler, classical form of straight-line
 53 assembly code [24]. To our knowledge, peephole rewriting along def-use chains [1] has not
 54 yet been formalized. As an example, consider the rewrite $(y = x + 1; z = y - 1) \mapsto (z = x)$.
 55 This pattern does *not* match the program $(y = x + 1; \mathbf{p} = \mathbf{y}; z = y - 1)$ in straight-line
 56 rewriting, due to the interleaved instruction $p = y$. On the other hand, by concentrating
 57 on the dataflow, we rewrite any subprogram of the form $(y = x + 1; \circ; z = y - 1)$ to
 58 $(y = x + 1; \circ; z = x)$, regardless of what fills the hole \circ . This is known as rewriting on
 59 the “def-use” chain, where the pattern matching is extended to semantic subexpressions in
 60 the program. Observe that the fact that addition and subtraction are pure, and that SSA
 61 does not allow mutating the value of y is critical for the correctness of this optimization.
 62 Domain-specific peephole rewrites within the MLIR compiler framework [17] rely on purity
 63 and referential transparency to enable this class of optimizations.

64 MLIR is a compiler framework for multi-level, domain-specific compiler IRs. It is
 65 widely used in the context of machine learning [34], quantum computing [28], and even as an
 66 alternative compiler for Lean [6], among other domains. MLIR lowers the cost of instantiating
 67 domain-specific IRs and encourages local transformations that exploit the value semantics
 68 (i.e., referential transparency) of specialized high-level IRs over global reasoning at a lower
 69 abstraction level. MLIR also introduces the concept of regions, which can model control flow
 70 and other structured IR operations as nested IRs that replace complex unstructured control.
 71 Existing formalizations of SSA do not cover domain-specific SSA-based IRs or regions.

72 In this paper, we propose a framework that is aimed at prototyping and verifying peephole
 73 optimizations for domain-specific SSA-based IRs. We formalize a core calculus for SSA-based
 74 IRs and mechanize it in the Lean [8] proof assistant to enable verification of peephole rewriting
 75 over SSA IRs based on value semantics with regions. Our framework is deliberately built to
 76 be interoperable with MLIR. This aims to streamline the verification of peephole rewrites for
 77 MLIR. Concretely, we contribute:

- 78 ■ A formalization of SSA with regions parametrized over a user-defined IR X and its
 79 mechanization in our framework² `LeanMLIR(X)` that exploits denotational-style value
 80 semantics for optimizing along the SSA use-def chain of an MLIR-style IR (Sections 2, 3)
- 81 ■ Evidence that our formalization of SSA allows for effective meta-theoretic reasoning:
 - 82 ■ A verified peephole rewriter, for which we prove that lifting a peephole rewrite to a
 83 rewrite on the entire program preserves semantics (Section 4.1)
 - 84 ■ Two verified implementations of generic SSA-based optimizations: dead code elimina-
 85 tion and common subexpression elimination (Section 4.2)
 - 86 ■ Proof automation for eliminating the abstraction overhead of our SSA calculus and
 87 exposing clean mathematical proof obligations for each rewrite (Section 4.3)
- 88 ■ An extension of our pure optimizations in a context with side effects (Section 5)

¹ Non-blank and non-comment lines of `.cpp` files in `llvm/lib/Transforms` on commit `f4f1cf6c3`.

² Our framework is open-source and available at <https://github.com/opencompl/ssa>.

```

inductive Ty
| r
| nat

inductive Op
| arith_const (x : Nat) -- with compile-time data `x`
| monomial -- build equivalence class of monomial
| add -- add op.

```

(a) User definitions for `QuotRing` in our framework. `Op` has three constructors, `add`, `monomial` and `(const x)`, for `x` an element of \mathbb{N} , matching the three operations of the IR. `Ty` has two constructors, `r` and `nat`.

```

instance : OpSignature Op Ty where signature
| .arith_const _ => { sig := [], outTy := .nat } -- takes no args, returns an `r`.
| .add          => { sig := [.r, .r], outTy := .r } -- takes two `r`s, returns an `r`.
| .monomial    => { sig := [.nat, .nat], outTy := .r } -- takes two `nat`s

```

(b) User-defined signatures of each `QuotRing` operation.

```

noncomputable def generator : (ZMod q) [X] := X^(2^n) + 1
abbrev R := (ZMod q) [X] / (span {generator q n})

```

```

instance : TyDenote Ty where
toType
| .r => R -- the denotation of `r` is an element of the ring `R`
| .nat => Nat

```

```

instance : OpDenote Op Ty where
denote
| .arith_const (x : Nat), _, _ => x -- Denotation of `(arith_const x)` is `x`
| .add, [(x : R), (y : R)]_h, _ => x + y
| .monomial, [(c : Nat), (i : Nat)]_h, _ =>
  Quotient.mk (span {generator q n}) (monomial i c)

```

(c) User-defined semantics of `QuotRing`. The `instance` syntax is used to define a typeclass instance, by specifying the corresponding members, which in this case are the denotation functions. The `noncomputable` annotation in Lean tells the compiler not to generate executable code for this function, since `mathlib` uses a noncomputable definition for quotients of polynomial rings. Note that our framework ensures that values are well-typed according to `OpSignature` and `TyDenote`.

■ **Figure 1** User definitions for `QuotRing`, which declares the operations and types of the IR, the type signatures of the operations, and the denotations of the types and operations into Lean types.

89 ■ Syntax, semantics, and local rewrites for three MLIR-based IRs: (1) arithmetic over
90 bitvectors, (2) structured control flow, and (3) fully homomorphic encryption (Section 6)

91 2 Motivation: Verifying Optimizations for High-Level IRs

92 Effective domain-specific optimizations are almost impossible when reasoning on traditional
93 LLVM-style compiler IRs. These offer a “universal” low-level abstraction, originally designed
94 to represent C-style imperative code. Such LLVM-style IRs are built around the concepts
95 of load/store/arithmetic/branching, which is ideal when optimizing at the level of scalar
96 arithmetic, instruction scheduling, and certain kinds of loop optimizations. However, this
97 level of abstraction is unsuitable for reasoning about high-level mathematical abstractions.

98 Consider a compiler for Fully Homomorphic Encryption (FHE) [9], a cryptographic
99 technique that uses algebraic structures to allow an untrusted third party to do computation
100 on encrypted data. In such a compiler, we might have a rewrite like $(a + X^{2^n} + 1 \mapsto a)$,

101 which is a simple identity on the corresponding quotient ring.³ Expressed in LLVM, the
 102 computation of this simple operation consists of multiple basic blocks forming a loop, each
 103 containing memory loads, pointer arithmetic, scalar operations, and branches. As a result,
 104 the algebraic structure is completely lost and exploiting simple algebraic identities turns into
 105 a heroic effort of reasoning about side effects and stateful program behavior. State-of-the-art
 106 compilers for FHE consequently use domain-specific IRs (often expressed with MLIR [35, 26])
 107 when generating optimized code for FHE, where algebraic optimizations can take place at an
 108 FHE-specific IR that has value-semantics (e.g., is referentially transparent) and is overall
 109 closer to the mathematical structure of the problem.

110 2.1 Defining LeanMLIR(QuotRing): Syntax and Semantics

111 As an example, we model an IR aimed at FHE that manipulates objects in the algebraic
 112 structure $R \equiv (\mathbb{Z}/q\mathbb{Z})[X]/(X^{2^n} + 1)$. To model it, we instantiate an IR `LeanMLIR(QuotRing)`
 113 in our framework. It has three simple operations: `arith_const` and `monomial`, to construct
 114 values in R , and `add` to add two values of R . To define the syntax and semantics of
 115 `LeanMLIR(QuotRing)`, we first declare the types and operations in the IR (Figure 1a).
 116 `QuotRing` has two types: `r`, which represents the ring R , and `nat` for naturals. Terms in `Op`
 117 represent the operations `arith_const`, `monomial` and `add`, and associated compile-time data.
 118 We then define the operation signatures by giving an instance of the `OpSignature` typeclass,
 119 which is offered by our framework to instantiate custom IRs (Figure 1b). That is, for each
 120 operation we specify: (1) the arity and types of arguments (`sig`), and (2) the type of the
 121 return value (`outTy`). The operation `arith_const` takes no arguments and returns a `nat`,
 122 `monomial` and `add` take two `nat/r`-valued arguments respectively, and both return an `r`.

123 The type denotation is also simple to express with the `TyDenote` typeclass (Figure 1c). `Ty`
 124 thus represents the embedded type in the IR and has only two inhabitants `r` and `nat`, whose
 125 denotation are `R` and `Nat`, the Lean (host) type that represents the mathematical objects R
 126 and \mathbb{N} respectively. The denotation of operations is a Lean function from the denotation of
 127 the input types (as recorded in the signature of that operation), to the denotation of the
 128 output type. Concretely, a `(arith_const n)` operation takes no arguments, so its denotation
 129 is an `nat`, while `add` takes two `r` arguments, so its denotation is a function from the product⁴
 130 of its arguments to its output, i.e., $R \times R \rightarrow R$. The same is true for `monomial` for $\mathbb{Z} \times \mathbb{Z} \rightarrow R$.
 131 We define the denotation of `(arith_const n)` to evaluate to `n`, `add(x, y)` to evaluate to `(x`
 132 `+ y)` and `monomial(a, i)` to `Quotient.mk (span generator p q (monomial a i))`, the
 133 equivalence class of aX^i . We also require a few lines of specific code to translate the MLIR
 134 abstract syntax tree (AST) to `Ty` (e.g., mapping `index` into `nat` or `R` to `r`) and `Op`, not shown
 135 here (details in Section 3.2). Together, these definitions instantiate `LeanMLIR(QuotRing)`.
 136 The `QuotRing` IR does not use regions. We will see examples of regions in (Section 6.2).

137 2.2 Defining and Executing Peephole Rewrites for QuotRing

138 We now verify the peephole rewrite $(a + X^{2^n} + 1 \mapsto a)$, where a is a variable and X^{2^n}
 139 is a constant in the ring. In $(\mathbb{Z}/q\mathbb{Z})[X]/(X^{2^n} + 1)$ this rewrite is simple to prove and,
 140 unsurprisingly, our custom `LeanMLIR(QuotRing)` IR enables us to rewrite at exactly this
 141 level. Any given peephole rewrite (of which Figure 2 is an example) consists of a context Γ of
 142 free variables in the search pattern of the peephole rewrite. The search pattern is called `lhs`,

³ We will discuss the underlying mathematical structure in more detail in Section 6.3

⁴ The mechanization uses a heterogeneous vector type `HVector`, which is coerced into the product type.

```

def a_plus_generator_eq_a : PeepholeRewrite Op [.r] .r := {
  lhs /- a + X^(2^n) + 1 -/ := [quotring_com q, n] {
    ^bb0(%a : !R):
      %one_int = arith.const 1 : i16
      %two_to_the_n = arith.const ${2**n} : index
      %x2n = poly.monomial %one_int, %two_to_the_n : (i16, index) -> !R
      %oner = poly.const 1 : !R
      %p = poly.add %x2n, %oner : !R
      %v1 = poly.add %a, %p : !R
      return %v1 : !R
  },
  rhs /- a -/ := [quotring_com q, n] {
    ^bb0(%a : !R):
      return %a : !R
  },
  correct := by
    funext Γv; simp_peephole [Nat.cast_one, Int.cast_one] at Γv 1
    /- ⊢ a + ((Quotient.mk (span {f q n})) ((monomial (2**n)) 1) + 1) = a -/
    ... /- simple proof with only definitions and theorems from Mathlib -/
}

```

■ **Figure 2** A peephole rewrite in `LeanMLIR(QuotRing)` asserts the semantic equivalence of two SSA programs given in MLIR syntax. Our proof automation through `simp_peephole` eliminates the framework overhead, such that closing a clean mathematical goal suffices to prove correctness.

143 and the replacement is `rhs`. The user has a proof obligation that the denotations of the left
 144 and right-hand sides are equal, which is given by the field `correct` of the peephole rewrite.

145 We declare our desired peephole rewrite (in Figure 2) by defining `a_plus_generator_eq_a`.
 146 Its type is `(PeepholeRewrite Op [.r] r)`, where the `Op` specifies the IR the rewrite belongs
 147 to and `[.r]` is the list of types of free variables in the program. For $(a + X^{2^n} + 1 \mapsto a)$, this is
 148 $(a : r)$. The final instruction we are matching yields a value of type `r`. The `lhs` is the program
 149 fragment we want to match on, with the free variable `%a` interpreted as being allowed to
 150 match any variable of type `r`. Observe that the type encapsulates exactly what is necessary
 151 for a well-typed match: the types of free variables (`r`) and the type of the instruction whose
 152 return value we are replacing (also `r` in this case). The rewritten program is the `rhs` field.

153 Both the left- and right-hand sides of the rewrite are written in MLIR syntax. (We
 154 use MLIR’s concise IR-specific syntax for readability here; our parser currently implements
 155 the slightly-more-verbose generic MLIR syntax). Note that we also include a custom
 156 quasiquotation `${2**n}` , to inline the symbolic (universally quantified) value n , even though
 157 the IR would require 2^n to be a concrete constant. Using MLIR syntax matches the LLVM
 158 community’s use of automation tooling, such as Alive: copy a code snippet and get a response.
 159 Our goal is to make the use of an interactive theorem prover part of the day-to-day workflow
 160 of compiler engineers. To enable this workflow, we implement a full MLIR syntax parser,
 161 along with facilities to convert from the generic MLIR abstract syntax tree (AST) into our
 162 framework type, such that we can use MLIR syntax in Lean.

163 To prove the correctness of `a_plus_generator_eq_a`, we use the `simp_peephole` 1
 164 tactic from our framework, which removes any overhead of our SSA encoding. We are
 165 left with: `⊢ a + ((Quotient.mk (span f q n)) ((monomial (2**n)) 1) + 1) = a`, a
 166 proof obligation in the underlying algebraic structure that, thanks to Lean’s `mathlib`, can
 167 be closed with a few (elided) lines of algebraic reasoning.

168 **2.3 Executing Peephole Rewrites**

169 Given a peephole rewrite `rw` and a source program `s`, we provide `rewritePeephole` to replace
 170 the pattern `rw.lhs` in the source program `s`. If the matching succeeds, we insert the target
 171 program `rw.rhs` (with appropriate substitutions) and replace all references to the original
 172 let-binding with a reference to the newly inserted `let`. Note that the matching is based on
 173 the def-use chain. Thus, a pattern need not be *syntactically* sequential in the program `s`.
 174 As long as the pattern `rw.lhs` can be found as *subprogram* of `s`, the subprogram will be
 175 rewritten. This makes our peephole rewriter an SSA peephole rewriter, which distinguishes
 176 it from a straight-line peephole rewriter that only matches a linear sequence of instructions.

177 Thanks to our intrinsically well-typed encoding, we know that the result of the rewriter
 178 is always a well-typed program, under the same context and resulting in the same type
 179 as the original program. Furthermore, the framework extends the local proof of semantic
 180 equivalence to a global proof, showing that the rewriter is semantics preserving:

```
177 /- The denotation of the rewritten program is equal to the source program. -/
178 theorem denote_rewritePeephole (fuel : ℕ) (pr : PeepholeRewrite Op Γ t)
179   (target : Com Op Γ₂ t₂) : (rewritePeephole fuel pr target).denote = target.denote
```

181 These typeclass definitions are all we need to define the `QuotRing` IR. Our framework
 182 takes care of building the intrinsically well-typed IR for `QuotRing` from this, and gives us a
 183 verified peephole rewriter, with other optimizations like CSE and DCE. We will now delve
 184 into the details of the framework and see how it achieves this.

185 **3 LeanMLIR(X): A Framework for Intrinsically Well-Typed SSA**

186 In this section, we describe the core design of the framework: the encoding of programs and
 187 their semantics in `LeanMLIR(X)` (Figure 3a). We review some dependently-typed tooling we
 188 use to define our IR. **Contexts:** Our encoding is intrinsically well-typed (i.e., each inhabitant
 189 of `Expr` or `Com` described below is, by construction, well typed). Thus, we need a *context* to
 190 track the types of variables that are allowed to occur (`Ctxt Ty`). A context is a list of types,
 191 where for example `[int, int, bool]` means that there are two variables of the (user-defined)
 192 type `int` and one variable of type `bool` we may refer to. **Variables:** The type `(Var Γ α)`
 193 encodes variables of type `α` in context `Γ`. We use De Bruijn indices [29] in the standard
 194 way, but, additionally, a variable with index `i` also carries a proof witness that the `i`-th entry
 195 of context `Γ` is the type `α`. **Heterogeneous Vectors:** To define an argument signature
 196 (`OpSignature.sig`), say, `[int, int, bool]`, we need an expression with this operation to
 197 store two variables of type `int` and one of type `bool`. We want to statically ensure that the
 198 types of these variables are correct, so we store them in a heterogeneous vector. A vector of
 199 type `HVector f [α₁, ..., αₙ]` is equivalent to a tuple `(f α₁ × ... × f αₙ)`.

200 **3.1 Semantics of LeanMLIR(X)**

201 The core types for programs are `Expr` and `Com`, shown in Figure 3a. The type `(Expr Γ`
 202 `α)` describes individual MLIR operations; we think of it as a function from values in the
 203 context `Γ`—also called a *valuation* for that context—to a value in the denotation of type `α`.
 204 Commands `(Com Γ α)` has a similar interpretation but represents sequences of operations.
 205 Each command binds a new value in the current context (the `var` constructor) until the
 206 sequence returns the value of one such variable `v` (the `ret` constructor). Thus, this encoding
 207 of SSA exploits the similarity to the ANF [2] and CPS [14] encodings. The semantics given


```

inductive Expr [OpSignature Op Ty] : Ctxt Ty → Ty → Type where
| mk (op : Op) -- op (arg1, arg2, ..., argn) : outTy op
  (args : HVector (Var Γ) (OpSignature.sig op)) : Expr Γ (OpSignature.outTy op)

inductive Com [OpSignature Op Ty] : Ctxt Ty → Ty → Type where
| ret (v : Var Γ α) : Com Γ α -- return v
| var (e : Expr Γ α) (body : Com (Γ.snoc α) β) : Com Γ β -- let v : α := e in body

```

(a) Core syntax of $\text{LeanMLIR}(X)$, polymorphic over Op . The arguments in square brackets are assumed typeclass instances. Type is the base universe of Lean types.

```

variable [TyDenote Ty] [OpDenote Op Ty] [DecidableEq Ty]

def Expr.denote : {ty : Ty} → (e : Expr Op Γ ty) → (Γv : Valuation Γ) → toType ty
| _, ⟨op, args⟩, Γv => OpDenote.denote op (args.map (fun _ v => Γv v))

def Com.denote : Com Op Γ ty → (Γv : Valuation Γ) → (toType ty)
| .ret e, Γv => Γv e
| .var e body, Γv => body.denote (Γv.snoc (e.denote Γv))

```

(b) Denotation of Expr and Com in $\text{LeanMLIR}(X)$, which extends the user's OpDenote to entire programs. Intrinsic well-typing of Com makes its denotation a well-typed function from the context valuation to the return type. The angled brackets are used to pattern match on a structure constructor anonymously.

■ **Figure 3** Definitions in $\text{LeanMLIR}(X)$ for Expr and Com , and their associated denotations.

208 by the user in `OpDenote` are extended to semantics for Expr and Com (Figure 3b) by the
 209 framework. An Expr evaluates its arguments by looking up their value in the valuation and
 210 then invokes the user-defined `OpDenote.denote` to evaluate the semantics of the `op`.

211 3.2 Writing $\text{LeanMLIR}(X)$ Programs Using MLIR Syntax

212 An important goal for our framework is to provide easy access to formalization for the MLIR
 213 community. Toward this goal, we have a deep embedding of MLIR's AST and a corresponding
 214 parser. This is developed using Lean's syntax extensions [33]. We augment this with a
 215 generic framework to build Expr and Com terms from a raw MLIR AST. This framework
 216 allows the user to pattern-match on the MLIR AST to build intrinsically well-typed terms,
 217 as well as to throw errors on syntactically correct, but malformed MLIR input. These are
 218 used by our framework to automatically convert MLIR syntax into our SSA encoding, along
 219 with the ability to provide precise error messages in cases of translation failure. This enables
 220 us to write all our examples in MLIR syntax, as demonstrated throughout the paper.

221 More concretely, we have an embedded domain-specific language (EDSL), which declares
 222 the MLIR grammar as a Lean syntax extension. As part of this work, we have found several
 223 inconsistencies with the MLIR language reference and contributed patches upstream to
 224 update them.⁵ Overall, this gives users the ability to write idiomatic MLIR code into our
 225 framework and receive an MLIR AST. Moreover, as we will showcase in the examples, our
 226 EDSL is idiomatically embedded into Lean, which allows us to quasiquote Lean terms. This
 227 will come in handy to write programs that are generic over constants, such as parameterizing
 228 a program by 2^n for any choice of n . We build our intrinsically well-typed data structures
 229 from this MLIR AST by writing custom elaborators.

⁵ reviews.llvm.org/{D122979, D122978, D122977, D119950, D117668}

23:8 Verifying Peephole Rewriting In SSA Compiler IRs

```
structure OpSignature (Tv : Type) where /- (1) Extending signature. -/  
  regSig : List (Ctxt Ty × Ty)  
  ...  
  
class OpDenote [TyDenote Ty] [OpSignature Op Ty] where /- (2) Extending denotation. -/  
  denote : (op : Op) → (args : HVector toType (OpSignature.sig op)) →  
    (regArgs : HVector (fun (ctx, t) => Valuation ctx → toType t) (OpSignature.regSig op)) →  
    (toType (OpSignature.outTy op))  
  
inductive Expr : (Γ : Ctxt Ty) → (ty : Ty) → Type where  
  | mk (op : Op)  
  | ...  
  (regArgs : HVector (fun (ctx, ty) => Com ctx ty) (OpSignature.regSig op)) :  
  Expr Γ ty  
  
mutual /- (3) extending expression denotation to recursively invoke regions. -/  
  def Expr.denote : {ty : Ty} → (e : Expr Op Γ ty) → (Γv : Γ.Valuation) → (toType ty)  
  | _, ⟨op, args, regArgs⟩, Γv =>  
  OpDenote.denote op (args.map (fun ty v => Γv v)) regArgs.denote  
  ...  
end
```

■ **Figure 4** Extending LeanMLIR(X) with regions. New fields are **in green**. In `OpDenote`, one can now access the sub-computation represented by the region when defining the semantics of `Op`.

230 3.3 Modelling Control Flow in LeanMLIR(X) With Regions

231 So far, our definition of `Com` only allows straight-line programs. To be able to model control
232 flow, we add regions to our IR. Regions add the syntactic ability to nest IR definitions,
233 thereby allowing syntactic encoding of concepts such as structured control flow. This is in
234 contrast with the approach of having a sea of basic blocks in a control-flow graph (CFG) that
235 are connected by branch instructions. More specifically, structured control flow with regions
236 allows modeling reducible control flow [1]. General CFGs allow us to represent more complex,
237 irreducible control flow, which makes them harder to reason about. Consequently, compiler
238 frameworks such as MLIR have moved toward directly representing structured control flow,
239 and we follow their approach. Notably, region arguments replace phi nodes in MLIR.

240 Intuitively, regions allow an `Op` to receive `Com`s as arguments, and choose to execute these
241 `Com` arguments zero, one, or multiple times. This allows us to model `if` conditions (by executing
242 the regions zero or once), loops (by executing the region n times), and complex operations
243 such as tensor contractions and convolutions by executing the region on the elements of
244 the tensor [34]. We implement this by extending `Expr` with a new field representing region
245 arguments (Figure 4). We also extend `OpSignature` with an extra argument for the input
246 types and output types of the region. In parallel, we add the denotation of regions as an
247 argument, extending `OpDenote`. Similarly, we extend the denotation of `Expr` to compute the
248 denotation of the region `Com`s in the `Expr`, before handing off to `OpDenote`.

249 This extension to our core calculus gives us the ability to model structured nesting of
250 programs. This is used pervasively in MLIR, to represent `if` conditions, `for` loops, and higher-
251 level looping patterns such as multidimensional strided array accesses over multidimensional
252 arrays (tensors). We show how to model control flow in Section 6.2.

253 4 Reasoning About LeanMLIR(X)

254 The correctness of peephole rewriting is a key aspect of the metatheory of LeanMLIR(X). We
255 begin by sketching the mechanized proof of correctness of peephole rewriting. We then discuss
256 how the infrastructure built for this proof is reused to prove two other SSA optimizations:

257 common subexpression elimination (CSE) and dead code elimination (DCE). Finally, we discuss
 258 our proof automation, which manipulates the IR encoding at elaboration time to eliminate
 259 all references to the framework and provide a clean goal to the proof engineer.

260 4.1 Verified SSA Rewriting With `rewritePeephole`

261 We now provide a sketch of the mechanized correctness proof of `rewritePeephole`. The
 262 key idea is that to apply a rewrite at location i , we open up the `Com` at location i in
 263 terms of a zipper [11]. This zipping and rewriting at a location i is implemented by
 264 `rewritePeepholeAt`. The zipper comprises of `Lets` to the left-hand side of i , and `Com` to
 265 the right: `let $x_2 = x_1$; (let $x_3 = x_2$; (let $x_4 = x_3$; (return x_3)))`: `Com [x_1] α` =

266 `((let $x_2 = x_1$); let $x_3 = x_2$); : Lets [x_1] [x_1, x_2, x_3]`
 267 `(let $x_4 = x_3$; (return x_3)) : Com [x_1, x_2, x_3] α`

268 The use of a zipper enables us to easily traverse the sequence of let-bindings during
 269 transformation and exposes the current `let` binding being analyzed. This exposing is
 270 performed by `Lets`, which unzips a `Com` such that the outermost binding of a `Lets` is the
 271 innermost binding of a `Com`. This forms the zipper, which splices the `Com` into a `Com` and
 272 a `Lets`. Also, while `Com` tracks only the return type α in the type index, `Lets` tracks the
 273 entire resulting context Δ . That is, in `(lets : Lets Γ Δ)`, the first context, Γ , lists all
 274 free variables (just as in `Com Γ τ`), but the second context, Δ , consists of all variables in Γ
 275 plus a new variable for each let-binding in the sequence `lets`. We can thus think of Δ as
 276 the context at the current position of the zipper. Another difference is the order in which
 277 these sequences grow. Recall that in `Com`, the outermost constructor represents the topmost
 278 let-binding. In `Lets`, the outermost constructor instead corresponds to the bottommost
 279 let-binding. This difference is what makes the zipper work.

280 We have two functions to go from a program to a zipper and back: (1) `(splitProgramAt`
 281 `pos prog)`, to create a zipper from a program `prog` by moving the specified number of
 282 bindings to a new `Lets` sequence, and (2) `(addComInMiddleOfLetCom top mid bot)`, to
 283 turn a zipper `top`, `bot` into the program, while inserting a program `mid : Com` in between.
 284 We also prove that the result of splitting a program with `splitProgramAt` is semantically
 285 equivalent to the original program. Similarly, we prove that stitching a zipper back together
 286 with `addComInMiddleOfLetCom` results in a semantically equivalent program.

287 Given a peephole rewrite (`matchCom`, `rewriteCom`), to rewrite at location i , we first split
 288 the target program into `top` and `bot`. We then attempt to match the def-use chain of the
 289 return variable in `matchCom` with the final variable in `top` (which is the target i , since we
 290 split the program there). This matching of variables recursively matches the entire expression
 291 tree.⁶ Upon successful matching, this returns a substitution σ for the free variables in
 292 `matchCom` in terms of (free or bound) variables of `top`. Using this successful matching, we
 293 stitch the program together as `top; σ (rewriteCom); τ (bot)`. Here, τ is another substitution
 294 that replaces the variable at location i with the return variable of `rewriteCom`. Since we
 295 derived a successful matching, we know that the semantics of variable i is equal to that
 296 of the return variable of `matchCom`. By assumption on the peephole rewrite, the variable
 297 i is equivalent to the return variable of `rewriteCom`. This makes it safe to replace all
 298 occurrences of the variable i in `bot` with the return variable of `rewriteCom`. This proves
 299 `denote_rewritePeephole`, which states that if a rewrite succeeds, then the semantics of

⁶ We match regions in expressions for structural equality. We *do not* recurse into regions during matching, and treat regions as black-boxes.

300 the program remain unchanged. In this way, we use a zipper as a key inductive reasoning
 301 principle to mechanize the proof of correctness of SSA-based peephole rewriting.

302 4.2 DCE & CSE: Folding Over Intrinsically Well Typed SSA

303 The classic optimizations enabled by SSA are peephole rewriting, dead code elimination
 304 (DCE), and common subexpression elimination (CSE). We implement these optimizations
 305 in our framework as a test of its suitability for metatheoretic reasoning. Our approach
 306 is different from previous approaches [41, 5] with our use of intrinsic well-typing, which
 307 mandates proofs of the structural rules on contexts to rewrite programs. We begin by
 308 building machinery to witness that a context Δ is equal to the context Γ , minus the variable
 309 x . This is spelled as `Deleted Γ x Δ` in `LeanMLIR(X)`. We then prove context-strengthening
 310 theorems to delete variables that do not occur in `Expr` and `Com` while preserving denotation.

311 Using this tooling, DCE is implemented in ≈ 400 LoC, which shows that our framework
 312 is well-suited to metatheoretic reasoning. The implementation is written in a proof-carrying
 313 style, interleaving function definitions with their proof of correctness. The recursive step of
 314 the dead code elimination takes a program $p : \text{Com } \Gamma \ \tau$ and a variable v to be deleted, and
 315 returns a new $p' : \text{Com } \Delta \ \tau$. The two contexts Γ and Δ are linked by a context morphism
 316 (`Hom Γ Δ`), to interpret p' (with the deleted variable) which lives in a strengthened context
 317 Δ in the old context Γ . We walk p recursively to eliminate dead values at each `let` binding.
 318 This produces a new p' with dead bindings removed, a proof of semantic preservation, and
 319 a context morphism from the context of p to the strengthened context of p' with all dead
 320 variables removed.

321 Similarly, the CSE implementation folds over `Com` recursively, maintaining data structures
 322 necessary to map variables and expressions to their canonical form. At each (`let $x =$
 323 $f(v_1, \dots, v_n)$ in b`) step, we canonicalize the variables v_i to find variables c_i . We then look up
 324 the canonicalized expression $f(c_1, \dots, c_n)$ in our data structure to find the canonical variable
 325 c_x if it exists and replace x with c_x . If such a canonical c_x does not exist, we add a new
 326 entry mapping $f(c_1, \dots, c_n)$ to x , thereby canonicalizing any further uses of this expression.

327 4.3 Proof Automation for Goal State Simplification in `LeanMLIR(X)`

328 The proof automation tactic `simp_peephole Γ` (used to eliminate framework definitions
 329 from the goal state) takes a context Γ , reduces its type completely, and abstracts out program
 330 variables to provide a theorem statement that is universally quantified over the variables of
 331 the program, with all framework definitions eliminated. It uses a set of equation theorems
 332 to normalize the type of Γ . This is necessary to extract the types of variables during
 333 metaprogramming. Once the type of Γ is known, we simplify away all framework definitions
 334 (such as `Expr.denote`). We then replace all occurrences of a variable accesses `$\Gamma[i]$` with a
 335 new (Lean, i.e., host) variable. We do this by abstracting terms of the form `$\Gamma[i]$` where i is
 336 the i -th variable. This gives us a proof state that is universally quantified over variables from
 337 the context. Finally, we clear the context away to eliminate all references to the context
 338 Γ . The set of definitions we simplify away is extensible, enabling us to add domain-specific
 339 simplification rewrites for the IR.

340 5 Pure Rewriting in a Side-Effectful World

341 While `LeanMLIR(X)` streamlines the verification of higher-level IRs that use only value
 342 semantics, typical IRs may interleave islands of pure operations (with value semantics) with

343 operations that carry side effects. An IR that is user-facing can usually be rephrased with
 344 high-level, side-effect-free semantics. Yet, each operation in such an IR is compiled through a
 345 sequence of IRs that are lower level and potentially side-effectful. For example, in the case of
 346 FHE, the pure FHE IR is compiled to a lower-level IR that encodes the coset representative
 347 of each ideal as an array, with control flow represented via structured control flow (`scf`).
 348 Eventually, this is compiled into LLVM which is rife with mutation and global state. In
 349 such a compilation flow, peephole rewrites are used at each intermediate IR to optimize pure
 350 fragments while leaving side-effectful fragments untouched. An effective compiler pipeline
 351 introduces the right abstractions to maximize rewrites on side effect-free fragments.

352 `LeanMLIR(X)` is designed to facilitate verification of peephole rewrites as they arise in such
 353 a compiler pipeline. The previous sections already presented how our framework supports
 354 the verification of peephole rewrites in a pure setting. Yet, our design also allows for the
 355 optimization of a pure fragment in a side-effectful context. We have a mechanized proof of
 356 the correctness of the extended framework with support for side effects and a rewrite theorem
 357 that performs pure rewrites in the presence of side effects. The key idea is to annotate each
 358 `Op` with an `EffectKind`, where `EffectKind.pure` changes the denotation of the `Expr` into
 359 the `Id` monad, while `EffectKind.impure` denotes into an arbitrary, user-chosen, IR-specific
 360 monad. We also introduce a new notion of monadic evaluation of `Lets`, which returns a
 361 valuation plus a proof that, for every variable v that represents a pure expression e in the
 362 sequence of let-bindings, the valuation applied to v agrees with the (pure) denotation of
 363 e . This proof-carrying definition allows us to use this invariant when reasoning inside a
 364 subexpression of a monadic bind.

365 With the above at hand, the overall rewriter construction and proof strategy remains
 366 unchanged, with the additional constraint of performing rewriting only on those operations
 367 marked as `EffectKind.pure`, and the surrounding monadic ceremony required to show that
 368 a pure rewrite indeed does not change the state of pure variables in various lemmas.⁷

369 **6 Case Studies**

370 We mechanize three IRs based on ones found in the MLIR ecosystem as case studies for
 371 `LeanMLIR(X)` and show how they benefit from the different aspects of our framework. Note
 372 that the core of our framework (definitions of `Expr`, `Com`, `PeepholeRewrite`, lemmas about
 373 these objects, and the peephole rewriting theorem) is $\approx 2.2k$ LoC. The case studies based
 374 on our framework together are $\approx 5.6k$ LoC, which stresses the framework to ensure that it
 375 scales to realistic formal verification examples.

376 **6.1 Reasoning About Bitvectors of Arbitrary Width**

377 We first demonstrate our ability to reason about a well-established domain of peephole
 378 rewrites: LLVM’s arithmetic operations over fixed-bitwidth integers. Using the Z3 SMT
 379 solver [7], the Alive project [21, 20] can efficiently and automatically reason about these.
 380 Notably, at the time of this writing, almost 700 LLVM patches have justified their correctness
 381 by referencing Alive. In this way, accessible proof tools can find a place in production compiler
 382 development workflows. However, Alive is limited by the capabilities of the underlying SMT
 383 solvers. SMT solvers are complex, heuristic-driven, and sometimes even have soundness

⁷ A limitation of our current mechanization is that we assume that all regions are potentially side-effecting. This is a simplification that shall be addressed in a newer version of the proof.

384 bugs [38]. They are also specialized to support very concrete theories. Among others, this
 385 means Alive can only reason about a given fixed bitwidth. Even recent work that specifically
 386 aims to generalize rewrites to arbitrary bitwidths, can only exhaustively test a concrete set
 387 of bitwidths [23]. Using our framework, we can reproduce Alive-style correctness proofs, and
 388 extend them to reason about arbitrary (universally quantified) bitwidths. This ability to
 389 handle arbitrary bitwidth is important in verification contexts that have wide bitvectors, as
 390 they can occur in real-life VLSI problems [12, 36]. MLIR itself has multiple IRs that require
 391 bitvector reasoning: `comb` for combinational logic in circuits, `arith` and `index` for integer
 392 and pointer manipulation, and `llvm` which embeds LLVM IR in MLIR. Our streamlined
 393 verification experience offers developers an Alive-style workflow for the `llvm` dialect, while
 394 allowing reasoning across bitwidths. As our framework is extensible, we believe we can also
 395 support other MLIR dialects that require bitvector reasoning, such as `comb` and `arith`.

396 6.1.1 Modeling a fragment of LLVM IR: Syntax and Semantics

397 To test our ability to reason about bitvectors in practice, we model the semantics of the
 398 arithmetic fragment of LLVM as the IR `LeanMLIR(LLVM)`. We support the (scalar) operators:
 399 `not`, `and`, `or`, `xor`, `shl`, `lshr`, `ashr`, `urem`, `srem`, `add`, `mul`, `sub`, `sdiv`, `udiv`, `select` and
 400 `icmp`. We support all `icmp` comparison flags, but not the strictness flags `nsw` and `nuw`.

401 At the foundation of our denotational semantics is Lean’s `BitVec` type, which models
 402 bitvectors of arbitrary width and offers `smtlib` [4] compatible semantics. However, when we
 403 started this work, most bitvector operations were not defined in the Lean ecosystem and the
 404 bitvector type itself was not fully fleshed out. Hence, we worked with the `mathlib` and Lean
 405 community to build and upstream a theory of bitvectors.⁸ After developing the core theory
 406 in `mathlib`, Lean’s mathematical library, development subsequently moved into Lean core,
 407 where we continue to evolve Lean’s bitvector support.

408 The semantics of LLVM’s arithmetic operations follow the semantics of `smtlib` (and
 409 consequently Lean’s) bitvectors closely. In case of integer wrapping or large shifts, for
 410 example, LLVM can produce so-called poison values [21], which capture undefined behavior
 411 as a special value adjoined to the bitvector domain. LLVM’s poison is designed not to be
 412 a side effect and, consequently, can be reasoned about in a pure setting. In contrast, `ub`
 413 is a side effect that triggers immediate undefined behaviour, and can be refined into any
 414 behavior. In LLVM, the following refinements are legal: `ub` \sqsubseteq `poison` \sqsubseteq `val`. Among the
 415 instructions we model, division and remainder can produce immediate undefined behavior `ub`.
 416 In our framework, we approximate these by collapsing the side-effectful undefined behavior
 417 and side-effect-free poison both into `Option.none`. We thus denote bitvectors into the type
 418 `Option (BitVec w)`. This is safe for the three kinds of rewrites we consider: (1) the left-
 419 and right-hand sides are both UB and poison free (arithmetic rewrites), (2) the left- and
 420 right-hand sides are both UB free (bitwise rewrites, where left and right shifts may produce
 421 poison), (3) the left-hand side may trigger UB, while the right hand side may only produce
 422 poison (e.g., refining division into arithmetic operations). We leave separating UB as a side
 423 effect distinct from poison, and reasoning about peephole rewrites which refine such side
 424 effects as interesting future work.

425 For side-effect-free programs, our semantics match the LLVM semantics. We perform

⁸ github.com/leanprover-community/mathlib4/pull/{5383,5390,5400,5421,5558,5687,5838,5896,7410,7451,8231,8241,8301,8306,8328,8345,8353},
github.com/leanprover/lean4/pull/{3487,3471,3461,3457,3445,3492,3480,3450,3436},
github.com/leanprover/std4/pull/{357,359,599,626,633-636,637,639,641,645-648,655,658-660,653}

426 exhaustive enumeration tests between our semantics and that of LLVM. We take advantage
 427 of the fact that an IR with computable semantics automatically defines an interpreter in our
 428 framework. We build an executable program that runs every instruction, with all possible
 429 input combinations upto bitwidth 8. We get LLVM’s ground truth by using LLVM’s optimizer,
 430 `opt` to transform the same instruction with constant inputs. This optimizes the program
 431 into a constant output, handling undefined behavior. By exhaustive enumeration, our tested
 432 executable semantics correspond to the LLVM semantics wherever the result is `Option.some`,
 433 and also soundly model undefined behavior whenever the result is `Option.none`. This gives
 434 us confidence our semantics correspond to LLVM’s.

435 6.1.2 Proving Bitvector Rewrites in our Framework

436 Effective automation for bitvector reasoning is necessary to resolve the proof obligations
 437 that `LeanMLIR(X)` derives automatically from peephole rewrites expressed as MLIR program
 438 snippets. While Lean does not yet have extensive automation for bitvectors, thanks to our
 439 work we can already use a decision procedure for commutative rings [10] and an extensionality
 440 lemma that establishes the equality of bitvectors given equality on an arbitrary bit index.

441 We test the available automation on a dataset of peephole optimizations from Alive’s
 442 test suite, consisting of theorems about addition, multiplication, division, bit-shifting and
 443 conditionals. Out of the 435 tests in Alive’s test suite, we translate 93 tests which are the
 444 ones that are supported by the LLVM fragment we model and without preconditions. We
 445 prove 54 of these rewrites from the Alive test suite automatically. Some rewrites cannot
 446 be handled automatically. Of those where automation struggles, we manually prove an
 447 additional 6, selecting the ones where an SMT solver takes long to prove them even for a
 448 specific bitwidth (e.g., 64). Our proofs are over arbitrary (universally quantified) bitwidth,
 449 save for some theorems that are only true at particular bitwidths.⁹ As an example, let us
 450 consider the rewrite:

```
example (w : Nat) :
  [llvm ( w )| {
    ^bb0(%X : _, %C1 : _, %C2 : _):
      %v1 = llvm.xor %X, %C1
      %v2 = llvm.and %v1, %C2
      llvm.return %v2
  }] ⊆ [llvm ( w )| {
    ^bb0(%X : _, %C1 : _, %C2 : _):
      %v1 = llvm.and %X, %C2
      %v2 = llvm.xor %X, %C1
      %v3 = llvm.and %C1, %C2
      %v4 = llvm.xor %v1, %v3
      llvm.return %v4
  }] := by simp_alive_peephole; alive_auto
```

451 Note that due to the support of MLIR syntax in our framework, this rewrite is specified
 452 in MLIR syntax. We use a custom extension with the placeholder syntax `_`, to stand for an
 453 arbitrary bitwidth w . After simplification of the framework code with `simp_peephole`, this
 454 yields the proof obligation:

```
455 (w : Nat) (X C1 C2 : BitVec w) ⊢ (X ^^^C1) &&& C2 = X &&& C2 ^^^C1 &&& C2
```

456 This proof obligation only concerns the semantics in the semantic domain of bitvectors,
 457 it does not feature MLIR and SSA anymore. This goal is automatically proven by our proof

⁹ e.g., $a + b = a \text{ xor } b$ is true only at bitwidth 1.

23:14 Verifying Peephole Rewriting In SSA Compiler IRs

```
/-- only control flow operations, parametric over another IR Op' -/  
inductive Op (Op': Type) [OpDenote Op' Ty'] : Type  
| coe (o : Op') -- coerce Op' to Op  
| for (ty : Ty') -- a for loop whose loop carried data is Ty'  
  
instance [I : HasTy Op' Int] : OpSignature (Op Op') Ty' where  
  signature  
  | .coe o => signature o  
  | .for t => <[/-start-/I.ty, /-step-/I.ty, /-nitters-/N.ty, /-v-/t],  
    /- region arguments: -/ [(/-i-/I.ty, /-v-/t], /-v'-/t)],  
    /-return-/t>  
instance [I : HasTy Op' Int] [OpDenote Op' Ty'] : OpDenote (Op Op') Ty' where  
  denote  
  | .coe o', args', regArgs' => OpDenote.denote o' args' regArgs' -- reuse denotation of o'  
  | .for ty, [istart, istep, niter, vstart]h, [f]h =>  
    let istart : ℤ := I.denote_eq ▶ istart -- coerce to `int`.  
    ... -- coerce other arguments  
    let loop_fn := ... -- build up the function that's iterated.  
    (loop_fn (istart, vstart)).2
```

■ **Figure 5** Simplified implementation of $\text{LeanMLIR}(scf(X))$. Observe that the IR is parametrized over another IR Op' , and that we add control flow to the other IR in a modular fashion.

458 automation for bitvectors, `alive_auto`. In the longer term, we aim to also connect our work
459 to a verified SAT checker that is under development.¹⁰

460 6.2 Structured Control Flow

461 The examples of IRs we have seen so far are all straight-line code. In this use case, we show
462 how we can add control flow to existing IRs, thanks to the parametricity of our framework.
463 We also demonstrate how encoding control flow structures as regions enable succinct proofs
464 for transformations, by exploiting the high-level structure of these operations. To this end,
465 we model structured control flow as a fragment of the `scf` IR in MLIR, by giving semantics
466 to two common kinds of control flow: `if` conditions and bounded `for` loops. Note that we
467 choose to model *bounded* for loops, since these are the loops that are used in MLIR to model
468 high-level operations such as tensor contractions. A pleasant upshot is that these guaranteed
469 to terminate, and can thus have a denotation as a Lean function without requiring modelling
470 of nontermination (which is side-effectful). Our sketch of the extended framework with side
471 effects will be used to pursue this line of research in the future. The conditionals and bounded
472 for loop allow us to concisely express loop canonicalizations and transformations from MLIR
473 in $\text{LeanMLIR}(scf)$. These operations allow us to concisely express loop canonicalizations and
474 transformations in $\text{LeanMLIR}(scf)$.

475 We built this parametrically over an existing IR X to allow these constructs to be added
476 to an existing IR X . The key idea is that the `Op` corresponding to `scf` is parametrized by the
477 `Op` corresponding to another IR X . Since the only datatypes `scf` requires are booleans and
478 natural numbers, we ask that the type domain of X contains these types. We then provide
479 denotations in `for` for booleans and integers from the type domain of X . Thus, what we encode
480 is $\text{LeanMLIR}(scf(X))$, which is an IR for structured control flow parametrized by another,
481 user-defined IR X .

482 The `scf.for` operation (Figure 5) has three arguments: the number of times the loop is
483 to be executed, a starting and step value for the iteration, and a seed value for the loop to

¹⁰ <https://github.com/leanprover/leansat>

484 iterate on. Note that in the definition, the IR `Op` is defined parametrically over another IR
 485 `Op'`, and the types of `Op` are the same as the types of other IR `Ty'`. We perform a similar
 486 construction for `if` conditions.

487 The denotation of the `for` loop, as well as theorems about loop transformations, follow
 488 from `mathlib`'s theory for iterating functions (`Nat.iterate`). The loop body in `scf.for`
 489 has a region that receives the current value of the loop counter and the current iterated value
 490 and returns the next iterated value. We prove the inductive invariant for loops using the
 491 standard theory of iterated function compositions ($f^0 = id$, $f^k \circ f^l = f^{k+l}$, $id^k = id$). We
 492 also prove common rewrites over loops: running a `for` loop for zero iterations is the same as
 493 not running a loop at all (dead loop deletion), two adjacent loops with the same body can
 494 be fused into one when the ending index of the first loop is the first index of the second loop
 495 (loop fusion), and a loop whose loop body does not depend on the iteration count can be
 496 reversed (loop reversal). Similarly, we prove that `if true e e' = e`, and `if false e e' = e'`.

497 These do *not* count as peephole rewrites in our framework, as they are universally
 498 quantified over the loop body (which is a region). This is unsupported — peephole rewrites
 499 in `LeanMLIR(X)` may only have free variables, not free region arguments. Increasing the
 500 power of peephole rewrites with arbitrary regions is an interesting question for future work.

501 Consider the loop optimization that converts iterated addition into a single multiplication.
 502 Its proof obligation is $(\vdash \lambda x. x + \delta)^n(c) = n \cdot \delta + c$ (a short proof by induction on n). This
 503 transformation is challenging to perform in a low-level IR, since there is no syntactic concept
 504 of a loop. However, this transformation *is* a valid peephole rewrite in our framework since it
 505 uses a *statically* known loop body. We showcase how regions permit MLIR (and, consequently,
 506 us) to easily encode and reason with commonplace loop transformations. Importantly, the
 507 parametricity of our framework allows us to prove a set of these as local peephole rewrites
 508 that are valid on all IR extensions `scf(X)`.

509 6.3 Fully Homomorphic Encryption

510 A key motivation for `LeanMLIR(X)` is to enable specifying formal semantics for high-level,
 511 mathematical IRs. These IRs require access to complex mathematical objects that are
 512 available in proof assistants, and verifying rewrites on such IRs is out of practical reach for
 513 today's SMT solvers. As a case study, we formalize the complete “Poly” IR.¹¹ This IR is a
 514 work in progress and is in flux, as it is part of the discussion of an upcoming open standard
 515 for homomorphic encryption, developed in collaboration by Intel and Google.¹² Contrary
 516 to what its naming implies, this IR does *not* model operations on polynomials.¹³ Instead,
 517 codewords are encoded as elements in a finitely-presented commutative ring, specifically,
 518 the ring $R \equiv (\mathbb{Z}/q\mathbb{Z})[x]/(x^{2^n} + 1)$, where $q, n \in \mathbb{N}$ are positive integers (q composite). The
 519 name “Poly” comes from the equivalence class representatives are polynomials, but not all
 520 IR operations are invariants of the equivalence class.

521 The “Poly” IR is, in fact, a superset of the `QuotRing` IR we defined in Section 2. It consists
 522 of the operations `add`, `sub`, `mul`, `mul_constant`, `leading_term`, `monomial`, `monomial_mul`,
 523 `from_tensor`, `to_tensor`, `arith.constant` and `constant`.¹⁴

524 Most of these operations are self-explanatory and derive from the (commutative) ring
 525 structure of R or are used to build elements in R , like the equivalence classes of constants

¹¹ as of commit 2db7701de

¹² <https://homomorphicencryption.org/>

¹³ In the same way that rationals \mathbb{Q} are not pairs of integers $\mathbb{Z} \times \mathbb{Z}$.

¹⁴ It also has distinct types for integers and naturals, which we unified in Section 2 for simplicity.

526 or monomials. Three operations, `to_tensor` and `from_tensor` and `leading_term` do not
 527 follow directly from the algebraic properties of the polynomial ring. Instead, they depend
 528 on a (non-canonical) choice of representatives for each ideal coset in the polynomial ring.
 529 More precisely, let $\pi : (\mathbb{Z}/q\mathbb{Z})[x] \rightarrow (\mathbb{Z}/q\mathbb{Z})[x]/(x^{2^n} + 1)$ be the canonical surjection into the
 530 quotient, taking a polynomial to its equivalence class modulo division by $x^{2^n} + 1$. Further let
 531 $\sigma : (\mathbb{Z}/q\mathbb{Z})[x]/(x^{2^n} + 1) \hookrightarrow \mathbb{Z}/q\mathbb{Z}[x]$ be the injection taking an equivalence class to its (unique)
 532 representative with degree $\leq 2^n$. This is a right-inverse of π , i.e. $\pi \circ \sigma = id$. Note that
 533 multiple right-inverses that could have been chosen for σ , as long as $\sigma(x)$ is a representative
 534 of the equivalence class of x for all $x \in (\mathbb{Z}/q\mathbb{Z})[x]/(x^{2^n} + 1)$, σ will be a right-inverse of
 535 π . The operation `to_tensor(p)` returns the vector $(\sigma(p)[i])_{i=0,\dots,2^n}$, where $a[i]$ represents
 536 the i -th coefficient, i.e. $\sigma(p) = \sum_{i=0}^{2^n} (\sigma(p)[i])x^i$, and `to_tensor` the converse. Similarly,
 537 `leading_term(p)` returns the equivalence class of the leading term of the representative $\sigma(p)$
 538 (which also depends on the choice of σ).

539 This allows us to define the semantics and prototype both the IR and rewrites in it.
 540 Rewrites like `mul(p, q) → mul(q, p)` follow immediately from the fact that R is a commutative
 541 ring. Other rewrites like `from_tensor(to_tensor(p)) → p`, or even `add(p, monomial(1, 2^n))`
 542 `→ sub(p, 1)`, on the other hand, are more specific to this IR and have a higher manual-proof
 543 overhead. We prove all of these.

544 We discussed the IR and potential semantics with the authors of the HEIR IR in the
 545 context of the upcoming open standard for homomorphic encryption. We believe that a
 546 framework like the one presented in this paper will allow standards like these to be defined
 547 with formal semantics from the ground up.

548 7 Related Work

549 Alive [21, 18] and Alive 2 [20] provide push-button verification for a subset of LLVM by
 550 leveraging SMT solvers. Alive-tv does the same for a set of concrete IRs for tensor operations
 551 in MLIR [3]. The semantics and correctness of compiling compositionally have been explored
 552 by multiple authors, like Pilsener [25] or many variants of CompCert [19]: like compositional
 553 CompCert [32], CompCertX [37], SepCompCert [13], CompCertM [31], and CompCertO [15].
 554 A great summary of the approaches to this problem (including the ones mentioned above),
 555 with their differences and similarities, is given by Patterson et al [27]. All of these use fixed
 556 languages but are reasonable ways of giving semantics to relevant IRs in `LeanMLIR(X)`. Our
 557 semantics is denotational and can be executed, like interaction trees [39, 40].

558 Our work differs from prior work on formalizing peephole rewrites by providing a framework
 559 for reasoning about SSA peephole rewrites. The closest similar work, Peek [24] defines
 560 peephole rewriting over an assembly instruction set. Their rewriter expects instructions to
 561 be adjacent to one another. Furthermore, their rewriter restricts source and target patterns
 562 to be of the same length, filling in the different lengths with `nop` instructions. Their patterns
 563 permit side effects, which we disallow since we are interested in higher-level, pure rewrites.
 564 Our patterns provide more flexibility since the source and target patterns are arbitrary
 565 programs, and are matched on sub-DAGs instead of a linear sequence.

566 8 Conclusion

567 Peephole rewrites represent a large and important class of compiler optimizations. We have
 568 seen how domain-specific IRs in SSA with regions greatly extend the scope of these peephole
 569 rewrites. They raise the level of abstraction both syntactically with def-use chains and

570 nesting, and semantically, with domain-specific abstractions. We have shown how to reason
571 effectively about such SSA-based compilers, and, specifically, local reasoning in the form
572 of peephole rewrites. We advocate building on top of a proof assistant with a small TCB,
573 an expressive language and a large library of mathematics. This increases the confidence
574 in our verification and extends its applicability to many domains where more specialized
575 methods don't exist. We also advocate proof automation and an intrinsically well-typed
576 mechanized core that can be designed to focus on the semantics of the domain. We incarnate
577 these principles in `LeanMLIR(X)`, a framework built on Lean and `mathlib` to reason about
578 domain-specific IRs in SSA with regions. We show how `LeanMLIR(X)` is simple to use,
579 amenable to automation, and effective for verifying IRs over complex domains.

580 ——— References ———

- 581 1 AV Aho, R Sethi, and JD Ullman. Compilers: Principles, techniques, and tools. 1985.
- 582 2 Andrew W Appel. SSA is functional programming. *Acm Sigplan Notices*, 33(4):17–20, 1998.
- 583 3 Seongwon Bang, Seunghyeon Nam, Inwhan Chun, Ho Young Jhoo, and Juneyoung Lee. SMT-
584 based translation validation for machine learning compiler. In Sharon Shoham and Yakir
585 Vizel, editors, *Computer Aided Verification - 34th International Conference, CAV 2022, Haifa,*
586 *Israel, August 7-10, 2022, Proceedings, Part II*, volume 13372 of *Lecture Notes in Computer*
587 *Science*, pages 386–407. Springer, 2022. doi:10.1007/978-3-031-13188-2_19.
- 588 4 Clark Barrett, Aaron Stump, Cesare Tinelli, et al. The smt-lib standard: Version 2.0. In
589 *Proceedings of the 8th international workshop on satisfiability modulo theories (Edinburgh,*
590 *UK)*, volume 13, page 14, 2010.
- 591 5 Gilles Barthe, Delphine Demange, and David Pichardie. Formal verification of an SSA-based
592 middle-end for CompCert. *ACM Transactions on Programming Languages and Systems*
593 *(TOPLAS)*, 36(1):1–35, 2014.
- 594 6 Siddharth Bhat and Tobias Grosser. Lambda the ultimate ssa: optimizing functional programs
595 in ssa. In *2022 IEEE/ACM International Symposium on Code Generation and Optimization*
596 *(CGO)*, pages 1–11. IEEE, 2022.
- 597 7 Leonardo De Moura and Nikolaj Bjørner. Z3: An efficient SMT solver. In *Tools and Algorithms*
598 *for the Construction and Analysis of Systems: 14th International Conference, TACAS 2008,*
599 *Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS*
600 *2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings 14*, pages 337–340. Springer,
601 2008.
- 602 8 Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming
603 language. In *International Conference on Automated Deduction*, pages 625–635. Springer,
604 2021.
- 605 9 Craig Gentry. *A fully homomorphic encryption scheme*. Stanford university, 2009.
- 606 10 Benjamin Grégoire and Assia Mahboubi. Proving equalities in a commutative ring done
607 right in Coq. In Joe Hurd and Thomas F. Melham, editors, *Theorem Proving in Higher*
608 *Order Logics, 18th International Conference, TPHOLs 2005, Oxford, UK, August 22-25, 2005,*
609 *Proceedings*, volume 3603 of *Lecture Notes in Computer Science*, pages 98–113. Springer, 2005.
610 doi:10.1007/11541868_7.
- 611 11 Gérard Huet. The zipper. *Journal of functional programming*, 7(5):549–554, 1997.
- 612 12 Petter Källström and Oscar Gustafsson. Fast and area efficient adder for wide data in recent
613 Xilinx FPGAs. In *2016 26th International Conference on Field Programmable Logic and*
614 *Applications (FPL)*, pages 1–4. IEEE, 2016.
- 615 13 Jeehoon Kang, Yoonseung Kim, Chung-Kil Hur, Derek Dreyer, and Viktor Vafeiadis.
616 Lightweight verification of separate compilation. In *Proceedings of the 43rd Annual ACM*
617 *SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 178–190,
618 2016.

- 619 **14** Richard A Kelsey. A correspondence between continuation passing style and static single
620 assignment form. *ACM SIGPLAN Notices*, 30(3):13–22, 1995.
- 621 **15** Jérémie Koenig and Zhong Shao. CompCertO: compiling certified open C components. In
622 *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language
623 Design and Implementation*, pages 1095–1109, 2021.
- 624 **16** Chris Lattner and Vikram Adve. LLVM: A compilation framework for lifelong program analysis
625 & transformation. In *International Symposium on Code Generation and Optimization, 2004.
626 CGO 2004.*, pages 75–86. IEEE, 2004.
- 627 **17** Chris Lattner, Jacques Pienaar, Mehdi Amini, Uday Bondhugula, River Riddle, Albert Cohen,
628 Tatiana Shpeisman, Andy Davis, Nicolas Vasilache, and Oleksandr Zinenko. MLIR: A compiler
629 infrastructure for the end of Moore’s law. *arXiv preprint arXiv:2002.11054*, 2020.
- 630 **18** Juneyoung Lee, Chung-Kil Hur, and Nuno P Lopes. AliveInLean: a verified LLVM peephole
631 optimization verifier. In *International Conference on Computer Aided Verification*, pages
632 445–455. Springer, 2019.
- 633 **19** Xavier Leroy, Sandrine Blazy, Daniel Kästner, Bernhard Schommer, Markus Pister, and
634 Christian Ferdinand. CompCert-a formally verified optimizing compiler. In *ERTS 2016:
635 Embedded Real Time Software and Systems, 8th European Congress*, 2016.
- 636 **20** Nuno P Lopes, Juneyoung Lee, Chung-Kil Hur, Zhengyang Liu, and John Regehr. Alive2:
637 bounded translation validation for LLVM. In *Proceedings of the 42nd ACM SIGPLAN
638 International Conference on Programming Language Design and Implementation*, pages 65–79,
639 2021.
- 640 **21** Nuno P Lopes, David Menendez, Santosh Nagarakatte, and John Regehr. Provably correct
641 peephole optimizations with alive. In *Proceedings of the 36th ACM SIGPLAN Conference on
642 Programming Language Design and Implementation*, pages 22–32, 2015.
- 643 **22** William M McKeeman. Peephole optimization. *Communications of the ACM*, 8(7):443–444,
644 1965.
- 645 **23** Manasij Mukherjee and John Regehr. Hydra: Generalizing peephole optimizations with
646 program synthesis. *Proceedings of the ACM on Programming Languages*, (OOPSLA), 2024.
- 647 **24** Eric Mullen, Daryl Zuniga, Zachary Tatlock, and Dan Grossman. Verified peephole optimiza-
648 tions for CompCert. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming
649 Language Design and Implementation*, pages 448–461, 2016.
- 650 **25** Georg Neis, Chung-Kil Hur, Jan-Oliver Kaiser, Craig McLaughlin, Derek Dreyer, and Viktor
651 Vafeiadis. Pilsner: A compositionally verified compiler for a higher-order imperative lan-
652 guage. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional
653 Programming*, pages 166–178, 2015.
- 654 **26** Sunjae Park, Woosung Song, Seunghyeon Nam, Hyeongyu Kim, Junbum Shin, and Juneyoung
655 Lee. HEaaN.MLIR: An optimizing compiler for fast ring-based homomorphic encryption.
656 *Proceedings of the 44th ACM SIGPLAN Conference on Programming Language Design and
657 Implementation*, 2023.
- 658 **27** Daniel Patterson and Amal Ahmed. The next 700 compiler correctness theorems (functional
659 pearl). *Proceedings of the ACM on Programming Languages*, 3(ICFP):1–29, 2019.
- 660 **28** Anurudh Peduri, Siddharth Bhat, and Tobias Grosser. QSSA: an SSA-based IR for quantum
661 computing. In *Proceedings of the 31st ACM SIGPLAN International Conference on Compiler
662 Construction*, pages 2–14, 2022.
- 663 **29** Benjamin C Pierce and C Benjamin. *Types and programming languages*. MIT press, 2002.
- 664 **30** Fabrice Rastello and Florent Bouchez Tichadou. *SSA-based Compiler Design*. Springer Nature,
665 2022.
- 666 **31** Youngju Song, Minki Cho, Dongjoo Kim, Yonghyun Kim, Jeehoon Kang, and Chung-Kil
667 Hur. CompCertM: CompCert with C-assembly linking and lightweight modular verification.
668 *Proceedings of the ACM on Programming Languages*, 4(POPL):1–31, 2019.

- 669 32 Gordon Stewart, Lennart Beringer, Santiago Cuellar, and Andrew W Appel. Compositional
670 CompCert. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on*
671 *Principles of Programming Languages*, pages 275–287, 2015.
- 672 33 Sebastian Ullrich and Leonardo de Moura. Beyond notations: Hygienic macro expansion for
673 theorem proving languages. In *International Joint Conference on Automated Reasoning*, pages
674 167–182. Springer, 2020.
- 675 34 Nicolas Vasilache, Oleksandr Zinenko, Aart J. C. Bik, Mahesh Ravishankar, Thomas Raoux,
676 Alexander Belyaev, Matthias Springer, Tobias Gysi, Diego Caballero, Stephan Herhut, Stella
677 Lorenzo, and Albert Cohen. Composable and modular code generation in MLIR: A structured
678 and retargetable approach to tensor compiler construction. *CoRR*, abs/2202.03293, 2022.
679 URL: <https://arxiv.org/abs/2202.03293>, arXiv:2202.03293.
- 680 35 Alexander Viand, Patrick Jattke, Miro Haller, and Anwar Hithnawi. HECO: Automatic code
681 optimizations for efficient fully homomorphic encryption. *arXiv preprint arXiv:2202.01649*,
682 2022.
- 683 36 Wei Wang and Xinming Huang. A novel fast modular multiplier architecture for 8,192-bit
684 RSA cryposystem. In *2013 IEEE High Performance Extreme Computing Conference (HPEC)*,
685 pages 1–5. IEEE, 2013.
- 686 37 Yuting Wang, Pierre Wilke, and Zhong Shao. An abstract stack based approach to veri-
687 fied compositional compilation to machine code. *Proceedings of the ACM on Programming*
688 *Languages*, 3(POPL):1–30, 2019.
- 689 38 Dominik Winterer, Chengyu Zhang, and Zhendong Su. Validating SMT solvers via semantic
690 fusion. In Alastair F. Donaldson and Emina Torlak, editors, *Proceedings of the 41st ACM*
691 *SIGPLAN International Conference on Programming Language Design and Implementation,*
692 *PLDI 2020, London, UK, June 15-20, 2020*, pages 718–730. ACM, 2020. doi:10.1145/
693 3385412.3385985.
- 694 39 Li-yao Xia, Yannick Zakowski, Paul He, Chung-Kil Hur, Gregory Malecha, Benjamin C Pierce,
695 and Steve Zdancewic. Interaction trees: representing recursive and impure programs in Coq.
696 *arXiv preprint arXiv:1906.00046*, 2019.
- 697 40 Yannick Zakowski, Calvin Beck, Irene Yoon, Ilia Zaichuk, Vadim Zaliva, and Steve Zdancewic.
698 Modular, compositional, and executable formal semantics for LLVM IR. *Proc. ACM Program.*
699 *Lang.*, 2021. URL: <https://doi.org/10.1145/3473572>.
- 700 41 Jianzhou Zhao, Santosh Nagarakatte, Milo MK Martin, and Steve Zdancewic. Formal
701 verification of SSA-based optimizations for LLVM. In *Proceedings of the 34th ACM SIGPLAN*
702 *conference on Programming language design and implementation*, pages 175–186, 2013.