CRELLVM: Verified Credible Compilation for LLVM

Jeehoon Kang* Yoonseung Kim* Youngju Song*
{jeehoon.kang, yoonseung.kim, youngju.song}@sf.snu.ac.kr
Seoul National University, Korea

Juneyoung Lee Sanghoon Park
Mark Dongyeon Shin Yonghyun Kim
{juneyoung.lee, sanghoon.park}@sf.snu.ac.kr
{dongyeon.shin, yonghyun.kim}@sf.snu.ac.kr
Seoul National University, Korea

Sungkeun Cho
skcho@ropas.snu.ac.kr
Seoul National University, Korea

Joonwon Choi
joonwonc@mit.edu
MIT CSAIL, USA

Chung-Kil Hur†
gil.hur@sf.snu.ac.kr

Kwangkeun Yi
kwang@ropas.snu.ac.kr
Seoul National University, Korea

Abstract

Production compilers such as GCC and LLVM are large complex software systems, for which achieving a high level of reliability is hard. Although testing is an effective method for finding bugs, it alone cannot guarantee a high level of reliability. To provide a higher level of reliability, many approaches that examine compilers' internal logics have been proposed. However, none of them have been successfully applied to major optimizations of production compilers.

This paper presents CRELLVM: a verified credible compilation framework for LLVM, which can be used as a systematic way of providing a high level of reliability for major optimizations in LLVM. Specifically, we augment an LLVM optimizer to generate translation results together with their correctness proofs, which can then be checked by a proof checker formally verified in Coq. As case studies, we applied our approach to two major optimizations of LLVM: register promotion (`mem2reg`) and global value numbering (`gvn`), having found four new miscompilation bugs (two in each).

CCS Concepts • **Theory of computation → Hoare logic**; • **Software and its engineering → Compilers**; **Formal software verification**;

Keywords LLVM, Coq, credible compilation, translation validation, compiler verification, relational Hoare logic

* The first three authors contributed equally to this work and are listed alphabetically.
† Hur is the corresponding author.

PLDI'18, June 18–22, 2018, Philadelphia, PA, USA

ACM ISBN 978-1-4503-5698-5/18/06…$15.00
https://doi.org/10.1145/3192366.3192377

ACM Reference Format:
Jeehoon Kang, Yoonseung Kim, Youngju Song, Juneyoung Lee, Sanghoon Park, Mark Dongyeon Shin, Yonghyun Kim, Sungkeun Cho, Joonwon Choi, Chung-Kil Hur, and Kwangkeun Yi. 2018. CRELLVM: Verified Credible Compilation for LLVM. In *Proceedings of 39th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'18).* ACM, New York, NY, USA, 28 pages. https://doi.org/10.1145/3192366.3192377

1 Introduction

Production compilers such as GCC and LLVM are large complex software systems, for which achieving a high level of reliability is hard. Their complexity comes in two fold. First, to generate efficient target code, they perform various complex optimizations. Second, to consume less time and memory during compilation, they are usually written in C/C++ using sophisticated data structures. Due to such complexity, it is hard to make mainstream compilers very reliable.

Although testing is an effective method for finding bugs, that alone hardly guarantees a high level of reliability. Recent random testing tools such as CSmith [53] and EMI [24] have shown their effectiveness by finding hundreds of bugs in GCC and LLVM. However, they missed bugs in the `gvn` and `mem2reg` passes of LLVM, which we discovered later (see §1.2 for details), since they treat compilers as black boxes without examining their internal logics.

In order to provide a higher level of reliability, many approaches that examine compilers' internal logics have been proposed, none of which, however, have been successfully applied to major optimizations of production compilers. For example, while compiler verification techniques have been applied to compilers such as CompCert [26] to guarantee their formal correctness, this approach is not readily applicable to production compilers since it requires compilers to be written in the language of a proof assistant such as Coq. As another example, Alive [30] is a domain-specific language (DSL) in which one can manually write a compiler's optimization logic and automatically verify its correctness or else generate a counterexample. Though this approach has been successfully applied to LLVM, its application is

limited to peephole optimizations because it is hard to faithfully translate the implementation of complex optimizations into Alive and, more importantly, Alive does not support cyclic control flows such as loop. As the last example, the credible compilation [16, 33, 34, 44] and verified translation validation [14, 19, 43, 50–52] approaches augment compilers to generate translation results together with their correctness proofs, which can then be checked by a (verified) proof checker. Since a correctness proof is generated and checked at each compilation time, it provides a formal correctness guarantee for the particular translation or else finds a bug (either in the compiler code or in the proof-generation code). However, there has been only a preliminary attempt to apply these approaches to production compilers so far. (See §9 for detailed comparison.)

This paper presents CRELLVM: a verified credible-compilation framework for LLVM, which can be used as a systematic way of providing a high level of reliability for major optimizations in LLVM. Specifically:

1. We design and develop a logic and its proof checker for reasoning about LLVM optimizations, called Extensible Relational Hoare Logic (ERHL), in the proof assistant Coq. This logic's novelty lies in its representation of relational predicates as mostly unary predicates (see §2.2 for details).
2. We fully verify a semantics-preservation result for our proof checker in the style of CompCert using the Coq formalization of LLVM IR (Intermediate Representation) from the VELLVM project [55].
3. As case studies, we wrote proof-generation codes (213 and 440 SLOC[1] in C++) for two major optimizations: register promotion in the `mem2reg` pass and global value numbering (GVN) with partial-redundancy elimination (PRE) in the `gvn` pass. Then we performed validation of the two optimizations for standard benchmarks, five large open-source projects and test files randomly generated by CSmith.
4. As a result, we found four new miscompilation bugs (two in each optimization). It is notable that all the four bugs had been hidden for 7-8 years until we found them.

1.1 Overview of CRELLVM

Framework The framework of CRELLVM works as follows. First, as shown in Fig. 1, we separate the compilation and validation phases. For compilation, as depicted in the left side of Fig. 1, we use the original optimizer to translate the source IR code `src.ll` to the target IR code `tgt.ll`. After the compilation, we can conduct validation, as depicted in the right side of Fig. 1. For this, we first run the optimizer extended with a proof-generation code that produces the target `tgt'.ll` together with the proof `Proof`. Then the proof checker validates `Proof` to see whether `src.ll` is correctly translated to `tgt'.ll`. If the validation fails, we can see a

[1]SLOC stands for significant lines of code *i.e.,* ignoring spaces and comments.

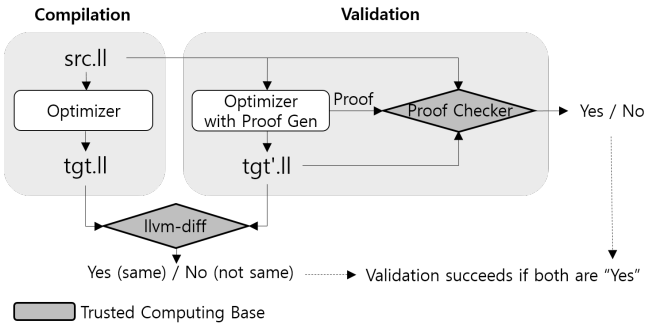

Figure 1. The CRELLVM Framework

logical reason for the failure, with which we can find a bug either in the compiler or in the proof-generation code. If the validation succeeds, we finally compare `tgt.ll` and `tgt'.ll` using the LLVM IR comparison tool `llvm-diff`.

There are two points to note about the framework. First, `llvm-diff` essentially performs alpha-equivalence checking, which is necessary because while `tgt.ll` may have unnamed IR registers, `tgt'.ll` has explicit names for all registers for proof-generation purposes. Second, since we just add proof-generation code without modifying existing compiler code except for giving names to unnamed registers, the original and proof-generating compilers are expected to generate alpha-equivalent programs, which is always checked using `llvm-diff` as described above. Therefore, programmers can use the original compiler in regular usage and then run the proof-generating one on occasion to check correctness because the former is much faster than the latter. On the other hand, compiler developers can use the latter for testing on regular basis to find bugs.

ERHL and Proof Checker For validation in CRELLVM, we develop ERHL, which is a variant of relational Hoare logic [16] specialized for LLVM IR. The logic and its proof checker is *extensible* because (*i*) the logic can be extended with any custom inference rules and (*ii*) the proof checker can be extended with any custom automation functions that try to fill in the gaps in incomplete proofs by automatically finding appropriate inference rules, like the `auto` tactic in Coq.

Verification of Proof Checker In the CRELLVM framework, the TCB (Trusted Computing Base) includes only the proof checker, the equality checker (`llvm-diff`) and custom inference rules. In particular, the proof-generation code in the compiler is not a part of the TCB because any incorrect proof would be invalidated by the proof checker.

We further remove the proof checker and inference rules from the TCB by implementing and verifying them in Coq. Though we currently use the (unverified) standard `llvm-diff` tool for comparing IR programs, it would also be possible to implement and verify it in Coq.

Note that verification of the proof checker and inference rules matters in practice. First, we found various corner-case

bugs in our proof checker during its verification. Second, we also found one of our two mem2reg bugs [9] during the verification of inference rules. See the example below.

```
p := alloca(); r := *p
foo(r)                          ⇝ foo(1 / ((int)G - (int)G))
*p := 1 / ((int)G - (int)G)
```

Here G is the constant address of a global variable.

To see why this translation is incorrect, suppose that the function foo(r) ignores r and repeatedly prints out 0 without returning to the caller. Then division-by-zero never happens in the source program, while it does in the target. The problem here is that the mem2reg pass assumes that constant expressions never raise any undefined behavior such as division-by-zero, which is not true since 1 / ((int)G - (int)G) forms a valid constant expression in LLVM. Following the logic of mem2reg, we also added such a custom inference rule, which we found unsound during the verification of the rule.

It is important to note that all the programs in this paper represent LLVM IR programs and we just use C syntax to help with understanding. For example, the source program in the above transformation is undefined as a C program but well-defined as an IR program. Thus, the transformation is only unsound as an IR-to-IR transformation. The LLVM community considers such an IR-to-IR miscompilation as a definite bug even when it does not cause any C-to-Assembly miscompilation since it can potentially cause an end-to-end miscompilation for other source languages such as Swift and Rust.

Results We wrote proof-generation codes for register promotion in the mem2reg pass and for GVN-PRE in the gvn pass; and also partly for loop-invariant code motion in the licm pass, and 139 micro-optimizations in the instcombine pass in order to demonstrate the generality of ERHL. We then conducted validation of the optimizations for the SPEC CINT2006 C Benchmarks [15], LLVM nightly test suite, and five open-source projects: sendmail, emacs, python, gimp, and ghostscript, in total 5.3 million LOC in C. As a result, we found four new miscompilation bugs.

We present the details of mem2reg validation in §3 and gvn validation in [1, §C].

1.2 Advantages of CRELLVM over Testing

CRELLVM checks whether optimizations are performed by *correct reasoning*, while testing simply checks *results* of the test programs. This can make a difference as follows.

First, an optimization performed by incorrect reasoning may still be correct for most programs including all the test programs. In this case, testing cannot uncover the bug, while CRELLVM can because it checks the underlying reasoning. For example, we found our first mem2reg bug [5] in this situation.

Specifically, the following optimization shows the bug.

```
p := alloca()
loop {                              loop {
    r := *p; foo(r); *p := 42  ⇝        foo(undef)
}                                   }
```

This translation is incorrect because only in the first iteration of the loop is r undef[2]; in the remaining iterations r is 42 according to the semantics of LLVM. The mem2reg pass performs this due to faulty reasoning.

However, this faulty reasoning is often not visible in the final compiled program. The reason is that, since the input to foo is sometimes undefined, for foo to be well behaved it often ignores its input r (*e.g.,* by using an operation like r & 0x0). Thus this transformation is actually correct in such a program since the value of r is never used in the program. Indeed, the SPEC benchmark that provoked this faulty reasoning behaved this way, and so the faulty reasoning never led to a faulty program, which is why the bug had been hidden for such a long time.

The fact that the faulty reasoning was inconsequential in this case does not mean the bug is unimportant. As we said before, the LLVM community cares about such an IR-to-IR miscompilation and immediately fixed the bug after we reported it. Moreover, visible miscompilations due to the bug could happen in a realistic situation (see [1, §B] for a concrete example).

Second, a potential flaw introduced by miscompilation may not be exploited by the current compiler and silently disappear during the compilation. Also in this case, CRELLVM can detect the bug because it checks the underlying reasoning. For example, we found the two gvn bugs [6, 7] in this situation, which had not been found for 8 years. Note that the two bugs are caused by the same reason but we counted them as two because they appear in two separate places.

Specifically, the following optimization shows the bug.

```
q1 := (p + 10) inbounds          q1 := (p + 10) inbounds
q2 := (p + 10)              ⇝
bar(q1, q2)                      bar(q1, q1)
```

In the source program, (p + 10) inbounds[3] is defined to be undef[4] when the index 10 is out of the bounds of p, while (p + 10) is always defined to be the computed address. Thus replacing q2 with q1 introduces more undefinedness, which is incorrect because it can be potentially exploited by subsequent optimizations. However, so far the LLVM compiler has not exploited such undefinedness, thereby causing no observable misbehaviors. Indeed this miscompilation happened many times during validation of the standard benchmarks but testing has failed to detect it.

[2] Since *p is uninitialized, it contains undef, which is a special value representing undefinedness

[3] This denotes the GetElementPtr (GEP) operation.

[4] Technically, it is defined to be poison but the difference does not matter here.

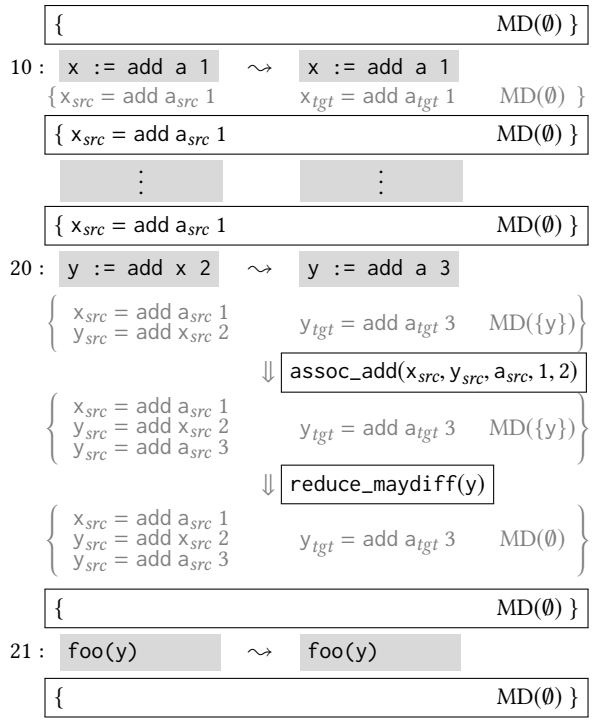

Figure 2. Validation of an `assoc-add` translation in ERHL

2 Overview

In this section, we give a more detailed overview of how CRELLVM works using the `assoc-add` optimization of the `instcombine` pass as a motivating example.

2.1 Translation Example

We first give an example translation of the `assoc-add` optimization, which is shown in the shaded part of Fig. 2. Here `y := add x 2` is replaced by `y := add a 3` at line 20. This translation can be beneficial because after it, the register `x` may no longer be used and thus `x := add a 1` at line 10 may be eliminated later. This translation is also sound because (*i*) the *assertion* "`x = add a 1`" holds throughout lines 10-20, since the registers `a` and `x` are not redefined between line 10 and 20 thanks to the Static Single Assignment (SSA) property [18][5]; and (*ii*) from this, we can *infer* that `add x 2 = add (add a 1) 2 = add a 3` holds at line 20.

2.2 Proof Validation

We now construct a proof for the `assoc-add` translation example and validate it in ERHL.

ERHL Proof A formal proof of the translation is given in the box of Fig. 2. Specifically, the proof consists of a set of assertions and a list of inference rules at each line. For example, at line 20, the set of assertions is { MD(∅) } and the list of inference rules is (`assoc_add`$(x_{src}, y_{src}, a_{src}, 1, 2)$, `reduce_maydiff(y)`).

This ERHL proof captures the assertion and the inference step of the intuitive reasoning above. First, the assertion MD(∅) at every line states that every register contains the same value in the source and target program states. Second, the additional assertion $x_{src} = \texttt{add}\ a_{src}\ 1$ between line 10 and line 20 states that in the source state, the value of the register `x` is equal to the result of `add a 1`. Finally, the inference rules `assoc_add`$(x_{src}, y_{src}, a_{src}, 1, 2)$ and `reduce_maydiff(y)` at line 20 are those that need to be applied for validation at line 20. The details of the rules will be given later when we discuss the validation process.

ERHL Assertions Before we proceed to the validation of the proof, we discuss ERHL assertions in more details. An ERHL assertion is a triple (S, T, M), where S is a set of assertions that should hold for the source state; T is for the target state; and M is an assertion relating the source and target states.

First, the source and target assertions, S and T, can contain various forms of predicates. For example, $x_{src} = \texttt{add}\ a_{src}\ 1$ is a source assertion and $x_{tgt} = \texttt{add}\ a_{tgt}\ 3$ is a target assertion. Here and henceforth, x_{src} and x_{tgt} represent the values of the register x in the source and target states, respectively. Though we only use the equality predicate for `assoc-add`, we will introduce various other predicates later. It is important to note that we do not allow arbitrary assertions relating the source and target states such as $x_{src} = y_{tgt} + 1$.

Second, the relational assertion M is a set of registers, called the *maydiff set*, that may contain different values in the source and target states. In other words, all the registers not in M should have the same value in the source and target states, which we denote by MD(M):

$$\text{MD}(M) \iff \forall x \notin M.\ x_{src} = x_{tgt}\ .$$

Note that the maydiff set is the only form of relational assertion relating the source and target states.

Finally, every ERHL assertion implicitly requires the public parts of the source and target memories to be equivalent. More precisely, we use the CompCert-style memory-injection relation [28]. Later we introduce predicates that allow private memory allocations that do not belong to the public part of memory (see §3.2).

The main novelty of ERHL assertions is that we can use the standard algorithm of (unary) Hoare logic to compute post relational assertions, because ERHL assertions are mainly unary (*i.e.*, only for the source state, or for the target state, not relating them) except for the maydiff set. This unary nature greatly simplifies the ERHL proof checker and its correctness proof. Though mainly unary, ERHL assertions can indirectly encode general forms of relational assertions (see §3.2 for details).

[5] The SSA property says that for every used register x, there is statically (*i.e.*, syntactically) exactly one instruction that defines x (*i.e.*, assigning a value to x), which moreover comes before every use of x.

Proof Validation The gray text in Fig. 2 shows the validation process performed by the ERHL proof checker, which proceeds as follows.

First, the proof checker checks that the initial assertion holds for all possible initial states. It accepts the initial assertion $\{\,\mathrm{MD}(\emptyset)\,\}$ in Fig. 2 since the source and target states are initially equivalent.

Second, the proof checker checks whether the Hoare triple $\{P\}\ I_{src} \sim I_{tgt}\ \{Q\}$ at each line is valid. This means that the assertion Q after the line holds for all program states resulted by executing the source and target instructions I_{src} and I_{tgt} at the line under any program states satisfying the assertion P before the line. In Fig. 2, we only explain validations at lines 10 and 20 in detail because the others are trivial.

At line 10, the proof checker first computes a strong post-assertion, $\{\,\mathsf{x}_{src} = \mathsf{add}\ \mathsf{a}_{src}\ 1, \mathsf{x}_{tgt} = \mathsf{add}\ \mathsf{a}_{tgt}\ 1, \mathrm{MD}(\emptyset)\,\}$, using our post-assertion computation algorithm. Here, the algorithm simply adds the equality predicates corresponding to the executed instructions. Then, the assertion after line 10, $\boxed{\{\,\mathsf{x}_{src} = \mathsf{add}\ \mathsf{a}_{src}\ 1, \mathrm{MD}(\emptyset)\,\}}$, follows from the computed strong post-assertion by a simple inclusion check.

At line 20, the checker also computes a strong post-assertion, $\{\,\mathsf{x}_{src} = \mathsf{add}\ \mathsf{a}_{src}\ 1, \mathsf{y}_{src} = \mathsf{add}\ \mathsf{x}_{src}\ 2, \mathsf{y}_{tgt} = \mathsf{add}\ \mathsf{a}_{tgt}\ 3, \mathrm{MD}(\mathsf{y})\,\}$. Here, the post-assertion computation adds the equality predicates corresponding to the executed instructions and also adds the register y to the maydiff set because the executed source and target instructions are not identical. Then, the proof checker applies the inference rules given by the proof. The rule $\mathtt{assoc\_add}(\mathsf{x}_{src}, \mathsf{y}_{src}, \mathsf{a}_{src}, 1, 2)$ derives $\mathsf{y}_{src} = \mathsf{add}\ \mathsf{a}_{src}\ 3$ from $\mathsf{x}_{src} = \mathsf{add}\ \mathsf{a}_{src}\ 1$ and $\mathsf{y}_{src} = \mathsf{add}\ \mathsf{x}_{src}\ 2$ by associativity:

$$\frac{(\mathtt{assoc\_add}(x, y, a, C_1, C_2))\quad x = \mathsf{add}\ a\ C_1 \qquad y = \mathsf{add}\ x\ C_2 \qquad C = C_1 + C_2}{\mathsf{add}\ \{\, y = \mathsf{add}\ a\ C\,\}}$$

The rule $\mathtt{reduce\_maydiff}(\mathsf{y})$ removes the register y from the maydiff set because $\mathsf{y}_{src} = \mathsf{add}\ \mathsf{a}_{src}\ 3$, $\mathsf{y}_{tgt} = \mathsf{add}\ \mathsf{a}_{tgt}\ 3$ and a is not in the maydiff set:

$$\frac{(\mathtt{reduce\_maydiff}(y, e))\quad y_{src} = e_{src} \qquad e_{tgt} = y_{tgt} \qquad \text{no registers in } e \text{ are in the maydiff set}}{\text{remove } y \text{ from the maydiff set}}$$

Then, the assertion after line 20, $\{\,\mathrm{MD}(\emptyset)\,\}$, easily follows by a simple inclusion check.

Finally, the proof checker checks whether the same observable events (*i.e.,* the same sequence of system calls) are produced at each line. It is the case in Fig. 2 because at line 20, no observable events are produced; and at the other lines, the source and target instructions are identical and the maydiff sets are empty implying that the source and target states are equivalent. In particular, at line 21, the proof checker explicitly checks that the same value is passed to the function `foo` because the function may produce observable events.

Algorithm 1 AssocAdd(F: Function)

```
A1: for l2: y := add (reg x) (const C2) in F do
A2:   if FindDef(F, x) is l1: x := add (reg a) (const C1) then
A3:     C := Simplify(add C1 C2)
A4:     ReplaceAt(F, l2, y := add (reg a) (const C))
A5:     [Assn(x_src = add a_src C1, l1, l2)]
A6:     [Inf(assoc_add(x_src, y_src, a_src, C1, C2), l2)]
A7:   end if
A8: end for
A9: [Auto(reduce_maydiff)]
```

2.3 Proof Generation

Now we explain how we generate proofs for `assoc-add`.

Algorithm Algorithm 1 shows the `assoc-add` optimization algorithm implemented in LLVM's `instcombine` pass, which is presented in a rather functional style for presentation purposes. Specifically, AssocAdd(F) optimizes each function definition F as follows (ignore the boxes for now, which are the proof-generation code).

[Line A1] Find an instruction of the form l_2: y := add x C_2 with C_2 constant. In Fig. 2, 20: y := add x 2 can be picked. **[Line A2]** Check if x is defined by an instruction of the form l_1: x := add a C_1 with C_1 constant. Here, FindDef(F, x) finds the instruction that defines the register x.[6] In Fig. 2, 10: x := add a 1 is picked. **[Lines A3-A4]** If it is the case, compute the constant $C = C_1 + C_2$ and replace the instruction at l_2 with y := add a C. In Fig. 2, the instruction at line 20 is replaced by y := add a 3.

Proof Generation Once we understand the `assoc-add` optimization algorithm, it is quite straightforward to write the proof-generation code given in the boxes of Algorithm 1.

[Line A5] Add the assertion x_{src} = add a_{src} C_1 at every line between l_1 and l_2. In Fig. 2, the assertion $\mathsf{x}_{src} = \mathsf{add}\ \mathsf{a}_{src}\ 1$ is added at every line between 10 and 20. **[Line A6]** Add the inference rule $\mathtt{assoc\_add}(x_{src}, y_{src}, a_{src}, C_1, C_2)$ at line l_2. In Fig. 2, the rule $\mathtt{assoc\_add}(\mathsf{x}_{src}, \mathsf{y}_{src}, \mathsf{a}_{src}, 1, 2)$ is added at line 20. **[Line A9]** Enable the custom automation function named `reduce_maydiff`, which tries to find and insert appropriate `reduce_maydiff` rules when necessary. In Fig. 2, it figures out that $\mathtt{reduce\_maydiff}(\mathsf{y})$ is needed at line 20.

Automation An automation function works as follows. When it remains to prove Q implies Q', the designated automation function examines the assertions Q and Q' and tries to find a sequence of inference rules that derives Q' from Q. For example, at line 20 in Fig. 2, after applying the `assoc_add` rule it remains to prove $Q = \{\,\mathsf{x}_{src} = \mathsf{add}\ \mathsf{a}_{src}\ 1, \mathsf{y}_{src} = \mathsf{add}\ \mathsf{x}_{src}\ 2, \mathsf{y}_{src} = \mathsf{add}\ \mathsf{a}_{src}\ 3, \mathsf{y}_{tgt} = \mathsf{add}\ \mathsf{a}_{tgt}\ 3, \mathrm{MD}(\mathsf{y})\,\}$ implies $Q' = \{\,\mathrm{MD}(\emptyset)\,\}$, from which the automation function finds the inference rule $\{\,\mathtt{reduce\_maydiff}(\mathsf{y})\,\}$.

[6] The instruction that defines x is unique thanks to the SSA property.

Automation functions can greatly simplify proof generation in certain cases. A good example is transitivity reasoning because it is much harder at proof generation time than at validation time. For instance, given a goal $x = y$, to prove it by transitivity, we have to figure out intermediate equations (*e.g.,* $x = a$, $a = b$, $b = y$). For this, at proof generation time, we have to write a code that (sometimes recursively) search through the compiler internal states, which is tightly coupled with the compiler code; while at validation time, since a concrete pre-assertion is given, we just need to search through the equations given in the pre-assertion, which is completely generic and can be easily automated.

It is important to note that automation functions do not need to be verified (*i.e.,* not a part of TCB) because all they do is to insert inference rules, which is a part of proof construction, not that of proof checking.

3 Register Promotion

Register-promotion optimization, the `mem2reg` pass of LLVM, transforms memory accesses to locally allocated memory locations into register accesses, provided that the memory location is only used for loads and stores (*i.e.,* never copied or escaped). This translation is important because register accesses are cheaper than memory accesses, and are subject to further optimizations.

The optimization also performs the SSA transformation so that the target program has the SSA property. This transformation is necessary because there can be statically multiple stores to a single location, and just transforming them to writes to a single register would break the SSA property.

In this section, we show how we generate and validate proofs for the `mem2reg` optimization.

3.1 Translation Example

The shaded part of Fig. 3 shows an example translation of the `mem2reg` optimization, where all memory accesses via `p` is promoted to register accesses to `p1` and uses of 42 and `x`. Note that `c`, `x`, and `q` are the function parameters.

More specifically, the allocation, load and store instructions to `p` are removed (ignore `lnop` for now), and every use of the result of a load from `p` is replaced by the value stored in `*p` at the time of the load. For example, in Fig. 3, the compiler figures out that `*p` contains 42 at line 20 (and so does the register `a`) due to the store of 42 in `*p` at line 11, and thus replaces the use of `a` with 42 at line 21. This translation is sound because (*i*) the assertion $*p_{src} = 42$ holds from line 11 to 20; and (*ii*) $a_{src} = 42$ holds from line 20 to 21. Note that we use the blue color for assertions about `*p` and the red color about the registers containing the value loaded from `*p`.

In a case where the value stored in `*p` depends on the control flow, the compiler inserts a *ϕ-node*, which is a unique construct in the SSA form and assigns different values to a register depending on the control flow. For example, at

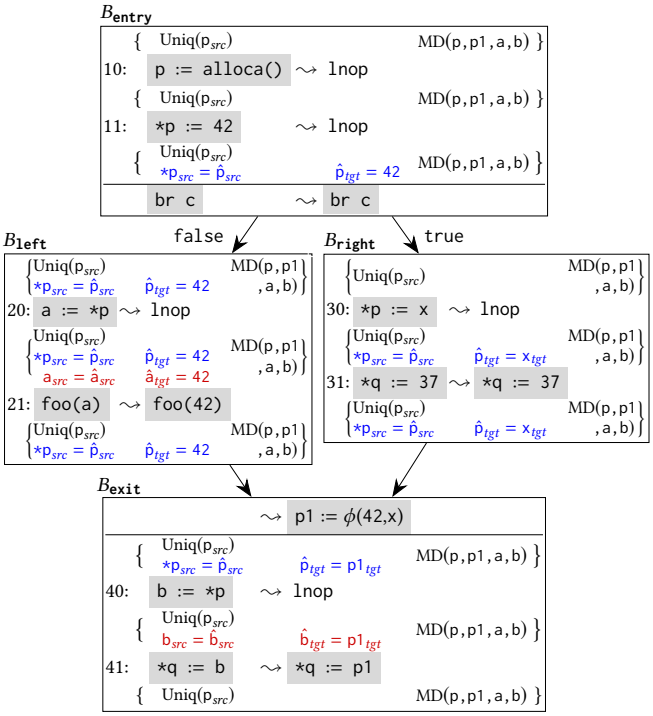

Figure 3. A register-promotion example

line 40, `*p` contains 42 if the control comes from $\mathbf{B}_{\texttt{left}}$, and `x` if it comes from $\mathbf{B}_{\texttt{right}}$. In this case, the compiler inserts a ϕ-node `p1 := φ(42,x)` at the beginning of $\mathbf{B}_{\texttt{exit}}$, which defines `p1` to be 42 when coming from $\mathbf{B}_{\texttt{left}}$ and `x` when coming from $\mathbf{B}_{\texttt{right}}$. Then, the use of the register `b` containing the loaded value from `*p` can be replaced by `p1` at line 41.

3.2 ERHL Proof

We show how to turn the intuition for soundness into a formal ERHL proof, which is given in the unshaded part of Fig. 3 including `lnop`. Here we omit the inference rules for simplicity, which will be shown later. We introduce interesting features of ERHL by explaining each part of the proof.

Logical No-Ops for Instruction Alignment Logical no-ops, denoted `lnop`, can be inserted as part of a proof in order to align source and target instructions when their alignment is broken by a translation. For example, in Fig. 3, `lnop` is inserted at lines 10, 11, 20, 30, 40 because the instructions there are removed by `mem2reg`.

Note that `lnop` is logical because it is absent from the real IR code and used only for validation purposes. During validation, it is interpreted as doing nothing (*i.e.,* no-op).

Ghost Registers for Relational Assertions For complex optimizations, we often need to state relational properties (*i.e.,* relating the source and target states) in a proof. For example, in Fig. 3, we need to state $*p_{src} = p1_{tgt}$ before line 40, which relates a value in the source ($*p_{src}$) with that in the target ($p1_{tgt}$).

Though not directly supported in ERHL, such relational properties can be encoded using *ghost registers*. Specifically, we can encode $e_{src} = e'_{tgt}$ using a fresh ghost register $\hat{g}$:

$$\{\, e_{src} = \hat{g}_{src},\ \hat{g}_{tgt} = e'_{tgt},\ \mathrm{MD}(M) \,\} \text{ with } \hat{g} \notin M$$

Since the ghost register $\hat{g}$ is not in the maydiff set, we have $\hat{g}_{src} = \hat{g}_{tgt}$, which, by transitivity, implies $e_{src} = e'_{tgt}$. For example, in Fig. 3, the assertion $\{*p_{src} = \hat{p}_{src}, \hat{p}_{tgt} = p1_{tgt}, \mathrm{MD}(p, p1, a, b)\}$ before line 40 effectively states $*p_{src} = p1_{tgt}$. Note that the ghost register $\hat{p}$ has nothing to do with the physical register p and we use $\hat{(-)}$ for ghost registers to distinguish them from the physical ones.

Ghost registers are *logical* ones that do not exist in physical program states. Instead, they are existentially quantified in the semantics of ERHL assertions. More specifically, a pair of source and target states $(\sigma_{src}, \sigma_{tgt})$ satisfies an ERHL assertion P, if there *exists* a pair of source and target ghost register files $(\hat{rs}_{src}, \hat{rs}_{tgt})$ such that the pair of σ_{src} extended with $\hat{rs}_{src}$ and σ_{tgt} extended with $\hat{rs}_{tgt}$ satisfies P.

Taking ghost registers into account, the proof in Fig. 3 has five relational assertions: $*p_{src} = 42_{tgt}$ between line 11 and the end of $\mathbf{B}_{\texttt{left}}$, $a_{src} = 42_{tgt}$ between line 20 and line 21, $*p_{src} = x_{tgt}$ between line 30 and the end of $\mathbf{B}_{\texttt{right}}$, $*p_{src} = p1_{tgt}$ between the beginning of $\mathbf{B}_{\texttt{exit}}$ and line 41, and $b_{src} = p1_{tgt}$ between line 40 and line 41. It is easy to see that these assertions correctly capture the relational properties caused by executing different instructions in the source and target.

Uniqueness Predicate for Isolation We can use the predicate Uniq in order to state that an address is completely isolated. For example, in Fig. 3, we have $\mathrm{Uniq}(p_{src})$ at every line. It means that in the source, if p contains an address ℓ, (i) ℓ is *not aliased* with any address stored in the other registers or in memory (*i.e.,* they point to disjoint memory blocks); and (ii) ℓ is *private* (*i.e.,* it is not in the public part of the memory injection) meaning that it has no corresponding equivalent address in the target. In other words, the address contained in p should point to a completely isolated block.

Note that ERHL also supports memory predicates weaker than $\mathrm{Uniq}(p)$: (i) the *privateness* predicate, $\mathrm{Priv}(p)$, which states that the address in p is private; and (ii) the *noalias* predicate, $p \perp q$, which states that the addresses in p and q point to disjoint memory blocks.

Maydiff Sets Finally, we have $\mathrm{MD}(\{\, p, p1, a, b \,\})$ at every line because these registers are removed or introduced so that they have different values in the source and target.

3.3 Proof Validation

We show how our proof checker validates the ERHL proof.

Entry The proof checker checks that the entry assertion, $\{\, \mathrm{Uniq}(p_{src}), \mathrm{MD}(\{\, p, p1, a, b \,\}) \,\}$, holds for initial states. It accepts the assertion $\mathrm{Uniq}(p_{src})$ since p is a local register and thus contains the `undef` value initially, which is not an address. It also accepts every maydiff set since the source and target registers initially contain equivalent values.

Allocation of the Promoted Location At line 10, the proof checker allows an allocation, `p := alloca()`, in the source and `lnop` in the target. In this case, it computes a post-assertion from the pre-assertion by (i) removing all assertions containing p_{src} because p_{src} is updated, (ii) adding $\{\, \mathrm{Uniq}(p_{src}), *p_{src} = \texttt{undef} \,\}$ because p contains a newly allocated address, and then (iii) adding p to the maydiff set. Thus, we have $\{\, \mathrm{Uniq}(p_{src}), *p_{src} = \texttt{undef}, \mathrm{MD}(\{\, p, p1, a, b \,\}) \,\}$, from which the assertion after line 10 trivially follows.

Stores to the Promoted Location At line 30 (and similarly at line 11), the proof checker allows a store, `*p := x`, in the source and `lnop` in the target because $*p_{src}$ is private (*i.e.,* has no corresponding target address) due to $\mathrm{Uniq}(p_{src})$ in the pre-assertion. In this case, it computes a post-assertion by (i) removing all and *only* the assertions containing $*p_{src}$ because $*p_{src}$ is updated and p_{src} has no alias with any other address due to $\mathrm{Uniq}(p_{src})$, and then (ii) adding $\{\, *p_{src} = x_{src} \,\}$. Thus, we have $\{\, \mathrm{Uniq}(p_{src}), *p_{src} = x_{src}, \mathrm{MD}(\{\, p, p1, a, b \,\}) \,\}$.

At this point, the proof gives the rule `intro_ghost`$(\hat{p}, x)$, which first makes $\hat{p}$ fresh by removing all assertions about $\hat{p}$ and removing $\hat{p}$ from the maydiff set and then adds $\{x_{src} = \hat{p}_{src}, \hat{p}_{tgt} = x_{tgt}\}$ when x is not in the maydiff set. Thus, we have $\{\mathrm{Uniq}(p_{src}), *p_{src} = x_{src}, x_{src} = \hat{p}_{src}, \hat{p}_{tgt} = x_{tgt}, \mathrm{MD}(\{p, p1, a, b\})\}$. Then, the proof gives the rule `transitivity`$(*p_{src}, x_{src}, \hat{p}_{src})$, which derives $*p_{src} = \hat{p}_{src}$ from $*p_{src} = x_{src}$ and $x_{src} = \hat{p}_{src}$. Then the assertion after line 30 trivially follows. (See [1, §I] for the definitions of `intro_ghost` and `transitivity`.)

ϕ-nodes At the ϕ-node of $\mathbf{B}_{\texttt{exit}}$, the proof checker validates the assertion separately for each incoming block. For the incoming block $\mathbf{B}_{\texttt{left}}$, the proof checker computes a post-assertion by (i) removing all assertions containing $p1_{tgt}$ because $p1_{tgt}$ is updated, (ii) adding $42 = p1_{tgt}$ because `p1 := 42` is executed in the target when control comes from $\mathbf{B}_{\texttt{left}}$, and then (iii) adding p1 to the maydiff set. Then the proof gives the inference rule `transitivity`$(\hat{p}_{tgt}, 42, p1_{tgt})$, which derives $\hat{p}_{tgt} = p1_{tgt}$, from which the assertion after the ϕ-node follows trivially. For the incoming block $\mathbf{B}_{\texttt{right}}$, validation succeeds similarly, where the proof gives the inference rule `transitivity`$(\hat{p}_{tgt}, x_{tgt}, p1_{tgt})$.

Note that for presentation purposes here we simplified the post-assertion computation for ϕ-nodes. ERHL performs a more general version to handle cyclic control flows (see §4).

Loads from the Promoted Location At line 40 (and similarly at line 20), the proof checker allows a load, `b := *p`, in the source and `lnop` in the target. In this case, it computes a post-assertion by (i) removing all assertions containing b_{src} because b_{src} is updated, (ii) adding $b_{src} = *p_{src}$ and then (iii) adding b to the maydiff set. Thus, we have $\{\mathrm{Uniq}(p_{src}), *p_{src} = \hat{p}_{src}, \hat{p}_{tgt} = p1_{tgt}, b_{src} = *p_{src}, \mathrm{MD}(\{p, p1, a, b\})\}$.

At this point, the proof gives the rule intro_ghost($\hat{b}, \hat{p}$), which adds $\{\hat{p}_{src} = \hat{b}_{src}, \hat{b}_{tgt} = \hat{p}_{tgt}\}$ because $\hat{p}$ is not in the maydiff set. Then the proof gives appropriate transitivity rules, which derives $b_{src} = *p_{src} = \hat{p}_{src} = \hat{b}_{src}$ and $\hat{b}_{tgt} = \hat{p}_{tgt} = p1_{tgt}$, from which the assertion after line 40 trivially follows.

Equivalence Checking At lines 21, 31 and 41, the proof checker checks that the behaviors of the source and target instructions are equivalent. Specifically, it checks that equivalent values are passed to the same function (at line 21) and stored at equivalent public locations (at lines 31,41) because these can be observed by other functions. These checks succeed thanks to the relational assertions ($\{ a_{src} = \hat{a}_{src}, \hat{a}_{tgt} = 42 \}$ at line 21, $\{ b_{src} = \hat{b}_{src}, \hat{b}_{tgt} = p1_{tgt} \}$ at line 41).

Alias Checking At lines 21, 31, and 41, the proof checker computes post-assertions using memory-alias information. In general, for a function call or store instruction, since it updates the public part of the memory, the proof checker removes all assertions about values stored in memory locations p (*i.e.,* those including $*p$) unless (*i*) p is in the private part of the memory (*i.e.,* Priv(p) or Uniq(p)), or (*ii*) p is not aliased with q (*i.e.,* $p \perp q$) in case $*q$ is updated by the store instruction. At lines 21, 31 and 41, thanks to Uniq(p_{src}), the assertions about $*p_{src}$ are preserved.

Note that in the example of Fig. 3, it suffices to use Priv(p_{src}) instead of Uniq(p_{src}). However, in general when more than one location is promoted, we need to know that those promoted locations are not aliased with each other, which follows from Uniq(p_{src}) for each promoted location p. Also for the sake of performance, we use Uniq instead of introducing $\perp$ between each pair of promoted locations.

3.4 Proof Generation

LLVM's mem2reg pass consists of three algorithms: the general register-promotion algorithm and two specialized ones optimized for efficiency: one for the case that the promotable location is stored at most once and the other for the case that the location is used only within a single block. In this section we explain the general algorithm and its proof-generation code. Note that we also validate the two specialized algorithms in the same way since they are just degenerate cases.

Algorithm 2 shows the general algorithm implemented in LLVM's mem2reg pass and the proof-generation code, given in the box, that we inserted. Note that we do not modify the existing compiler code at all and only add the proof-generation code. In detail, the overall algorithm including proof generation works as follows.

Promotable Allocation **[Line A1]** We find a promotable allocation p at line l_a. **[Line A2]** Then we insert empty ϕ-nodes wherever needed[7], and add them to the maydiff set globally (*i.e.,* at every line). **[Line A3]** We also remove the allocation, insert lnop at that line, and add Uniq(p_{src}) and MD(p) globally. **[Line A4]** In addition, we add the rule intro_ghost($\hat{p}$,undef) because the initial value undef in $*p$ may be used by some load from $*p$ (though it does not happen in Fig. 3). In that case, the code at line A13 would introduce $\{ *p_{src} = \hat{p}_{src}, \hat{p}_{tgt} = \text{undef} \}$ at line l_a, which will be inferred with the help of intro_ghost($\hat{p}$,undef).

For example, in Fig. 3, the empty ϕ-node p1 := $\phi(-,-)$ is inserted in $\mathbf{B}_{exit}$ and p1 is added to the maydiff set globally; then the allocation at line 10 is removed, lnop is inserted, Uniq(p_{src}) is added and p is added to the maydiff set globally; and finally intro_ghost($\hat{p}$,undef) is added at line 10.

Block Traversal **[Lines A5-A7]** We traverse the blocks in DFS order starting from the entry block using the worklist WL. An element of WL consists of triple (B,v,l), where B is

[7]The optimization uses the "dominance frontier" algorithm [18] in order to list up the blocks that require a ϕ-node. We omit the details for brevity.

Algorithm 2 RegisterPromotion(F:Function)

```
A1: for l_a: p := alloca() in F if p's uses are loads/stores only do
A2:   InsertEmptyPhinodesFor(F, p)
                  // Add the φ-nodes to the maydiff set globally
A3:   Remove(l_a), Nop(l_a,tgt),Assn({Uniq(p_src),MD(p)},global)
A4:   Inf(intro_ghost(p̂,undef),l_a)
A5:   WL := [(Entry(F), undef, l_a)], MarkVisited(Entry(F))
A6:   while NonEmpty(WL) do
A7:     (B, v, l) :: WL := WL
A8:     for (l_i : i) in Instr(B) do
A9:       if i is a store instruction (*p := w) then
A10:        Remove(l_i), Nop(l_i, tgt),Inf(intro_ghost(p̂,w),l_i)
A11:        v := w, l := l_i
A12:      else if i is a load instruction (x := *p) then
A13:        Assn({*p_src = p̂_src, p̂_tgt = v_tgt}, l, l_i)
A14:        Inf(intro_ghost(x̂,p̂),l_i)
A15:        for (l_j : j) in Use(x) do
A16:          Replace(F,l_j,x,v), Assn({x_src = x̂_src, x̂_tgt = v_tgt},l_i,l_j)
A17:        end for
A18:        Remove(l_i), Nop(l_i, tgt),Assn({MD(x)}, global)
A19:      end if
A20:    end for
A21:    for B' in Successor(B) do
A22:      if B' has a φ-node (z := φ(· ·)) inserted at line A2 then
A23:        z[B] := v, Assn({*p_src = p̂_src, p̂_tgt = v_tgt}, l, End(B))
A24:        if not IsVisited(B') then WL := (B',z,Begin(B')) :: WL
A25:      else
A26:        if not IsVisited(B') then WL := (B', v, l) :: WL
A27:      end if
A28:      MarkVisited(B')
A29:    end for
A30:  end while
A31: end for
A32: Auto(transitivity)
```

the block to visit, v is the value in $*p$ at the beginning of B, and l is the line number where the value v is stored in $*p$. **[Line A5]** Initially, we put (Entry(F), undef, line l_a) in WL and mark the entry block Entry(F) as visited. **[Lines A6-A7]** Then we process the blocks in WL one by one. For example, in Fig. 3, $\mathbf{B}_{\texttt{entry}}$, $\mathbf{B}_{\texttt{left}}$, $\mathbf{B}_{\texttt{exit}}$, and $\mathbf{B}_{\texttt{right}}$ are visited in order.

Instruction Traversal **[Line A8]** Given a work (B, v, l), we traverse each instruction $(l_i : i)$ in the block B as follows.

Store Instructions **[Lines A9-A11]** If i is a store instruction $*p := w$ (line A9), then we remove the instruction (line A10) and update v with the stored value w (line A11). The proof-generation code inserts `lnop`, adds `intro_ghost`$(\hat{p}, w)$ (line A10) and updates l with the store location l_i (line A11).

For example, in Fig. 3, when i is 11: `*p := 42`, the store i is replaced by `lnop`; `intro_ghost`$(\hat{\texttt{p}}, 42)$ is added at line 11; and v and l are updated to be 42 and line 11.

Load Instructions **[Lines A12-A18]** If i is a load instruction $x := *p$ (line A12), then we replace all the uses of x with the current value v (lines A15-A17), and remove the load instruction (line A18). The proof-generation code adds the relational assertion $*p_{src} = v_{tgt}$ from the store site l to the load site l_i (line A13) and the rule `intro_ghost`$(\hat{x}, \hat{p})$ at l_i (line A14). Then it adds $x_{src} = v_{tgt}$ from the load site l_i to every use site l_j (line A16). It also inserts `lnop` at l_i in the target and adds x to the maydiff set globally (line A18).

For example, in Fig. 3, when i is 20: `a := *p`, the load i is replaced by `lnop`; the use of `a` is replaced by the current value 42 at line 21; $*\texttt{p}_{src} = 42_{tgt}$ is added from 11 to 20; `intro_ghost`$(\hat{\texttt{a}}, \hat{\texttt{p}})$ is added at line 20; $\texttt{a}_{src} = 42_{tgt}$ is added from 20 to 21; and `a` is added to the maydiff set globally.

Successors **[Lines A21-A28]** Now we handle the successor (*i.e.,* outgoing) blocks of the current block B. **[Line A21]** We traverse each successor block B' as follows.

- If B' has a ϕ-node $(z := \phi(\cdots))$ that is inserted by the code at line A2 (line A22), then we update the ϕ-node z's component for the incoming block B with the value v of $*p$ at the end of B (line A23). In addition, if B' has not been visited yet, we add $(B', z, \text{Begin}(B'))$ to the worklist WL (line A24). Since the value v is used at the ϕ-node z, we add $*p_{src} = v_{tgt}$ from store location l to the end of B (line A23). For example, in Fig. 3, when (B, B') is $(\mathbf{B}_{\texttt{left}}, \mathbf{B}_{\texttt{exit}})$, the ϕ-node `p1 := `$\phi(-,-)$ is updated to `p1 := `$\phi(42,-)$ and $(\mathbf{B}_{\texttt{exit}}, \texttt{p1}, \text{Begin}(\mathbf{B}_{\texttt{exit}}))$ is added to the worklist WL. Also $*\texttt{p}_{src} = 42_{tgt}$ is added from line 11 to the end of $\mathbf{B}_{\texttt{left}}$.
- If B' has no such ϕ-node (line A25), then we simply add (B', v, l) to the worklist WL if B' has not been visited yet (line A26). For example, when (B, B') is $(\mathbf{B}_{\texttt{entry}}, \mathbf{B}_{\texttt{right}})$, the triple $(\mathbf{B}_{\texttt{right}}, 42, \text{line 11})$ is added to the worklist.

[Line A28] Finally the successor B' is marked as visited.

Inference Rules As shown in §3.3, the complete proof for `mem2reg` contains two inference rules, `intro_ghost` and `transitivity`. The `intro_ghost` rules are explicitly added by the proof-generation code shown in Algorithm 2, while the `transitivity` rules are automatically added by the automation function `transitivity` (line A32).

4 Reasoning about Cyclic Control Flows

In this section, using an example of `fold-`ϕ optimization, we discuss a challenge in ERHL validation arising from cyclic control flows and show how to address it.

Fold-ϕ Optimization Consider the translation below performed by the `fold-`ϕ optimization of `instcombine`, and its ERHL proof. This translation basically replaces `z :=` ϕ`(x, y)` with `z :=` ϕ`(a, z)+1` using the temporary variable `t :=` ϕ`(a, z)`. This removes the dependence of `z` on `x` and `y`, thereby allowing `x` and `y` to be eliminated away by a subsequent optimization unless they are used elsewhere. This translation is correct because we have $\texttt{z}_{src} = \phi(\texttt{x}_{src}, \texttt{y}_{src}) = \phi(\texttt{a}_{src}+1, \texttt{z}_{src}+1) = \phi(\texttt{a}_{tgt}+1, \texttt{z}_{tgt}+1) = \phi(\texttt{a}_{tgt}, \texttt{z}_{tgt}) + 1 = \texttt{t}_{tgt} + 1 = \texttt{z}_{tgt}$.

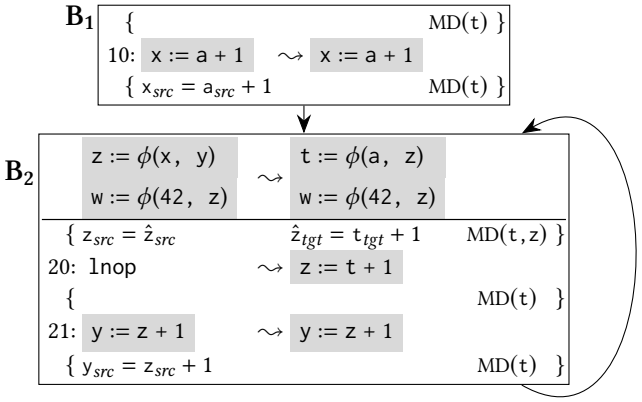

Note that a set of ϕ-nodes can appear at the beginning of a block and are executed *simultaneously*. For example, in the source program above, when control flows from $\mathbf{B}_2$ to itself, the ϕ-nodes `z` and `w` are set to the *old values* of `y` and `z` just before executing the ϕ-nodes, respectively. In particular, `w` is set to the old value of `z`, not the new value stored in `z` at the first ϕ-node, and thus `w` contains the same value in the source and target programs.

Challenge The challenge here is that we should be able to express and reason about both old and new values of `z`. This is because `z` is used and defined at the same time in the ϕ-nodes, which is only possible due to cyclic control flows in the SSA form. Specifically, the proof checker should derive something like $\texttt{z}_{src} = \texttt{y}_{src}$ and $\texttt{w}_{src} = \mathit{old}(\texttt{z}_{src})$ as part of the strong post-condition after the ϕ-nodes when control flows from $\mathbf{B}_2$.

We address this challenge by expressing the old value of register $\mathit{old}(\texttt{z}_{src})$ using a ghost variable. Specifically, we reserve a set of ghost registers, denoted $\bar{r}$ and called *old registers*, for all registers r to represent the old value of r. Note, however, that old registers are just normal ghost registers

and technically have nothing to do with physical old values of the corresponding registers.

Proof Validation We show how the ERHL proof checker systematically uses the old registers by validating the above proof in the most interesting case: the ϕ-nodes of $\mathbf{B}_2$ when control comes from itself.

First, it computes a post-assertion from the pre-assertion $\{\, \mathtt{y}_{src} = \mathtt{z}_{src} + 1, \mathrm{MD}(\mathtt{t}) \,\}$ as follows.

1. It removes all assertions about old registers from the pre-assertion and copies all assertions about current registers into those about old ones.

$$\{\, \mathtt{y}_{src} = \mathtt{z}_{src} + 1, \bar{\mathtt{y}}_{src} = \bar{\mathtt{z}}_{src} + 1, \mathrm{MD}(\mathtt{t}, \bar{\mathtt{t}}) \,\}.$$

2. It computes a post-condition from this new assertion as if the assignments `z := ȳ, w := z̄` are executed in the source and `t := z̄, w := z̄` in the target. Specifically, it (*i*) removes source assertions about `z,w` and target ones about `t,w` because those registers are updated; (*ii*) adds `t,z` to the maydiff set because they are updated differently in the source and target (note that `w` is updated equivalently since $\bar{\mathtt{z}}$ is not in the maydiff set); and (*iii*) adds the equalities corresponding to the executed assignments. Thus we have

$$\{\, \bar{\mathtt{y}}_{src} = \bar{\mathtt{z}}_{src} + 1, \mathtt{z}_{src} = \bar{\mathtt{y}}_{src}, \mathtt{w}_{src} = \bar{\mathtt{z}}_{src}, \\ \mathtt{t}_{tgt} = \bar{\mathtt{z}}_{tgt}, \mathtt{w}_{tgt} = \bar{\mathtt{z}}_{tgt}, \mathrm{MD}(\mathtt{t}, \bar{\mathtt{t}}, \mathtt{z}) \,\}.$$

Then the proof gives the rule $\mathtt{intro\_ghost}(\hat{\mathtt{z}}, \bar{\mathtt{z}} + 1)$, which adds $\{\, \bar{\mathtt{z}}_{src} + 1 = \hat{\mathtt{z}}_{src}, \hat{\mathtt{z}}_{tgt} = \bar{\mathtt{z}}_{tgt} + 1 \,\}$ because $\bar{\mathtt{z}}$ is not in the maydiff set. Then the automation function derives $\{\, \mathtt{z}_{src} = \hat{\mathtt{z}}_{src}, \hat{\mathtt{z}}_{tgt} = \mathtt{t}_{tgt} + 1 \,\}$ by transitivity: $\mathtt{z}_{src} = \bar{\mathtt{y}}_{src} = \bar{\mathtt{z}}_{src} + 1 = \hat{\mathtt{z}}_{src}$ and $\hat{\mathtt{z}}_{tgt} = \bar{\mathtt{z}}_{tgt} + 1 = \mathtt{t}_{tgt} + 1$. Then it eliminates $\bar{\mathtt{t}}$ from the maydiff set after eliminating all assertions about $\bar{\mathtt{t}}$, which is sound because $\bar{\mathtt{t}}$ is just a ghost variable that has nothing to do with a physical value of the register `t`. Finally, the assertion after the ϕ-nodes $\{\, \mathtt{z}_{src} = \hat{\mathtt{z}}_{src}, \hat{\mathtt{z}}_{tgt} = \mathtt{t}_{tgt} + 1, \mathrm{MD}(\mathtt{t}, \mathtt{z}) \,\}$ trivially follows by a simple inclusion check.

5 ERHL Proof Checker and Logic

In this section, we explain the proof checker in terms of the ERHL logic, and describe the soundness of the proof checker using the semantic interpretation of the logic. All our results are formally verified in Coq (see [1, §H] for details).

Proof Rules The checker is based on the proof rules presented in Fig. 4. The checker is given the source and target programs Prg_{src}, Prg_{tgt} and a translation proof Ψ, and tries to deduce $Prg_{src} \sim Prg_{tgt}$ using the (Sim) rule. Here, $Entry(F)$ denotes the entry block of the function F; $Prg[F].\zeta[B,i]$ the i-th instruction of the block B in F; and $Prg[F].\phi[B,B']$ the assignment instructions of the ϕ-nodes of B' when control comes from B (*e.g.,* in the source program in §4, $Prg[F].\phi[\mathbf{B_1},\mathbf{B_2}] = \{\, \mathtt{z := x, w := 42} \,\}$). Also, $\Psi[F].\alpha[B,i]$ denotes the assertion in the proof Ψ just before the i-th instruction of B in F (it denotes the last assertion when $i = -1$).

$$\boxed{Prg_{src} \sim Prg_{tgt}}$$

$$\frac{\begin{array}{c}\mathrm{CheckCFG}(Prg_{src}, Prg_{tgt}) \quad \forall F \in Prg_{src}.\ \mathrm{CheckInit}(\Psi[F].\alpha[\mathrm{Entry}(F), 0]) \\ \forall F,B,i.\ \{\Psi[F].\alpha[B,i]\}\ Prg_{src}[F].\zeta[B,i] \sim Prg_{tgt}[F].\zeta[B,i]\ \{\Psi[F].\alpha[B,i+1]\} \\ \forall F,B,B'.\ \{\Psi[F].\alpha[B,-1]\}\ Prg_{src}[F].\phi[B,B'] \sim Prg_{tgt}[F].\phi[B,B']\ \{\Psi[F].\alpha[B',0]\}\end{array}}{Prg_{src} \sim Prg_{tgt}}\ (\textsc{Sim})$$

$$\boxed{\{P\}\ I_{src} \sim I_{tgt}\ \{Q\}}$$

$$\frac{\mathrm{CheckEquivBeh}(P, I_{src}, I_{tgt}) \quad Q = \mathrm{CalcPostAssn}(P, I_{src}, I_{tgt})}{\{P\}\ I_{src} \sim I_{tgt}\ \{Q\}}\ (\textsc{PostAssn})$$

$$\frac{\{P\}\ I_{src} \sim I_{tgt}\ \{Q\} \quad Q \Rrightarrow Q'}{\{P\}\ I_{src} \sim I_{tgt}\ \{Q'\}}\ (\textsc{Consequence})$$

$$\boxed{Q \Rrightarrow Q'}$$

$$\frac{Q \Rrightarrow Q' \quad Q' \Rrightarrow Q''}{Q \Rrightarrow Q''}\ (\textsc{Trans}) \qquad \frac{rule \in \mathrm{CustomRules} \quad Q' = \mathrm{ApplyInf}(rule, Q)}{Q \Rrightarrow Q'}\ (\textsc{ApplyInf}) \qquad \frac{\mathrm{CheckIncl}(Q, Q')}{Q \Rrightarrow Q'}\ (\textsc{Incl})$$

Figure 4. Proof Rules of ERHL

The checker first checks if Prg_{src} and Prg_{tgt} have the same CFG (CheckCFG), the assertion in the entry is satisfied by the initial states for each function (CheckInit), and the Hoare triple $\{P\}\ I_{src} \sim I_{tgt}\ \{Q\}$ is valid for all matching intra-block commands I_{src} and I_{tgt} and their pre- and post-assertions P and Q given by Ψ. For example, in Fig. 2, it checks at line 20 if $\{\mathtt{x}_{src} = a_{src} + 1, MD(\emptyset)\}\ \mathtt{y := x + 2} \sim \mathtt{y := a + 3}\ \{MD(\emptyset)\}$ is valid. It also checks for each inter-block edge from B to B' that $\{P\}\ Prg_{src}.\phi[B,B'] \sim Prg_{tgt}.\phi[B,B']\ \{Q\}$ is valid, where P is the last assertion in B and Q is the first assertion in B'.

To validate a Hoare triple $\{P\}\ I_{src} \sim I_{tgt}\ \{Q\}$, the checker first computes a post-assertion Q_0 with $\{P\}\ I_{src} \sim I_{tgt}\ \{Q_0\}$ using the rule PostAssn (see [1, §H] for the definition of CheckEquivBeh and CalcPostAssn). Then it suffices to validate $Q_0 \Rrightarrow Q$ by the rule Consequence.

For this, using the rules ApplyInf and Trans, the checker iteratively applies a sequence of inference rules $rule_1, \ldots, rule_n$ (either given by Ψ or generated by an automation function) and deduces $Q_0 \Rrightarrow Q_n$, where $Q_i = \mathrm{ApplyInf}(rule_i, Q_{i-1})$.

Finally, the checker validates $Q_n \Rrightarrow Q$ using the rule Incl, where CheckIncl performs a simple inclusion check.

Semantic Interpretation For the soundness of the proof checker, we give the semantic interpretation of the top-level judgment as semantics preservation, or behavior refinement:

$$[\![Prg_{src} \sim Prg_{tgt}]\!] \stackrel{\text{def}}{=} Beh(Prg_{src}) \supseteq Beh(Prg_{tgt})\ .$$

The soundness of (Sim) is proved using a local simulation in the style of [22], which is a simplification of parametric bisimulation [21]. First, we show that CheckInit(P) implies:

$$\forall \sigma_{src}, \sigma_{tgt}, \alpha.\ FInit(\sigma_{src}) \wedge FInit(\sigma_{tgt}) \implies [\![P]\!]_\alpha(\sigma_{src}, \sigma_{tgt})\ .$$

Here, $FInit(\sigma)$ means σ is a possible initial state of a function call, $[\![P]\!]$ is the semantic interpretation of the assertion P (see [1, §G] for details), and α is a CompCert-style memory injection [28], which basically maps a memory block in the source to an equivalent one in the target.

Second, we give the semantic interpretation of the Hoare triple for non-call instructions I_{src}, I_{tgt} as a simulation step:

$$\begin{aligned}[\![\{P\}\ I_{src} \sim I_{tgt}\ \{Q\}]\!] \stackrel{\text{def}}{=} \forall \sigma_{src}.\ Instr(\sigma_{src}) = I_{src} \implies \\ \forall \sigma_{tgt}.\ Instr(\sigma_{tgt}) = I_{tgt} \implies \\ \forall \alpha, \sigma'_{tgt}, \varepsilon.\ [\![P]\!]_\alpha(\sigma_{src}, \sigma_{tgt}) \ \wedge\ \sigma_{tgt} \xrightarrow{\varepsilon} \sigma'_{tgt} \implies \\ \exists \sigma'_{src}, \alpha'.\ [\![Q]\!]_{\alpha'}(\sigma'_{src}, \sigma'_{tgt}) \ \wedge\ \sigma_{src} \xrightarrow{\varepsilon} \sigma'_{src} \ \wedge\ \alpha \sqsubseteq \alpha'\end{aligned}$$

where, $Instr(\sigma)$ is the next instruction to execute in the program state σ, and $\sigma \xrightarrow{\varepsilon} \sigma'$ means the state σ steps to σ' emitting an observable event ε. Also, $\sqsubseteq$ is the extension relation of memory injection.

For call instructions I_{src}, I_{tgt}, $[\![\{P\}\ I_{src} \sim I_{tgt}\ \{Q\}]\!]$ basically states that Q is satisfied by all possible equivalent returns states when an arbitrary function is called from states satisfying P (see [1, §H] for details). We followed the basic approach of parametric bisimulation [21].

The semantic interpretation of $\Rrightarrow$ is as follows:

$$\begin{aligned}[\![Q \Rrightarrow Q']\!] \stackrel{\text{def}}{=} \forall \sigma_{src}, \sigma_{tgt}, \alpha.\ [\![Q]\!]_\alpha(\sigma_{src}, \sigma_{tgt}) \implies \\ \exists \alpha'.\ [\![Q']\!]_{\alpha'}(\sigma_{src}, \sigma_{tgt}) \ \wedge\ \alpha \sqsubseteq \alpha'\ .\end{aligned}$$

For the soundness of (APPLYINF), every custom *rule* should satisfy that $[\![Q \Rrightarrow \text{ApplyInf}(rule, Q)]\!]$ holds for all Q.

6 Implementation

We developed the CRELLVM framework for LLVM 3.7.1.

Coverage We wrote proof-generation code for register promotion in the `mem2reg` pass and GVN-PRE in the `gvn` pass implemented in the following files respectively:

- `lib/Transforms/Utils/PromoteMemoryToRegister.cpp`
- `lib/Transforms/Scalar/GVN.cpp`

For `mem2reg`, we covered the entire file, and for `gvn`, we covered all functions except for the following functions: `SimplifyInstruction`, `processLoad`, `splitCriticalEdges` and `MergeBlockIntoPredecessor`. These functions are not part of the main GVN-PRE algorithm because they are not technically related to value numbering (*i.e.,* neither using nor constructing value numbering). Other reasons why we omitted them are because `SimplifyInstruction` is a common function that just consists of many peephole optimizations and the others use features that are not currently supported by CRELLVM: `processLoad` uses the alias analysis module and `splitCriticalEdges` and `MergeBlockIntoPredecessor` change control-flow graphs. Note that the reason why those functions are used by the `gvn` pass is because they transform programs in such a way that opportunities for GVN-PRE optimizations are increased.

To demonstrate the generality of ERHL logic and the proof checker, we also covered a part of the loop-invariant code motion (`licm`) pass that can be currently supported by CRELLVM and 139 micro-optimizations of the instruction combining (`instcombine`) pass (see [1, §D] for details).

	`mem2reg`	`gvn`	`licm`	`instcombine`
Compiler (Covered)	568	1,092	706	702
Proof Generation	213	440	286	1,357

Figure 5. SLOC of Proof-Generation Code

Proof-Generation Code We explicitly mark as "not supported" for translations using operations not supported by VELLVM, or relying on deep analyses such as division-by-zero and alias analyses.

Fig. 5 shows the SLOC in C++ of the compiler and proof-generation code for each pass. The SLOC ratio of the proof-generation code to that of the corresponding compiler code is 37.5% for `mem2reg`, 40.3% for `gvn`, 40.5% for `licm`, and 193.3% for `instcombine`. The CRELLVM infrastructure for proof-generation consists of 1,708 lines for common library and 15,980 lines for JSON serialization library, of which 72.2% is automatically generated from 2,079 SLOC in a simple DSL.

Inference Rules In the proof checker we installed 221 custom inference rules, of which 202 are arithmetic rules like `assoc_add`. All 9 non-arithmetic rules used for `mem2reg`, `gvn`, and `licm`, including `transitivity` and `intro_ghost`, are formally verified in Coq (see [1, §I] for details).

Verification of Proof Checker In order to reduce TCB, we formally verified the soundness of the proof checker in Coq (see §5). It is worth noting that we achieved the same kind of guarantee as CompCert for the translations that are validated by the proof checker using only verified inference rules.

We used the formal semantics of LLVM IR from the VELLVM project [55], but significantly upgraded the semantics in various ways. In particular, VELLVM used the CompCert memory model [28] version 1.9 and we upgraded it to version 2.4 in order to use the notion of *permission* in the LLVM semantics; and added the `switch` instruction to the formalization of LLVM IR. Note that VELLVM has a simpler memory model than the LLVM's informal official one (*e.g.,* pointer-equality tests and pointer-integer casts are more undefined).

In total, our Coq development consists of 25,970 SLOC. The proof checker is 2,987 SLOC, and its verification is 18,934 SLOC. The 221 inference rules are 2,193 SLOC, and the verification of 9 rules took 1,856 SLOC. Note that the underlying semantics of VELLVM consists of 39,307 SLOC.

Experience Writing proof-generation code was an iterative process: we had to repeat bug-fix processes many times. When proof checking fails, it tells us a logical reason for the failure so that we could easily identify the bug in proof-generation code (or else in the compiler). We believe the iteration could be shortened if we collaborated with LLVM developers.

	Results			Time (sec.)			
	#V	#F	#NS	Orig	PCal	I/O	PCheck
mem2reg	76.79K	10	10.58K	8.59	322.18	13.16K	21.26K
gvn	365.99K	453	7.92K	41.81	249.85	41.96K	37.89K
licm	168.20K	0	24.93K	22.42	895.93	56.44K	11.36K
instcombine	1593.84K	0	528.75K	184.49	442.85	152.51K	105.40K

Figure 6. Experimental Results

Custom functions for automatically finding inference rules are greatly helpful for developing proof-generation algorithms. Using such automation, we could develop much simpler proof-generation algorithms for `mem2reg` and `gvn`, compared to our initial development, by making the code size less than half and speeding up more than twice.

CRELLVM is less cost-effective for peephole optimizations in `instcombine`. We had to write 1.9 lines of proof-generation code for each line of the corresponding compiler code, and we did not verify arithmetic inference rules. Even though CRELLVM achieves higher level of reliability, we think more automated approaches using an SMT solver such as Alive [30] would be more cost-effective for peephole optimizations.

7 Experiment

Benchmarks Using CRELLVM, we validated the compilation of the SPEC CINT2006 C Benchmarks [15], LLVM nightly test suite, and five open-source projects written in C (the biggest benchmarks used in [37][8]), totaling 5.3 million LOC in C. We omitted 3 files from the benchmarks because they contain instructions currently not supported by VELLVM, including the `indirectbr` instruction.

Fig. 6 summarizes the validation results and the time spent on running the proof-generation codes and the proof checker for each optimization pass. In the experiment, we compiled each benchmark program with the `-O2` flag, and validated the intermediate translations with the generated proofs. For more detailed results, see [1, §A].

We show the total number of translation steps (**#V**), the number of not-supported translations (**#NS**), and the number of translations failed at validation (**#F**). The rest of the translations (*i.e.,* **#V** − **#F** − **#NS**) succeeded in validation. Also, all the successful translations were shown to be equivalent to the original translations using the `llvm-diff` tool. During the experiment, we also found and reported a bug in `llvm-diff`, which has been confirmed and fixed [8].

Out of 2,205K validations in total, 1632K (74.0%) are successfully validated. All 463 (0.01%) failures (**#F**) are due to compiler bugs: 10 are due to the `mem2reg` bug [5] we discussed in §1.2, 295 are due to the two `gvn` bugs [6, 7] we found, and 158 are due to a known `gvn` bug [11] that is currently fixed in the LLVM trunk. Note that there is no failure due to the other `mem2reg` bug [9] we found.

The other 572.2K (26.0%) translations (**#NS**) are currently not supported in our validator. Among them, 555.9K (97.1%) use instructions not supported by VELLVM: vector operations 515.1K (90.0%), aggregate type operations 30.4K (5.31%), debug attributes 8.7K (1.52%), and atomic operations 1.7K (0.29%). 13.0K (2.27%) use the alias and division-by-zero analysis modules of LLVM; 2.3K (0.41%) alter type declarations; and 0.7K (0.12%) require deeper analysis on functions such as read-only function analysis.

We measured the time spent on performing each optimization in the original compiler (**Orig**); on performing each optimization and calculating validation proofs in the modified compiler (**PCal**); on writing and reading the source and target programs with the proofs via files (**I/O**); and on validating the proofs by the proof checker (**PCheck**). The table shows total times aggregated over the entire run.

In the experiment, we embarrassingly parallelized compilation and validation jobs and fully utilized the 96 hardware threads from four identical workstations with Intel Xeon E5-2630 CPU (2.6GHz, 12 cores, 2 hardware threads per core), 128GB RAM, and 1TB SSD (Samsung 850 PRO). The whole experiment took about three hours in wall clock.

Validating Randomly Generated Programs We randomly generated 1,000 C programs using CSmith [53], compiled them with `-O2` flag, and validated the intermediate translations with the generated proofs. All 55,008 validations for `gvn` are successfully validated, except for one due to the `gvn` bug [6] we found. Out of 42,584 validations for `mem2reg`, 11,816 (27.7%) are currently not supported due to LLVM lifetime intrinsics, which is not supported by VELLVM. The other 30,768 (72.3%) are all successfully validated.

Performance Proof checking takes much more time than regular compilation, but we believe it is still reasonable for compiler writers to use CRELLVM for stabilizing compilers. Also, as we have shown in the experiment, compiler writers can further reduce runtime by checking proofs in parallel. Furthermore, there is still a large room for performance improvement as we have not done any serious performance analysis and tuning for the proof checker. In particular, we believe we can significantly reduce I/O time, which is one of the current bottlenecks, by writing proofs in binary format rather than in plain-text JSON format and also by writing only the changes made between IR files rather than writing full IR files. In our benchmark, the CLANG frontend generated 4,885 IR files with average size of 187.63 KB, from which 2,205K validations with average proof size of 17.5 KB were generated.

Bug Reports By November 2016 when we completed our initial implementation of CRELLVM for LLVM 3.7.1, we reported three miscompilation bugs, one in `mem2reg` [5] and two in `gvn` [6, 7], which were immediately confirmed and subsequently fixed. Around July 2017 when we verified selected inference rules, we reported another miscompilation bug in `mem2reg` [9], which was immediately confirmed but has not been fixed yet (as of 14 April 2018) because it is

[8] We omitted Linux, since it is currently not compiled with LLVM (see [29]).

unlikely to occur in practice (it did not occur in our benchmark either) and there is no consensus on how to fix it. Around March 2018, we additionally covered the function `performScalarPREInsertion` in `gvn`, which was omitted initially because it is loosely related to value numbering: deciding whether to perform the transformation, not the transformation itself, depends on value numbering. The reason for this coverage is because we were informed of a new bug [11] found in the function. As we have seen above, CRELLVM successfully detected the bug by failing at 158 validations.

8 Discussion

8.1 Reliability

In order to see how effectively CRELLVM improved reliability of LLVM, we investigated all bug reports about miscompilation in `mem2reg` and `gvn` since the release of LLVM 3.7.1. To the best of our knowledge, other than the five bugs [5–7, 9, 11] detected by CRELLVM, there is no confirmed miscompilation bug that is (*i*) due to the code we covered in `mem2reg` and `gvn` and (*ii*) not related to any LLVM feature that is currently not supported by CRELLVM (as of 14 April 2018).

Specifically, we conducted our investigation as follows. We checked all relevant bug reports in the LLVM bug tracker [4] and OSS-Fuzz bug tracker [3]. Moreover, we asked the `llvm-dev` mailing list about relevant bugs [2]. We also posted a draft of this paper on our website in February 2018 and received comments. One of the most important comments was about the `gvn` bug [11] in the code we newly covered (*i.e.,* the function `performScalarPREInsertion`). The bug was discovered and fixed in October 2017 by Azul Systems via fuzz testing of the company's LLVM-based Java JIT compiler, using JavaFuzzer [10] (private communication with Philip Reames, March 2018).

8.2 Maintainability

To evaluate maintenance cost, we ported our full development of CRELLVM to LLVM 5.0.1 just omitting `instcombine` because it is not our main target. After the initial porting, which took two days, we found one validation fail in `gvn` due to insufficient proof generation. We fixed it by adding an automation function, which took 5 days by one person including analysis of the problem. After applying the `gvn` bug fix [11] in the main trunk to LLVM 5.0.1, our benchmark experiment produces no validation failures except for not-supported ones (see [1, §A] for details).

8.3 Limitations and Future Work

We discuss current limitations of CRELLVM, which also indicate a direction of future research.

Semantics VELLVM does not fully formalize the LLVM IR semantics. First, it does not support several features of LLVM IR, including atomic operations for concurrency, vector operations and attributes like `noalias`, `readonly` and `nsw`.

Second, VELLVM does not properly formalize casts between integers and pointers, which itself is a challenging research topic. Applying the idea of Kang *et al.* [22] would be interesting future research.

Finally, VELLVM does not properly formalize the *undef* and *poison*
Lee *et al.* [25] proposed a possible solution to this problem using a new instruction, called *freeze*. Applying it to VELLVM would be interesting work.

Analyses Our proof checker does not support various analysis passes such as division-by-zero analysis, alias analysis, read-only function analysis, and memory dependence analysis. We believe it would be possible to support them by adding appropriate predicates and inference rules in the underlying logic of proof checker.

CFG-Changing Optimizations CRELLVM relies on the condition that the source and target programs can be aligned line-by-line by inserting logical no-op instructions. While we think this condition holds for majority of LLVM optimizations, there are several important optimizations that break the condition by changing the control-flow graph. Examples include loop unrolling, loop unswitching and loop splitting. We believe it would be possible to support them by generalizing the proof checker following the ideas from existing translation validation works [36, 49–52, 57].

9 Related Work

A large number of prior work on improving reliability of compiler are roughly classified into the following categories.

Credible Compilation Rinard *et al.* [44], who coined the term *credible compilation*, proposed the framework of credible compilation and presented a relational Hoare logic, in which one can reason about register allocation and instruction scheduling optimizations in the presence of pointer aliasing. Independently, Benton [16] proposed a relational Hoare logic for a functional language. However, their logics are designed for simple languages, and the framework has not been implemented and applied to compilers.

Namjoshi *et al.* [33, 34] presented a "proof of concept" implementation of credible compilation (or a witnessing compiler in their terminology) for LLVM optimizations such as constant propagation, dead-code elimination, and LICM. However, the work can be seen as rather preliminary for the following reasons. First, their proof checker supports a small subset of LLVM IR, most notably ignoring memory operations. Second, it assumes that main functions of the compiler are correct. For example, it assumes that the constant-folding function of LLVM is correct.

Verified translation validation is similar to verified credible compilation but differs in that it develops a verified

validator specialized for a particular optimization, rather than developing a proof checker for a general logic. Various verified translation validators have been developed for CompCert: instruction scheduling [50], lazy code motion [51], software pipelining [52]; register allocation [43]; SSA transformation [14]; and GVN and sparse conditional constant propagation (SCCP) [19].

(Foundational) proof carrying code (PCC) [12, 35] is similar to (verified) credible compilation, but it employs a (verified) unary logic for validating safety properties of the generated target program.

Translation Validation This approach develops a general validator that checks correctness of any given translation between IR programs without requiring any proof. Compared to credible compilation, translation validation is more scalable (*i.e.,* more easily applicable to different optimizations) because it requires much less manual effort due to no need for writing proof-generation code. On the other hand, though it can be used to guarantee correctness of certain compilations, it can hardly be used to find compiler bugs due to many false positives. The reason for false positives is that such a general validator is inherently incomplete since it is agnostic to the compiler's internal logic.

Due to such incompleteness, a variety of translation validators with different heuristics and trade-offs were proposed [20, 36, 38, 39, 45–47, 49, 54, 57, 58]. In particular, Tristan *et al.* [49] and Stepp *et al.* [46] developed translation validators for LLVM optimization passes, including dead-code elimination, GVN-PRE, constant propagation, and LICM. However, they failed at about 20% of the validations, most of which are likely to be false positives.

Compiler Verification Verified compilers provide the highest level of reliability by proving the semantics-preservation property for all possible source programs in a proof assistant. CompCert [26, 27] is the most sophisticated formally verified optimizing C compiler, whose correctness is proved in Coq [13], and CakeML [23] is an optimizing ML compiler formally verified in the HOL4 theorem prover [40]. However, verifying a full-fledged compiler is highly costly and verified compilers are usually much less performant than production compilers.

Zhao *et al.* [55, 56] implemented and verified the `vmem2reg` pass for LLVM in Coq, but its algorithm is significantly simplified compared to that in LLVM. Their simplified algorithm is based on a rewriting logic in which each rewriting step preserves semantics and each intermediate program is type-checked. On the other hand, LLVM's register-promotion algorithm temporarily breaks the semantics-preservation property and even the intermediate programs are not type-checked, because ill-formed empty ϕ-nodes are inserted in the middle and their arguments are filled later. According to the authors, this renders the formal verification hard for the register-promotion implementation in LLVM.

DSL for Optimizations Lopes *et al.* [30–32] presented Alive, a DSL for writing peephole optimizations using the SMT solver Z3 [41]. With Alive, one can either prove the correctness of an optimization or find a counterexample. They ported 300 micro-optimizations of `instcombine` to Alive, and in doing so they found 8 bugs in `instcombine`. However, the Alive DSL is not expressive enough to describe complex algorithms such as `mem2reg` and `gvn`, and limited to supporting only peephole optimizations that do not involve reasoning about cyclic control flows. In addition, Alive makes simplifying assumptions on the LLVM semantics, and their encoding of an optimization into SMT queries is a part of the TCB. Furthermore, since there is a gap between an actual implementation in C++ and a corresponding algorithm description in Alive DSL, implementation bugs cannot be detected. Tatlock and Lerner [48] also presented a DSL for writing CompCert optimizations based on a rewriting logic, but it is not general enough to support register promotion and GVN-PRE.

Compiler Testing Random testing tools such as CSmith [17, 42, 53] and EMI [24] have been very successful. They have found hundreds of bugs in GCC and LLVM. However, most of them are found in the `instcombine` pass and none of them are miscompilation bugs in `mem2reg` and `gvn`.

10 Conclusion

We have demonstrated that the credible-compilation approach scales to the production compiler LLVM by developing our Crellvm framework. We also empirically demonstrated that Crellvm can be an effective tool for achieving high reliability of major optimizations by discovering four long-standing bugs in the `mem2reg` and `gvn` passes.

Acknowledgments

We thank Daniel Berlin, Davide Italiano, Yeonwoo Kim, Philip Reames, John Regehr, and anonymous reviewers for very helpful feedback, and Sung-hwan Lee for his contribution to early development of Crellvm. This research was supported by Samsung Research Funding Center of Samsung Electronics under Project Number SRFC-IT1502-07. Jeehoon Kang, Yoonseung Kim, and Juneyoung Lee have been supported by Korea Foundation for Advanced Studies Scholarships.

References

[1] Supplementary material for this paper, available at http://sf.snu.ac.kr/crellvm/.
[2] http://lists.llvm.org/pipermail/llvm-dev/2018-April/122482.html.
[3] https://bugs.chromium.org/p/oss-fuzz.
[4] https://bugs.llvm.org/.
[5] https://bugs.llvm.org/show_bug.cgi?id=24179.
[6] https://bugs.llvm.org/show_bug.cgi?id=28562.
[7] https://bugs.llvm.org/show_bug.cgi?id=29057.
[8] https://bugs.llvm.org/show_bug.cgi?id=33623.
[9] https://bugs.llvm.org/show_bug.cgi?id=33673.

[10] https://github.com/AzulSystems/JavaFuzzer.
[11] https://reviews.llvm.org/D38619.
[12] Andrew W. Appel. 2001. Foundational Proof-Carrying Code *(LICS '01)*.
[13] The Coq Proof Assistant. https://coq.inria.fr/.
[14] Gilles Barthe, Delphine Demange, and David Pichardie. 2014. Formal Verification of an SSA-Based Middle-End for CompCert. *ACM Trans. Program. Lang. Syst.* 36, 1 (March 2014).
[15] The SPEC CINT2006 Benchmark. https://www.spec.org/cpu2006/CINT2006/.
[16] Nick Benton. 2004. Simple Relational Correctness Proofs for Static Analyses and Program Transformations *(POPL '04)*.
[17] Yang Chen, Alex Groce, Chaoqiang Zhang, Weng-Keen Wong, Xiaoli Fern, Eric Eide, and John Regehr. 2013. Taming Compiler Fuzzers *(PLDI '13)*.
[18] Ron Cytron, Jeanne Ferrante, Barry K. Rosen, Mark N. Wegman, and F. Kenneth Zadeck. 1991. Efficiently Computing Static Single Assignment Form and the Control Dependence Graph. *ACM Trans. Program. Lang. Syst.* 13, 4 (Oct. 1991).
[19] Delphine Demange, David Pichardie, and Léo Stefanesco. 2016. Verifying Fast and Sparse SSA-Based Optimizations in Coq *(CC '16)*.
[20] Chris Hawblitzel, Shuvendu K. Lahiri, Kshama Pawar, Hammad Hashmi, Sedar Gokbulut, Lakshan Fernando, Dave Detlefs, and Scott Wadsworth. 2013. Will You Still Compile Me Tomorrow? Static Cross-version Compiler Validation *(ESEC/FSE '13)*.
[21] Chung-Kil Hur, Derek Dreyer, Georg Neis, and Viktor Vafeiadis. 2012. The Marriage of Bisimulations and Kripke Logical Relations. In *POPL*.
[22] Jeehoon Kang, Chung-Kil Hur, William Mansky, Dmitri Garbuzov, Steve Zdancewic, and Viktor Vafeiadis. 2015. A Formal C Memory Model Supporting Integer-pointer Casts *(PLDI '15)*.
[23] Ramana Kumar, Magnus O. Myreen, Michael Norrish, and Scott Owens. 2014. CakeML: A Verified Implementation of ML *(POPL '14)*.
[24] Vu Le, Mehrdad Afshari, and Zhendong Su. 2014. Compiler Validation via Equivalence Modulo Inputs *(PLDI '14)*.
[25] Juneyoung Lee, Yoonseung Kim, Youngju Song, Chung-Kil Hur, Sanjoy Das, David Majnemer, John Regehr, and Nuno P. Lopes. 2017. Taming Undefined Behavior in LLVM *(PLDI '17)*.
[26] Xavier Leroy. 2006. Formal Certification of a Compiler Back-end or: Programming a Compiler with a Proof Assistant *(POPL '06)*.
[27] Xavier Leroy. 2009. Formal verification of a realistic compiler. *Commun. ACM* (2009).
[28] Xavier Leroy, Andrew W. Appel, Sandrine Blazy, and Gordon Stewart. 2012. *The CompCert Memory Model, Version 2*. Research report RR-7987. INRIA.
[29] LLVM Linux. http://llvm.linuxfoundation.org.
[30] Nuno P. Lopes, David Menendez, Santosh Nagarakatte, and John Regehr. 2015. Provably Correct Peephole Optimizations with Alive *(PLDI '15)*.
[31] David Menendez and Santosh Nagarakatte. 2017. Alive-Infer: Data-driven Precondition Inference for Peephole Optimizations in LLVM *(PLDI '17)*.
[32] David Menendez, Santosh Nagarakatte, and Aarti Gupta. 2016. Alive-FP: Automated Verification of Floating Point Based Peephole Optimizations in LLVM *(SAS '16)*.
[33] Kedar S. Namjoshi, Giacomo Tagliabue, and Lenore D. Zuck. 2013. A Witnessing Compiler: A Proof of Concept *(RV '13)*.
[34] Kedar S. Namjoshi and Lenore D. Zuck. 2013. Witnessing Program Transformations *(SAS '13)*.
[35] George C. Necula. 1997. Proof-carrying Code *(POPL '97)*.
[36] George C. Necula. 2000. Translation Validation for an Optimizing Compiler *(PLDI '00)*.
[37] Hakjoo Oh, Kihong Heo, Wonchan Lee, Woosuk Lee, Daejun Park, Jeehoon Kang, and Kwangkeun Yi. 2014. Global Sparse Analysis Framework. *ACM Trans. Program. Lang. Syst.* 36, 3 (Sept. 2014).
[38] Amir Pnueli, Michael Siegel, and Eli Singerman. 1998. Translation Validation *(TACAS '98)*.
[39] Amir Pnueli, Ofer Strichman, and Michael Siegel. 1998. The Code Validation Tool CVT: Automatic Verification of a Compilation Process *(STTT '98)*.
[40] HOL Interactive Theorem Prover. https://hol-theorem-prover.org/.
[41] The Z3 Theorem Prover. https://github.com/Z3Prover/z3.
[42] John Regehr, Yang Chen, Pascal Cuoq, Eric Eide, Chucky Ellison, and Xuejun Yang. 2012. Test-case reduction for C compiler bugs *(PLDI '12)*.
[43] Silvain Rideau and Xavier Leroy. 2010. Validating Register Allocation and Spilling *(CC '10)*.
[44] Martin C. Rinard and Darko Marinov. 1999. Credible Compilation with Pointers *(RRV '99)*.
[45] Hanan Samet. 1978. Proving the Correctness of Heuristically Optimized Code *(ACM '78)*.
[46] Michael Stepp, Ross Tate, and Sorin Lerner. 2011. Equality-based Translation Validator for LLVM *(CAV '11)*.
[47] Ross Tate, Michael Stepp, Zachary Tatlock, and Sorin Lerner. 2009. Equality Saturation: A New Approach to Optimization *(POPL '09)*.
[48] Zachary Tatlock and Sorin Lerner. 2010. Bringing Extensibility to Verified Compilers *(PLDI '10)*.
[49] Jean-Baptiste Tristan, Paul Govereau, and Greg Morrisett. 2011. Evaluating Value-graph Translation Validation for LLVM *(PLDI '11)*.
[50] Jean-Baptiste Tristan and Xavier Leroy. 2008. Formal Verification of Translation Validators: A Case Study on Instruction Scheduling Optimizations *(POPL '08)*.
[51] Jean-Baptiste Tristan and Xavier Leroy. 2009. Verified Validation of Lazy Code Motion *(PLDI '09)*.
[52] Jean-Baptiste Tristan and Xavier Leroy. 2010. A Simple, Verified Validator for Software Pipelining *(POPL '10)*.
[53] Xuejun Yang, Yang Chen, Eric Eide, and John Regehr. 2011. Finding and Understanding Bugs in C Compilers *(PLDI '11)*.
[54] Anna Zaks and Amir Pnueli. 2008. CoVaC: Compiler Validation by Program Analysis of the Cross-Product *(FM '08)*.
[55] Jianzhou Zhao, Santosh Nagarakatte, Milo M.K. Martin, and Steve Zdancewic. 2012. Formalizing the LLVM Intermediate Representation for Verified Program Transformations *(POPL '12)*.
[56] Jianzhou Zhao, Santosh Nagarakatte, Milo M.K. Martin, and Steve Zdancewic. 2013. Formal Verification of SSA-based Optimizations for LLVM *(PLDI '13)*.
[57] Lenore Zuck, Amir Pnueli, Benjamin Goldberg, Clark Barrett, Yi Fang, and Ying Hu. 2002. Translation and Run-Time Validation of Loop Transformations *(RV '02)*.
[58] Lenore D. Zuck, Amir Pnueli, and Benjamin Goldberg. 2003. VOC: A Methodology for the Translation Validation of Optimizing Compilers *(J. UCS '03)*.

	LOC	Result											
		Register Promotion			GVN-PRE			LICM			InstCombine		
		#V	#F	#NS	#V	#F	#NS	#V	#F	#NS	#V	#F	#NS
400.perlbench	168.16K	1.75K	0	1	11.90K	17	0	2.39K	0	105	59.34K	0	6.94K
401.bzip2	8.29K	90	0	0	1.63K	0	0	443	0	36	4.52K	0	1.84K
403.gcc	517.52K	5.43K	0	5	37.03K	21	0	8.35K	0	1.10K	140.67K	0	4.87K
429.mcf	2.69K	24	0	0	149	0	0	29	0	2	487	0	53
433.milc	15.04K	235	0	2	2.05K	0	0	1.78K	0	311	3.69K	0	471
445.gobmk	196.24K	2.64K	0	1	7.19K	0	0	2.50K	0	448	15.97K	0	1.50K
456.hmmer	35.99K	558	0	0	3.53K	3	2	2.54K	0	179	11.06K	0	3.41K
458.sjeng	13.85K	130	0	0	1.75K	0	0	355	0	69	3.78K	0	224
462.libquantum	4.36K	123	0	79	353	0	261	337	0	269	1.04K	0	782
464.h264ref	51.58K	532	1	0	12.81K	27	0	8.77K	0	1.49K	22.45K	0	5.31K
470.lbm	1.16K	19	0	0	77	0	0	61	0	2	174	0	51
482.sphinx3	25.09K	364	0	0	1.65K	0	2	1.05K	0	120	6.24K	0	1.04K
sendmail-8.15.2	138.68K	536	0	403	4.64K	4	107	1.81K	0	163	14.12K	0	396
emacs-25.1	463.54K	5.15K	0	4	28.67K	23	25	7.76K	0	622	112.94K	0	9.42K
python-3.4.1	486.38K	8.78K	0	89	28.02K	130	26	9.13K	0	381	89.11K	0	13.09K
gimp-2.8.18	1004.20K	19.45K	6	528	38.36K	135	315	24.37K	0	3.85K	150.89K	0	44.14K
ghostscript-9.14.0	797.65K	13.00K	0	9.18K	67.80K	27	6.95K	37.49K	0	7.79K	246.21K	0	82.75K
LLVM nightly test	1358.76K	17.98K	3	291	118.38K	66	240	59.04K	0	7.99K	711.14K	0	352.46K
Total	5289.18K	76.79K	10	10.58K	365.99K	453	7.92K	168.20K	0	24.93K	1593.84K	0	528.75K

Figure 7. Validation Results for LLVM 3.7.1

	Time (sec.)															
	mem2reg				GVN				LICM				InstCombine			
	Orig	PCal	I/O	PCheck	Orig	PCal	I/O	PCheck	Orig	PCal	I/O	PCheck	Orig	PCal	I/O	PCheck
400.perlbench	0.28	15.00	0.17K	0.70K	1.59	8.14	1.48K	4.16K	0.29	10.55	0.88K	0.32K	6.29	12.84	5.91K	9.13K
401.bzip2	0.01	1.46	0.01K	0.11K	0.09	0.48	0.14K	0.23K	0.06	1.46	0.08K	0.02K	0.40	1.19	0.24K	0.32K
403.gcc	0.71	41.22	5.12K	9.26K	5.66	33.42	6.35K	8.58K	1.56	30.32	3.14K	0.62K	19.47	59.12	32.51K	35.18K
429.mcf	<0.01	0.10	<0.01K	<0.01K	0.01	0.05	<0.01K	<0.01K	0.01	0.03	<0.01K	<0.01K	0.05	0.07	<0.01K	<0.01K
433.milc	0.03	0.94	0.01K	0.01K	0.11	0.51	0.08K	0.04K	0.08	0.37	0.07K	0.03K	0.47	0.59	0.11K	0.05K
445.gobmk	0.20	8.36	1.90K	0.95K	0.80	4.30	0.46K	0.42K	0.35	3.59	0.35K	0.09K	3.35	4.23	1.78K	0.98K
456.hmmer	0.07	3.06	0.03K	0.07K	0.33	1.72	0.11K	0.09K	0.18	1.79	0.16K	0.05K	1.54	2.05	0.24K	0.13K
458.sjeng	0.02	0.75	0.01K	0.02K	0.14	0.71	0.07K	0.06K	0.05	0.55	0.04K	0.01K	0.57	0.83	0.15K	0.14K
462.libquantum	0.01	0.25	<0.01K	<0.01K	0.02	0.14	<0.01K	<0.01K	0.01	0.10	0.01K	<0.01K	0.14	0.22	0.01K	<0.01K
464.h264ref	0.12	7.23	0.06K	0.39K	0.81	4.75	1.16K	0.95K	0.37	8.33	1.05K	0.47K	3.02	6.08	1.60K	0.99K
470.lbm	<0.01	0.19	<0.01K	<0.01K	0.01	0.04	<0.01K	<0.01K	0.02	0.05	<0.01K	<0.01K	0.03	0.04	<0.01K	<0.01K
482.sphinx3	0.04	1.54	0.01K	0.03K	0.17	0.82	0.05K	0.04K	0.09	0.64	0.05K	0.02K	0.90	1.13	0.16K	0.08K
sendmail-8.15.2	0.15	2.93	0.02K	<0.01K	0.53	2.51	0.49K	0.38K	0.22	3.70	0.45K	0.17K	2.31	3.50	1.44K	0.79K
emacs-25.1	0.64	36.25	1.06K	2.24K	3.18	22.35	4.74K	3.26K	1.00	25.40	4.07K	0.94K	16.10	36.35	19.08K	8.65K
python-3.4.1	0.82	35.31	1.12K	0.91K	3.88	25.66	4.23K	2.11K	0.77	19.01	3.50K	0.77K	16.29	27.26	12.48K	3.91K
gimp-2.8.18	1.50	48.09	1.07K	0.73K	4.98	28.47	2.14K	1.01K	2.02	20.30	2.20K	0.48K	28.76	40.68	5.48K	2.30K
ghostscript-9.14.0	2.03	25.46	0.32K	0.06K	7.68	40.24	4.52K	3.23K	4.44	52.43	5.29K	1.57K	27.66	52.16	12.35K	5.55K
LLVM nightly test	1.94	94.05	2.25K	5.77K	11.83	75.54	15.93K	13.33K	10.89	717.31	35.09K	5.79K	57.15	194.52	58.98K	37.19K
Total	8.59	322.18	13.16K	21.26K	41.81	249.85	41.96K	37.89K	22.42	895.93	56.44K	11.36K	184.49	442.85	152.51K	105.40K

Figure 8. Time Spent on Running the Proof-Generation Codes and the Proof Checker for LLVM 3.7.1

	Results			Time (sec.)			
	#V	#F	#NS	Orig	PCal	I/O	PCheck
mem2reg	76.84K	0	10.63K	13.20	388.49	13.81K	20.62K
gvn	285.82K	134	8.68K	49.93	214.29	37.99K	31.80K
licm	181.53K	0	30.70K	24.32	900.57	64.48K	12.99K

Figure 9. Experimental Results for LLVM 5.0.1 before GVN patch

A Experimental Results

Fig. 7 and Fig. 8 shows the validation results and the time spent on running the proof-generation codes and the proof checker for each benchmark program and optimization pass for LLVM 3.7.1. Fig. 9, Fig. 10, and Fig. 11 show the experimental results for LLVM 5.0.1 without gvn bug[7] Patch. Fig. 9 is a summary of the entire experimental results. Fig. 10 and Fig. 11 represents the validation results and the time spent on the proof-generation codes and the proof checker for each benchmark program

	LOC	Result								
		mem2reg			GVN			LICM		
		#V	**#F**	**#NS**	**#V**	**#F**	**#NS**	**#V**	**#F**	**#NS**
400.perlbench	168.16K	1.75K	0	1	9.24K	0	0	2.47K	0	142
401.bzip2	8.29K	90	0	0	1.60K	0	0	489	0	42
403.gcc	517.52K	5.43K	0	7	25.28K	2	0	9.11K	0	1.28K
429.mcf	2.69K	24	0	0	76	0	0	46	0	8
433.milc	15.04K	235	0	2	1.60K	0	0	1.70K	0	135
445.gobmk	196.24K	2.64K	0	1	4.96K	0	0	2.56K	0	361
456.hmmer	35.99K	558	0	0	2.86K	0	0	2.72K	0	187
458.sjeng	13.85K	130	0	0	1.22K	0	0	407	0	68
462.libquantum	4.36K	123	0	79	247	0	178	321	0	252
464.h264ref	51.58K	532	0	0	12.17K	0	0	9.38K	0	1.45K
470.lbm	1.16K	21	0	0	22	0	0	55	0	0
482.sphinx3	25.09K	365	0	0	1.49K	0	2	1.33K	0	244
sendmail-8.15.2	138.68K	536	0	403	3.60K	5	106	1.88K	0	212
emacs-25.1	463.54K	5.16K	0	12	19.14K	8	26	7.45K	0	538
python-3.4.1	486.38K	8.79K	0	74	22.77K	12	26	9.97K	0	540
gimp-2.8.18	1004.20K	19.44K	0	550	27.74K	88	2.55K	25.84K	0	5.75K
ghostscript-9.14.0	797.65K	13.03K	0	9.20K	50.11K	2	5.57K	38.68K	0	7.93K
LLVM nightly test	1358.76K	17.99K	0	297	101.70K	17	217	67.12K	0	11.56K
Total	5289.18K	76.84K	0	10.63K	285.82K	134	8.68K	181.53K	0	30.70K

Figure 10. Validation Results for LLVM 5.0.1 before GVN patch

	Time (sec.)											
	mem2reg				GVN				LICM			
	Orig	**PCal**	**I/O**	**PCheck**	**Orig**	**PCal**	**I/O**	**PCheck**	**Orig**	**PCal**	**I/O**	**PCheck**
400.perlbench	0.38	17.81	0.19K	0.74K	1.70	6.87	1.25K	3.09K	0.30	9.65	0.94K	0.31K
401.bzip2	0.02	1.81	0.01K	0.12K	0.12	0.56	0.16K	0.23K	0.07	2.11	0.12K	0.04K
403.gcc	1.04	48.06	5.31K	9.05K	6.47	27.71	3.10K	3.73K	1.72	29.33	3.42K	0.65K
429.mcf	<0.01	0.12	<0.01K	<0.01K	0.01	0.05	<0.01K	<0.01K	0.01	0.04	<0.01K	<0.01K
433.milc	0.05	1.08	0.01K	0.01K	0.12	0.46	0.07K	0.03K	0.08	0.39	0.08K	0.03K
445.gobmk	0.29	9.84	1.88K	0.93K	0.82	3.50	0.33K	0.25K	0.32	3.45	0.40K	0.09K
456.hmmer	0.11	3.40	0.03K	0.07K	0.36	1.49	0.10K	0.08K	0.20	1.89	0.18K	0.05K
458.sjeng	0.03	0.95	0.01K	0.02K	0.16	0.53	0.05K	0.05K	0.06	0.51	0.04K	0.01K
462.libquantum	0.01	0.31	<0.01K	<0.01K	0.03	0.12	<0.01K	<0.01K	0.01	0.10	0.01K	<0.01K
464.h264ref	0.16	8.21	0.06K	0.38K	1.04	4.76	1.26K	1.00K	0.41	9.74	1.26K	0.48K
470.lbm	<0.01	0.20	<0.01K	<0.01K	0.01	0.03	<0.01K	<0.01K	0.01	0.05	<0.01K	<0.01K
482.sphinx3	0.06	1.68	0.02K	0.02K	0.18	0.80	0.05K	0.03K	0.10	0.71	0.07K	0.02K
sendmail-8.15.2	0.22	4.45	0.03K	<0.01K	0.57	1.89	0.43K	0.32K	0.24	3.41	0.51K	0.17K
emacs-25.1	0.92	42.93	1.16K	2.26K	3.56	15.84	3.92K	2.78K	1.05	24.75	4.34K	0.94K
python-3.4.1	1.23	43.10	1.21K	0.89K	4.46	19.89	3.98K	1.83K	0.87	19.13	3.96K	0.83K
gimp-2.8.18	2.39	55.38	1.11K	0.73K	5.49	23.77	1.44K	0.66K	2.35	21.85	2.38K	0.46K
ghostscript-9.14.0	3.27	37.32	0.34K	0.06K	8.75	35.12	4.05K	2.80K	4.67	56.15	6.10K	1.82K
LLVM nightly test	3.02	111.84	2.44K	5.31K	16.08	70.89	17.80K	14.94K	11.84	717.31	40.68K	7.09K
Total	13.20	388.49	13.81K	20.62K	49.93	214.29	37.99K	31.80K	24.32	900.57	64.48K	12.99K

Figure 11. Time Spent on Running the Proof-Generation Codes and the Proof Checker for LLVM 5.0.1 before GVN patch

and optimization pass. Fig. 12, Fig. 13, and Fig. 14 shows the corresponding result for Fig. 9, Fig. 10, and Fig. 11 for LLVM 5.0.1 with gvn bug[7] Patch. **TODO: Add 5.0.1 summary data**

	Results			Time (sec.)			
	#V	**#F**	**#NS**	**Orig**	**PCal**	**I/O**	**PCheck**
mem2reg	76.84K	0	10.63K	13.20	385.15	13.73K	20.62K
gvn	285.64K	0	8.75K	49.93	214.90	37.62K	31.90K
licm	181.53K	0	30.70K	24.32	891.71	64.18K	13.26K

Figure 12. Experimental Results for LLVM 5.0.1 after GVN patch

	LOC	**Result**								
		mem2reg			GVN			LICM		
		#V	**#F**	**#NS**	**#V**	**#F**	**#NS**	**#V**	**#F**	**#NS**
400.perlbench	168.16K	1.75K	0	1	9.24K	0	0	2.47K	0	142
401.bzip2	8.29K	90	0	0	1.60K	0	0	489	0	42
403.gcc	517.52K	5.43K	0	7	25.28K	0	1	9.11K	0	1.28K
429.mcf	2.69K	24	0	0	76	0	0	46	0	8
433.milc	15.04K	235	0	2	1.60K	0	0	1.70K	0	135
445.gobmk	196.24K	2.64K	0	1	4.96K	0	1	2.56K	0	361
456.hmmer	35.99K	558	0	0	2.85K	0	0	2.71K	0	187
458.sjeng	13.85K	130	0	0	1.22K	0	0	407	0	68
462.libquantum	4.36K	123	0	79	247	0	178	321	0	252
464.h264ref	51.58K	532	0	0	12.15K	0	0	9.38K	0	1.45K
470.lbm	1.16K	21	0	0	22	0	0	55	0	0
482.sphinx3	25.09K	365	0	0	1.49K	0	2	1.33K	0	244
sendmail-8.15.2	138.68K	536	0	403	3.60K	0	112	1.88K	0	212
emacs-25.1	463.54K	5.16K	0	12	19.12K	0	28	7.45K	0	538
python-3.4.1	486.38K	8.79K	0	74	22.77K	0	36	9.97K	0	540
gimp-2.8.18	1004.20K	19.44K	0	550	27.65K	0	2.58K	25.84K	0	5.75K
ghostscript-9.14.0	797.65K	13.03K	0	9.20K	50.11K	0	5.57K	38.68K	0	7.93K
LLVM nightly test	1358.76K	17.99K	0	297	101.65K	0	236	67.12K	0	11.56K
Total	5289.18K	76.84K	0	10.63K	285.64K	0	8.75K	181.53K	0	30.70K

Figure 13. Validation Results for LLVM 5.0.1 after GVN patch

B Miscompilation of a Realistic Program due to a Mem2reg Bug

One of the bug we found in `mem2reg` miscompiles the following program:

```
#include <stdio.h>
int sqr(int i, int prev, int cur) {
  return cur * cur;
}
int diffsqr(int i, int prev, int cur) {
  if (i==0) return 0; else return (cur-prev) * (cur-prev);
}
void foo(int arr[], int n) {
  int sqrsum = 0, diffsqrsum = 0;
  int i, prev, cur;
  for (i = 0; i < n; ++i) {
    prev = cur; cur = arr[i];
    sqrsum += sqr(i, prev, cur);
    diffsqrsum += diffsqr(i, prev, cur);
  }
  printf ("square sum=%d, diff sqr sum=%d \n", sqrsum, diffsqrsum);
}
```

	Time (sec.)											
	mem2reg				GVN				LICM			
	Orig	**PCal**	**I/O**	**PCheck**	**Orig**	**PCal**	**I/O**	**PCheck**	**Orig**	**PCal**	**I/O**	**PCheck**
400.perlbench	0.38	18.16	0.19K	0.76K	1.70	6.69	1.24K	3.10K	0.30	9.58	0.94K	0.32K
401.bzip2	0.02	1.71	0.01K	0.12K	0.12	0.52	0.16K	0.23K	0.07	2.03	0.12K	0.04K
403.gcc	1.04	47.11	5.32K	9.07K	6.47	27.12	3.16K	3.73K	1.72	29.15	3.40K	0.66K
429.mcf	<0.01	0.10	<0.01K	<0.01K	0.01	0.05	<0.01K	<0.01K	0.01	0.04	<0.01K	<0.01K
433.milc	0.05	1.10	0.01K	0.01K	0.12	0.47	0.06K	0.03K	0.08	0.37	0.08K	0.03K
445.gobmk	0.29	9.44	1.88K	0.94K	0.82	3.49	0.34K	0.26K	0.32	3.42	0.41K	0.10K
456.hmmer	0.11	3.33	0.03K	0.07K	0.36	1.57	0.11K	0.07K	0.20	1.92	0.19K	0.05K
458.sjeng	0.03	0.92	0.01K	0.02K	0.16	0.50	0.05K	0.05K	0.06	0.44	0.04K	0.01K
462.libquantum	0.01	0.30	<0.01K	<0.01K	0.03	0.12	<0.01K	<0.01K	0.01	0.10	0.01K	<0.01K
464.h264ref	0.16	8.14	0.06K	0.39K	1.04	4.80	1.24K	0.96K	0.41	9.75	1.24K	0.48K
470.lbm	<0.01	0.21	<0.01K	<0.01K	0.01	0.03	<0.01K	<0.01K	0.01	0.05	<0.01K	<0.01K
482.sphinx3	0.06	1.70	0.02K	0.02K	0.18	0.77	0.05K	0.04K	0.10	0.67	0.07K	0.02K
sendmail-8.15.2	0.22	4.44	0.03K	<0.01K	0.57	1.92	0.43K	0.32K	0.24	3.35	0.52K	0.18K
emacs-25.1	0.92	42.31	1.13K	2.25K	3.56	16.12	3.85K	2.79K	1.05	24.88	4.25K	0.96K
python-3.4.1	1.23	42.86	1.22K	0.87K	4.46	20.03	3.91K	1.80K	0.87	19.04	3.97K	0.84K
gimp-2.8.18	2.39	54.88	1.11K	0.74K	5.49	24.08	1.45K	0.61K	2.35	21.82	2.35K	0.45K
ghostscript-9.14.0	3.27	37.51	0.33K	0.06K	8.75	35.35	4.00K	2.89K	4.67	55.54	6.16K	1.84K
LLVM nightly test	3.02	110.94	2.38K	5.30K	16.08	71.29	17.57K	15.04K	11.84	709.57	40.44K	7.27K
Total	13.20	385.15	13.73K	20.62K	49.93	214.90	37.62K	31.90K	24.32	891.71	64.18K	13.26K

Figure 14. Time Spent on Running the Proof-Generation Codes and the Proof Checker for LLVM 5.0.1 after GVN patch

```
int main () {
  int a[3] = {1, 2, 5};
  foo(a, 3);
  return 0;
}
```

The function `foo()` takes an array `a` and its size `n` and prints the sum of the squares of the numbers in `a`, and the sum of the squares of the differences of adjacent numbers in `a`, *i.e.*, :

$$\sum_{i=0}^{n-1} \texttt{a[i]}^2 \quad \text{and} \quad \sum_{i=1}^{n-1} (\texttt{a[i]} - \texttt{a[i-1]})^2 \,.$$

The function `foo()` calculates the two summations in a single loop. Note that the function `diffsqr(i, prev, cur)` returns $(\texttt{cur} - \texttt{prev})^2$ if $\texttt{i} > 0$, and zero otherwise.

Clang 3.7.1 with `-O2` flag miscompiles this program. The compilation result prints out 30 and 0 instead of the correct answer 30 ($= 1^2 + 2^2 + 5^2$) and 10 ($= (2-1)^2 + (5-2)^2$). This is due to the `mem2reg` bug we found on the special case that all stores to a promotable location is in a single block. In essence, the loads from the local variable `prev` are *illegally* promoted to `undef`, the function call to `diffsqr()` is inlined, and exploiting the `undef` semantics, the result of the inlined call is optimized to 0.

Even though the program invokes undefined behavior according to the C standard, we believe it is still likely to be written in the real-world: a programmer may logically conclude that the `undef` value is never used, and the program is safe.

C Global Value Numbering with Partial Redundancy Elimination

Global Value Numbering optimization (GVN), which is implemented in the `gvn` pass of LLVM, detects and removes redundant instructions. GVN algorithm first groups expressions and values into equivalent classes and assigns a unique "value number" to each class. Then, a leader value is chosen for each class, and all the non-leader, redundant, instructions are substituted with the leader value. The `gvn` pass also does Partial Redundancy Elimination optimization (PRE), which eliminates instructions that are partially redundant depending on the control flow.

In this section, we show how we generate and validate proofs for GVN-PRE optimization. It is worth noting that even though the GVN and PRE algorithms are separately written in the `gvn` pass, their validation logics are so similar that the resemblance of their validation logics enabled us to write a single proof generation code that uniformly works for both GVN and PRE.

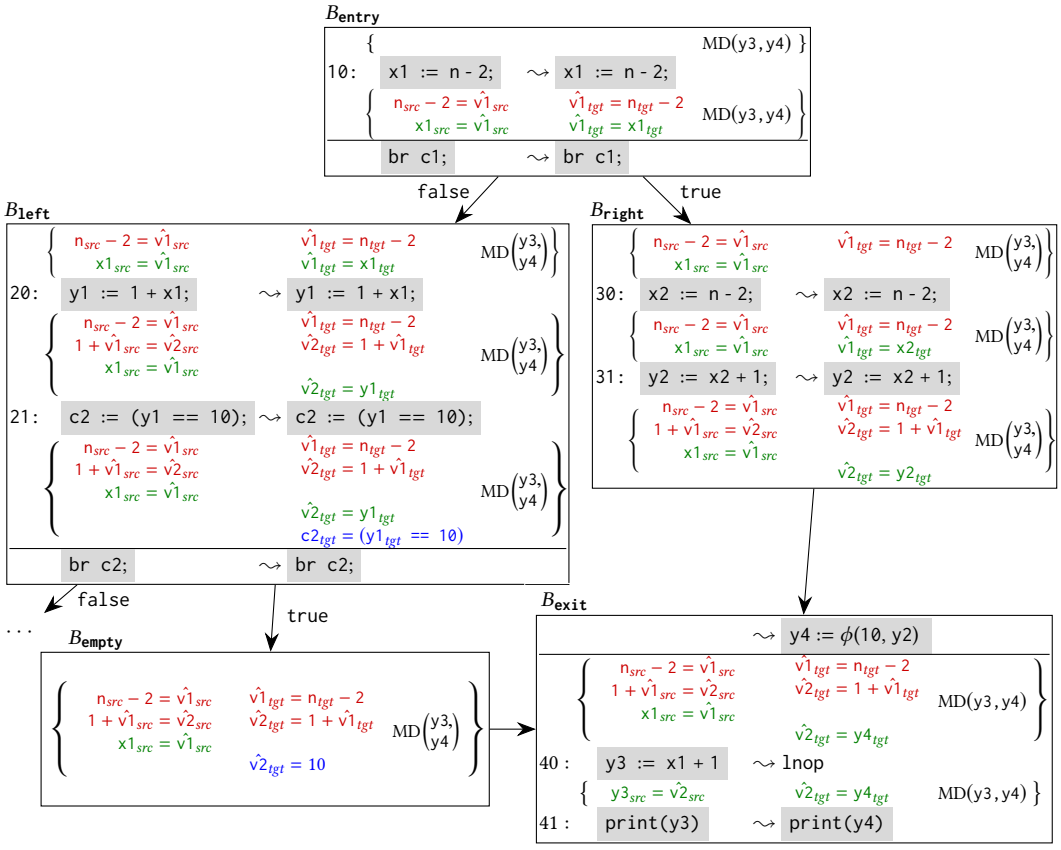

Figure 15. A PRE Example

C.1 Translation Example

We use a PRE translation example because it is more interesting than a GVN example. The shaded part of Fig. 15 shows an example translation of the PRE optimization, where the registers n and c1 are the function parameters.

First, by analyzing the source program, GVN assigns unique value numbers to the classes of equivalent values and expressions. For example, in Fig. 15, GVN constructs the following tables, VT, ET and LT.

$$VT = [\mathtt{x1}, \mathtt{x2} \mapsto ①; \quad \mathtt{y1}, \mathtt{y2}, \mathtt{y3} \mapsto ②]$$
$$ET = [n - 2 \mapsto ①; \quad 1 + ① \mapsto ②]$$
$$LT = [\ldots; \quad \mathbf{B}_{\mathtt{empty}} \mapsto [① \mapsto \mathtt{x1};\ ② \mapsto 10]; \quad \mathbf{B}_{\mathtt{right}} \mapsto [① \mapsto \mathtt{x2};\ ② \mapsto \mathtt{y2}]; \quad \ldots]$$

Originally, The gvn pass also assigns value numbers to singleton classes, but in this example, we only consider those classes with more than one value for brevity.

Here, the *value table* VT assigns value number ① to x1 and x2, and ② to y1, y2 and y3; and the *expression table* ET assigns value number ① to the expression $n - 2$, and ② to $1 + ①$. This means that, at any point of execution, there *exist* some values for ① and ② such that x1, x2 and $n - 2$ evaluate to the value ① and y1, y2, y3 and $1 + ①$ evaluate to the value ② whenever they are well-defined. It is easy to see that indeed this property holds for the source program in Fig. 15. Finally, the *leader table* LT determines the leader value for each value number among the values in each block. Note that the leader values can be different for each block and not every block necessarily has a leader value for every value number. For example, in the block $\mathbf{B}_{\mathtt{entry}}$, the leader value for ② does not exist because none of the values with value number ② are defined in the block.

In the example, PRE detects partial redundancies of the instruction `y3 := x1 + 1` in $\mathbf{B}_{\mathtt{exit}}$ because y3 belongs to ② and $\mathbf{B}_{\mathtt{exit}}$'s incoming blocks have a leader value for ②: 10 in $\mathbf{B}_{\mathtt{empty}}$ and y2 in $\mathbf{B}_{\mathtt{right}}$. In other words, depending on the control flow, the expression y3 is equivalent to either 10 or y2. To eliminate the redundant instruction, PRE (*i*) inserts the phinode `y4 := φ(10,y2)` at the beginning of $\mathbf{B}_{\mathtt{exit}}$; (*ii*) replaces all uses of y3 with y4 at line 41; and (*iii*) eliminates `y3 := x1 + 1` at line 40. This translation is beneficial because the inserted phinode is compiled down to a move instruction, which is more cost-effective than the eliminated addition instruction.

C.2 ERHL Proof

We can turn our intuition for soundness into a formal proof. The unshaded part of Fig. 15 shows the proof, each assertion of which can be split into three parts.

Expression Assertions The red equalities of the form $e_{src} = \hat{v}i_{src}$ or $\hat{v}i_{tgt} = e_{tgt}$ relate expressions with their value numbers according to the expression table ET. For example, the assertion after line 21 states that there exist some values $\hat{v}1$ and $\hat{v}2$ (for ① and ②) such that $\hat{v}1 = n - 2$ and $\hat{v}2 = 1 + \hat{v}1$ hold for both the source and target states.

Value Assertions The green equalities of the form $x_{src} = \hat{v}i_{src}$ or $\hat{v}i_{tgt} = x_{tgt}$ relate values with their value numbers according to the value table VT. For example, the assertion after line 21 also states that x1 has the value $\hat{v}1$ in the source and y1 has the value $\hat{v}2$ in the target.

Branching Assertions The blue equalities show properties derived from branching conditions. For example, the equality $\texttt{c2}_{tgt} = (\texttt{y1}_{tgt} \texttt{ == } \texttt{10})$ at line 21 is derived from the definition of the branching register c2, and the equality $\hat{v}2_{tgt} = 10$ in the block $\mathbf{B}_{\texttt{empty}}$ is derived from the fact that c2 must be true in $\mathbf{B}_{\texttt{empty}}$ and thus y1 is 10 and so is its associated value $\hat{v}2$.

C.3 Proof Validation

The proof in Fig. 15 is validated as follows.

Adding Value Assertions At line 40 (and similarly at line 30), the proof checker computes a post-assertion by adding $\{\, \texttt{y3}_{src} = \texttt{x1}_{src} + 1 \,\}$, and deduces, by applying the inference rules from the proof, that:

$$\begin{aligned}
\texttt{y3}_{src} &= \texttt{x1}_{src} + 1 && \text{(post-assertion)}\\
&= \hat{v}1_{src} + 1 && \text{(substitution}(\texttt{x1}_{src} + 1, \texttt{x1}_{src}, \hat{v}1_{src})\text{)}\\
&= 1 + \hat{v}1_{src} && \text{(commutativity_add}(\hat{v}1_{src}, 1)\text{)}\\
&= \hat{v}2_{src} && \text{(expression assertion)}
\end{aligned}$$

from which the next assertion trivially follows.

Adding Expression Assertions At line 31 (and similarly at lines 10 and 20), the proof checker computes a post-assertion by adding $\{\, \texttt{y2}_{src} = \texttt{x2}_{src} + 1, \texttt{x2}_{tgt} + 1 = \texttt{y2}_{tgt} \,\}$. Then the proof gives the inference rule $\texttt{intro\_ghost}(1 + \hat{v}1, \hat{v}2)$, which adds $\{\, 1 + \hat{v}1_{src} = \hat{v}2_{src}, \hat{v}2_{tgt} = 1 + \hat{v}1_{tgt} \,\}$. Then the proof checker deduces $\hat{v}2_{tgt} = \texttt{y2}_{tgt}$ similarly as in line 40.

Adding Phinode Assertions At the phinode of $\mathbf{B}_{\texttt{exit}}$, the assertion is separately validated for each incoming block. For the incoming block $\mathbf{B}_{\texttt{left}}$ (and similarly for $\mathbf{B}_{\texttt{right}}$), the proof checker computes a post-assertion by adding $\texttt{10} = \texttt{y4}_{tgt}$ and adding y to the maydiff set, and derives $\hat{v}2_{tgt} = \texttt{10} = \texttt{y4}_{tgt}$ by applying the transitivity rule from the proof.

Adding Branching Assertions At line 21, the proof checker computes a post-assertion by adding $\{\, \texttt{c2}_{src} = (\texttt{y1}_{tgt} \texttt{ == } \texttt{10}), \texttt{c2}_{tgt} = (\texttt{y1}_{tgt} \texttt{ == } \texttt{10}) \,\}$, from which the next assertion trivially follows.

Using Branching Assertions At the beginning of $\mathbf{B}_{\texttt{empty}}$, the proof checker computes a post-assertion by adding $\{\, \texttt{true} = \texttt{c2}_{src}, \texttt{true} = \texttt{c2}_{tgt} \,\}$ from the branching condition, and deduces, by applying the inference rules from the proof, that $\texttt{true} = \texttt{c2}_{tgt} = (\texttt{y1}_{tgt} \texttt{ == } \texttt{10})$ and

$$\begin{aligned}
\hat{v}2_{tgt} &= \texttt{y1}_{tgt} && \text{(value assertion)}\\
&= \texttt{10} && \text{(from } \texttt{true} = (\texttt{y1}_{tgt} \texttt{ == } \texttt{10}) \text{ by } \texttt{icmp_to_eq}(\texttt{true}, \texttt{y1}_{tgt}, \texttt{10})\text{)}
\end{aligned}$$

from which the next assertion trivially follows.

C.4 Proof Generation

Now we explain how we generate proofs for the GVN-PRE optimization.

Algorithm 3 ProofGen(F: Function, r: Register, v: Value, VT, RET, RPT, BCT)

A1: $(l_r: r := \_)$:= FindDef(F, r)
A2: WL := $[(l_r, r, VT[r], src), (l_r, v, VT[r], tgt)]$
A3: **while** NonEmpty(WL) **do**
A4: $((l_0, x, \hat{n}, side) :: WL) := WL$
A5: **match** FindDef(F, x) **with**
A6: | Some $(l: x := e) \Rightarrow$
A7: Assn($RET[x]$, l, l_0), Assn($x_{side} = \hat{n}_{side}$, l, l_0)
A8: **for** $(l', y, \hat{m})$ **in** MatchExpr(e, $RPT[x]$) **do**
A9: $WL := (l', y, \hat{m}, side) :: WL$
A10: **end for**
A11: | Some (ConstantOrParameter(C)) $\Rightarrow$
A12: **match** FindBranchingCondition($BCT[(C, \hat{n})]$, l_0, F) **with**
A13: | Some $(v, y, c, l) \Rightarrow$
A14: **if** FindDef(F, c) = Some $(l_c: c := e)$ **then** Assn($c_{side} = e_{side}$, l_c, l) **end if**
A15: Inf(`icmp_to_eq`(v,y_{tgt},C), l)
A16: Assn($RET[y]$, l, l_0), Assn($C_{side} = \hat{n}_{side}$, l, l_0)
A17: $WL := (l, y, \hat{n}, side) :: WL$
A18: **end match**
A19: **end match**
A20: **end while**
A21: Auto(`GVN_PRE`)

We first add code to GVN and PRE algorithms that generate auxiliary data, RET, RPT and BCT. For example, we compute the following for Fig. 15:

$$\begin{aligned} RET &= [\mathtt{x1}, \mathtt{x2} \mapsto \{\, n_{src} - 2 = \hat{\mathtt{v1}}_{src}, \hat{\mathtt{v1}}_{tgt} = n_{tgt} - 2 \,\}; \\ &\qquad \mathtt{y1}, \mathtt{y2}, \mathtt{y3}, \mathtt{y4} \mapsto \{\, n_{src} - 2 = \hat{\mathtt{v1}}_{src}, \hat{\mathtt{v1}}_{tgt} = n_{tgt} - 2, 1 + \hat{\mathtt{v1}}_{src} = \hat{\mathtt{v2}}_{src}, \hat{\mathtt{v2}}_{tgt} = 1 + \hat{\mathtt{v1}}_{tgt} \,\}] \\ RPT &= [\mathtt{x1}, \mathtt{x2} \mapsto (\_, \_); \quad \mathtt{y1} \mapsto (\_, ①); \quad \mathtt{y2}, \mathtt{y3} \mapsto (①, \_); \quad \mathtt{y4} \mapsto (②, ②)] \\ BCT &= [(\mathtt{10}, ②) \mapsto \{\, (\mathbf{B}_{\mathtt{left}}, \mathtt{true}, \mathtt{y1}) \,\}] \,. \end{aligned}$$

The *register expression table* RET contains sufficient expression assertions for reasoning about each register. For example, $RET(\mathtt{x1})$ contains $\{\, n_{src} - 2 = \hat{\mathtt{v1}}_{src}, \hat{\mathtt{v1}}_{tgt} = n_{tgt} - 2 \,\}$, which is sufficient for reasoning about `x1`. The *register parameter table* RPT contains the value numbers of the instruction parameters for each register. For example, $RPT[\mathtt{y1}]$ contains $(\_, ①)$, which means that the second parameter `x1` of the instruction 20: `y1 := 1 + x1` has the value number ①. The *branching condition table* BCT maps pairs of constants/function parameters and their value numbers to their associated branching information. For example, $BCT[(\mathtt{10}, ②)]$ contains $(\mathbf{B}_{\mathtt{left}}, \mathtt{true}, \mathtt{y1})$, which means that (i) `y1 = 10` holds when control flows from $\mathbf{B}_{\mathtt{left}}$ with the branching condition being `true` (*i.e.,* `c2 = true`), and (ii) `y1`'s value number is ②. We can easily construct RET, RPT, and BCT following the construction of VT and ET.

Algorithm 3 presents the common proof-generation code that generate assertions for both GVN and PRE, which is given in a rather functional style for presentation purposes. Specifically, ProofGen(F, r, v, VT, RET, RPT, BCT) generates a proof for the replacement of r with v in function F. In essence, the proof generation code adds assertions, starting from those for r_{src} and v_{tgt} at r's definition point and recursively down to the arguments, using a worklist and auxiliary data RET, RPT, and BCT. The recursion stops when the target value is the only value of its value number. More concretely, the algorithm works for the example of Fig. 15 as follows.

Initialize Worklist The code finds where r is defined (line A1), and add a work for the source and another for the target to the worklist (line A2). A work $(l_0, x, \hat{n}, side)$ means that (i) x's value number is $\hat{n}$, and (ii) line l_0 needs the assertion $x_{side} = \hat{n}_{side}$ (where $side$ is either src or tgt) and the expression assertions. For the translation in Fig. 15, x = `y3` is defined at line 40 in the source, so the initial works are $(40, \mathtt{y3}, \hat{\mathtt{v2}}, src)$ and $(40, \mathtt{y4}, \hat{\mathtt{v2}}, tgt)$. The code processes each work $(l_0, x, \hat{n}, side)$ as follows (lines A3-A4).

Processing Registers If x is a register defined as e at l (line A6), the code adds the expression assertions $RET[x]$ for x from l to l_0, and the value assertion $x = \hat{n}$ at $side$ (line A7). Consider the work $(40, \mathtt{y4}, \hat{\mathtt{v2}}, tgt)$, for example. The register `y4` is defined by the phinode of $\mathbf{B}_{\mathtt{exit}}$, so l is the phinode of $\mathbf{B}_{\mathtt{exit}}$ and $e = \phi(10, \mathtt{y2})$. Hence $RET[\mathtt{y4}]$ and $\hat{\mathtt{v2}}_{tgt} = \mathtt{y4}_{tgt}$ are inserted from the phinode of $\mathbf{B}_{\mathtt{exit}}$ to line 40. For the work $(40, \mathtt{y3}, \hat{\mathtt{v2}}, src)$, since `y3` is defined at line 40, no assertions are inserted.

In order to justify the inserted assertions at line l, the code adds sub-works for the values in e (lines A8-A9). Note that MatchExpr(e, $RPT[x]$) matches e against the register parameter table in order to know which value, say y, should have which value number, say $\hat{m}$, at line l' (line A8). For example, in order to deduce $\texttt{y3}_{src} = \hat{\texttt{v2}}_{src}$ after line 40, $\texttt{x1}_{src} = \hat{\texttt{v1}}_{src}$ should hold before the line, so the code adds (40, x1, $\hat{\texttt{v1}}$, *src*) to the worklist; for justifying $\hat{\texttt{v2}}_{tgt} = \texttt{y4}_{src}$ after the phinode of $\mathbf{B}_{\texttt{exit}}$, the code adds ([edge from $\mathbf{B}_{\texttt{empty}}$ to $\mathbf{B}_{\texttt{exit}}$]), 10, $\hat{\texttt{v2}}$, *tgt*) and ([edge from $\mathbf{B}_{\texttt{right}}$ to $\mathbf{B}_{\texttt{exit}}$], y2, $\hat{\texttt{v2}}$, *tgt*) to the worklist.

Processing Constants and Parameters If x is a constant or parameter C (line A11), the code looks for a branching condition in $BCT[(C, \hat{n})]$, which justifies that C has the value number $\hat{n}$ at l_0 (line A12). The result (v, y, c, l) means $y = C$ holds at l thanks to the branching condition c being equal to v.

The code adds the branching assertion when necessary (line A14), and also adds an `icmp_to_eq` inference rule (line A15). Also, similarly to the case of registers, the code adds expression and branching assertions (line A16) to the proof, and adds sub-works for the value y (line A17).

For example, consider the work (E_e, 10, $\hat{\texttt{v2}}$, *tgt*), where E_e is the edge from $\mathbf{B}_{\texttt{empty}}$ to $\mathbf{B}_{\texttt{exit}}$. The code finds a branching condition in $BCT[(10, \hat{\texttt{v2}})]$ which guarantees that $10 = \hat{\texttt{v2}}$ holds at E_e (line A12). Block $\mathbf{B}_{\texttt{left}}$ has such a branch when c2 equals to true ($(v, y, c, l) = (\texttt{true}, \texttt{y1}, \texttt{c2}, E_l)$ at line 13 where E_l is the edge $\mathbf{B}_{\texttt{left}}$ to $\mathbf{B}_{\texttt{empty}}$). Hence the code adds $\texttt{c2}_{tgt} = (\texttt{y1}_{tgt} == 10)$ from line 21 to E_l (line A14), and adds the inference rule $\texttt{icmp\_to\_eq}(\texttt{true}, \texttt{y1}_{tgt}, 10)$ at E_l (line A15). Then the code adds the expression assertions $RET[\texttt{y1}]$ for y1 and the branching assertion $\hat{\texttt{v2}}_{tgt} = 10$ from E_l to E_e (line A16), and adds the work (E_l, y1, $\hat{v2}$, *tgt*) to the worklist (line A16).

Automation Function Throughout the proof we use the `GVN_PRE` automation function, which adds `intro_ghost`, `commutativity`, and `substitution` in a specific way for GVN-PRE, and adds `transitivity` and `reduce_maydiff` in the same way as for `assoc-add` and `mem2reg`.

D Validation Coverage

***Code Coverage in* `licm`** We wrote proof-generation code for `licm` pass implemented in `lib/Transforms/Scalar/LICM.cpp`. We covered all functions except `promoteLoopAccessesToScalars`, because it uses alias analysis.

***Micro-Optimizations in* `instcombine`** We validated the following 139 micro-optimizations in `instcombine`:
`add-comm-sub`, `add-const-not`, `add-dist-sub`, `add-mask`, `add-onebit`, `add-or-and`, `add-select-zero`, `add-shift`, `add-signbit`, `add-sub`, `add-xor-and`, `add-zext-bool`, `and-de-morgan`, `and-mone`, `and-not`, `and-or-const2`, `and-or-not1`, `and-or`, `and-same`, `and-undef`, `and-xor-const`, `and-zero`, `bitcast-bitcast`, `bitcast-fpext`, `bitcast-fptosi`, `bitcast-fptoui`, `bitcast-fptrunc`, `bitcast-inttoptr`, `bitcast-ptrtoint`, `bitcast-sametype`, `bitcast-sext`, `bitcast-sitofp`, `bitcast-trunc`, `bitcast-uitofp`, `bitcast-zext`, `bop-associativity`, `dead-code-elim`, `dead-store-elim`, `fold-phi-bin-const`, `fold-phi-bin`, `fpext-bitcast`, `fpext-fpext`, `fptosi-bitcast`, `fptosi-fpext`, `fptoui-bitcast`, `fptoui-fpext`, `fptrunc-bitcast`, `fptrunc-fpext`, `icmp-eq-add-add`, `icmp-eq-srem`, `icmp-eq-sub-sub`, `icmp-eq-sub`, `icmp-eq-xor-not`, `icmp-eq-xor-xor`, `icmp-ne-add-add`, `icmp-ne-srem`, `icmp-ne-sub-sub`, `icmp-ne-sub`, `icmp-ne-xor-xor`, `icmp-ne-xor`, `icmp-sge-or-not`, `icmp-sgt-and-not`, `icmp-sle-or-not`, `icmp-slt-and-not`, `icmp-swap`, `icmp-uge-or-not`, `icmp-ugt-and-not`, `icmp-ule-or-not`, `icmp-ult-and-not`, `inttoptr-bitcast`, `inttoptr-ptrtoint`, `mul-bool`, `mul-mone`, `mul-neg`, `mul-shl`, `or-and-xor`, `or-and`, `or-mone`, `or-not`, `or-or2`, `or-or`, `or-same`, `or-undef`, `or-xor2`, `or-xor3`, `or-xor4`, `or-xor`, `or-zero`, `ptrtoint-bitcast`, `ptrtoint-inttoptr`, `sdiv-mone`, `select-bop-fold`, `select-icmp-eq-xor1`, `select-icmp-eq-xor2`, `select-icmp-eq`, `select-icmp-gt-const`, `select-icmp-lt-const`, `select-icmp-ne-xor1`, `select-icmp-ne-xor2`, `select-icmp-ne`, `select-icmp-sgt-xor1`, `select-icmp-sgt-xor2`, `select-icmp-slt-xor1`, `select-icmp-slt-xor2`, `sext-bitcast`, `sext-sext`, `sext-trunc-ashr`, `sext-zext`, `shift-undef1`, `shift-undef2`, `shift-zero1`, `shift-zero2`, `sitofp-bitcast`, `sitofp-sext`, `sitofp-zext`, `sub-add`, `sub-const-add`, `sub-const-not`, `sub-mone`, `sub-onebit`, `sub-or-xor`, `sub-remove`, `sub-shl`, `sub-sub`, `trunc-bitcast`, `trunc-onebit`, `trunc-sext`, `trunc-trunc`, `trunc-zext`, `uitofp-bitcast`, `uitofp-zext`, `xor-same`, `xor-undef`, `xor-zero`, `zext-bitcast`, `zext-trunc-and-xor`, `zext-trunc-and`, `zext-xor`, `zext-zext`

Note that we gave these names and they are not officially used in LLVM.

E Program Points between Two Lines

Assertion(P, l_1, l_2) in the proof generation code means predicate P should be added to the assertions between l_1 and l_2. More specifically, the proposition P should be added at every program point appearing in a path from l_1 to l_2 that does not visit l_1 but may visit l_2 in-between. Since l_1 is the source of the proposition P so that we can get P as a post-assertion every time we visit l_1, there is no need to add P along a path from l_1 to l_1. For example, consider the following program.

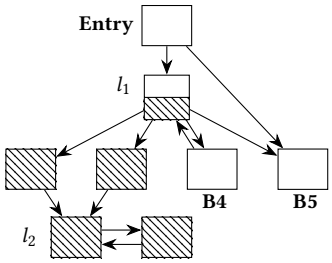

The marked area between l_1 and l_2 is where we should add the proposition P.

Thanks to the SSA property [18], we can efficiently calculate the program points between l_1 and l_2. First, we can assume that l_1 *dominates* l_2 (*i.e.,* one should have visited l_1 to reach l_2 from the entry point), since the proposition P created at l_1 should hold at l_2. Then a program point l is on a path from l_1 to l_2 that does not visit l_1 in-between if and only if (i) l_1 dominates l and (ii) l_2 is reachable from l without visiting l_1. We efficiently check the first condition using the *dominator tree* [18] and the second condition by a backward BFS search from l_2.

In the above example, any program point in the marked area is such that l_1 dominates it and l_2 is reachable from it without visiting l_1. For example, l_2 is not reachable from the block B_4 without visiting l_1, and l_1 does not dominate the block B_5.

Note that this algorithm is not a part of TCB: validation may fail but cannot succeed incorrectly due to bugs in this algorithm.

F Lessdef Predicates

LLVM and CompCert has the notion of `undef` value, which is designated as the result of erroneous operations. The compilers are allowed to replace `undef` by an arbitrary value. For example, LLVM's InstCombine performs the following translation, where the register `y` is replaced by 1, presumably because $\mathtt{y}_{src} = \mathtt{a}_{src} - \mathtt{x}_{src} = \mathtt{a}_{src} - (\mathtt{a}_{src} - 1) = 1$ holds for any integer value of $\mathtt{a}_{src}$:

```
x := a - 1;        x := a - 1;
y := a - x;   ⇝    y := a - x;
z := y + 1;        z := 1 + 1;
```

However, if $\mathtt{a}_{src}$ is `undef`, then we have $\mathtt{z}_{src} = (\mathtt{a}_{src} - (\mathtt{a}_{src} - 1)) + 1 =$ `undef` `undef` is propagated in the arithmetic. Thus the equation $\mathtt{z}_{src} = \mathtt{z}_{tgt}$ no longer holds, breaking the equality relationship between the source and target values.

In order to reason such optimizations, we use the CompCert-style *lessdef* relation [28] throughout this work instead of the equality as assertion predicates. Concretely, x is less defined than y, denoted $x \sqsupseteq y$, if x is `undef` or it equals to y. For example, we have $\mathtt{y}_{src} = \mathtt{a}_{src} - \mathtt{x}_{src} = \mathtt{a}_{src} - (\mathtt{a}_{src} - 1) \sqsupseteq 1$ and thus $\mathtt{z}_{src} = \mathtt{y}_{src} + 1 \sqsupseteq 1 + 1 = \mathtt{z}_{tgt}$ regardless of whether $\mathtt{a} =$ `undef` or not, which justifies the above translation.

So far we used the equality instead of lessdef relation for simplicity of presentation. However, all the assertions presented in this paper works even when equalities are replaced by lessdef relations. Note that the post-assertion generator should be adapted to lessdef relations so that they introduce both $x \sqsupseteq e$ and $e \sqsupseteq x$ after $x := e$.

G Semantic Interpretation of Assertions

The syntax of assertions is as follows:

$$
\begin{aligned}
\mathit{Reg} \ni r &::= \cdots \\
\mathit{Const} \ni c &::= \cdots \\
\mathit{Typ} \ni \mathit{typ} &::= \mathbf{int} \mid *\,\mathit{typ} \mid \cdots \\
\mathit{SVal} \ni v &::= r \mid c \\
\mathit{Tag} \ni \mathit{tag} &::= \mathit{Phy} \mid \mathit{Ghost} \mid \mathit{Old} \\
\mathit{Reg} \times \mathit{Tag} \ni rT &::= (r, \mathit{tag}) \\
\mathit{SVal} \times \mathit{Tag} \ni vT &::= (v, \mathit{tag}) \\
\mathit{Expr} \ni e &::= \mathsf{add}\ vT\ vT \mid \mathsf{load}\ vT\ \mathit{typ}\ \mathit{align} \mid \cdots \\
\mathit{Pred} \ni \mathit{pred} &::= e \sqsupseteq e \mid \mathrm{Uniq}(r) \mid \mathrm{Priv}(rT) \mid vT \perp vT \\
\mathit{MD} \ni M &::= \{rT, \cdots\} \\
\mathit{AssnU} \ni S, T &::= \{\mathit{pred}, \cdots\} \\
\mathit{Assn} \ni P, Q &::= (S, T, M)
\end{aligned}
$$

Reg, *Const*, and *Typ* are the types of LLVM registers, constants, and types. A (static) value is either a register or a constant. Registers and values may be tagged (see /coq/def/Exprs.v:106): the tag *Phy* means physical registers, *Ghost* means regular ghost registers (§3.2), and *Old* means old ghost registers (§4). We write r for (r, Phy), $\hat{r}$ for (r, Ghost), and $\bar{r}$ for (r, Old).

An expression is basically the right-hand side of a side-effect-free LLVM instructions with operand values being tagged (see /coq/def/Exprs.v:366 in the Coq development[9]). Notice that load is side-effect-free (except for undefined behavior) so there are load expressions, while store is side-effectful and there are no store expressions. Recall that $e_1 \sqsupseteq e_2$ means either $e_1 = e_2$ or $e_1 = \texttt{undef}$, and the predicate $vT_1 \perp vT_2$ means the addresses in vT_1 and vT_2 point to disjoint memory blocks. A unary assertion is a set of predicates (see /coq/def/Hints.v:34), and a maydiff set is a set of tagged registers. An assertion consists of a unary assertion for source, another for target, and a maydiff set (see /coq/def/Hints.v:41).

We use the following semantic domains:

$$\begin{aligned} Val \ni V &::= \cdots \\ RF \ni rs &::= \{r \mapsto V, \cdots\} \\ State \ni \sigma &::= \cdots \\ StateT \ni \sigma T &::= (\sigma, rs, rs) \\ Meminj \ni \alpha &::= \cdots \end{aligned}$$

Let Val be the set of (dynamic) values, RF be the set of register files, $State$ be the set of states, $StateT$ be the set of *extended states*, which are tuples of a state, a register file for ghost registers, and another for old ghost registers. $Meminj$ is the set of CompCert-style memory injections, which basically maps a source memory block to the equivalent target block [28].

The semantics of assertions is as follows:

$$\begin{aligned} [\![r]\!], [\![c]\!], [\![v]\!] &: State \to Val \stackrel{\text{def}}{=} \cdots \\ [\![rT]\!], [\![vT]\!] &: StateT \to Val \stackrel{\text{def}}{=} \cdots \end{aligned}$$

$$\begin{aligned} [\![pred]\!] &: 2^{Blk} \to StateT \to \mathbb{P} \\ [\![e_1 \sqsupseteq e_2]\!](priv, \sigma T) &\stackrel{\text{def}}{=} [\![e_1]\!](\sigma T) \sqsupseteq [\![e_2]\!](\sigma T) \\ [\![\text{Uniq}(r)]\!](priv, \sigma T) &\stackrel{\text{def}}{=} \forall b, b', o, o'.\ [\![r]\!](\sigma T) = (b, o) \land (b', o') \in \sigma(T \setminus r) \implies b \neq b' \\ [\![\text{Priv}(rT)]\!](priv, \sigma T) &\stackrel{\text{def}}{=} \forall b, o.\ [\![rT]\!](\sigma T) = (b, o) \implies b \in priv \\ [\![vT_1 \perp vT_2]\!](priv, \sigma T) &\stackrel{\text{def}}{=} \forall b_1, b_2, o_1, o_2.\ [\![vT_1]\!](\sigma T) = (b_1, o_1) \land [\![vT_2]\!](\sigma T) = (b_2, o_2) \implies b_1 \neq b_2 \end{aligned}$$

$$\begin{aligned} [\![M]\!](\alpha, \sigma T_{src}, \sigma T_{tgt}) &\stackrel{\text{def}}{=} \forall rT \notin M.\ [\![rT]\!](\sigma T_{src}) \sim_\alpha [\![rT]\!](\sigma T_{tgt}) \\ [\![S]\!](priv, \sigma T) &\stackrel{\text{def}}{=} \forall pred \in S.\ [\![pred]\!](priv, \sigma T) \\ [\![(S, T, M)]\!](\alpha, \sigma_{src}, \sigma_{tgt}) &\stackrel{\text{def}}{=} \exists \sigma T_{src}, \sigma T_{tgt}. [\![M]\!](\alpha, \sigma T_{src}, \sigma T_{tgt}) \land \\ &\quad \sigma T_{src}.0 = \sigma_{src} \land [\![\sigma]\!](priv_{src}(\alpha), \sigma T_{src}) \land \\ &\quad \sigma T_{tgt}.0 = \sigma_{tgt} \land [\![T]\!](priv_{tgt}(\alpha), \sigma T_{tgt}) \end{aligned}$$

Registers, constants, static values have obvious semantics that maps a state to a dynamic value. Tagged registers and values map an extended state to a dynamic value. The semantics of *pred* is a predicate over a set of blocks, which represents the set of private blocks, and an extended state (see /coq/proof/InvState.v:346). In particular, $[\![\text{Uniq}(r)]\!](priv, \sigma T)$ means r is not alised with any other values in σT (see /coq/proof/InvState.v:314), and $[\![\text{Priv}(rT)]\!](priv, \sigma T)$ means if rT represents a pointer, it is not in $priv$ (see /coq/proof/InvState.v:332). $[\![vT_1 \perp vT_2]\!]$ means if two values are pointers, they points to different memory block (b for block, o for offset, see /coq/proof/InvState.v:244). The semantics of a maydiff set is that all corresponding values are injected except for those in the maydiff set. The semantics of an assertion is that there exist ghost and old register files of the source and target such that, together with the source and target states, satisfy the semantics of the source and target assertions and the maydiff set (see /coq/proof/InvState.v:430). Here, $priv_{src}(\alpha)$ is the set of those source blocks that are not mapped to a target block, and $priv_{tgt}(\alpha)$ is the set of those target blocks that are not mapped from a source block.

[9] We submitted the Coq proof scripts as supplementary material.

Algorithm 4 CheckEquivBeh(P, I_{src}, I_{tgt}: Command)

match I_{src}, I_{tgt} **with**
// call
| $(x_{src} := \texttt{call}\ f_{src}\ args_{src})$, $(x_{tgt} := \texttt{call}\ f_{tgt}\ args_{tgt}) \Rightarrow$ **return** $(f_{src} \sim_P f_{tgt}) \wedge (args_{src} \sim_P args_{tgt})$
| $(\_ := \texttt{call}\ \_\ \_), \_ \mid \_, (\_ := \texttt{call}\ \_\ \_) \Rightarrow$ **return** `false`
// alloca
| $(p_{src} := \texttt{alloca}\ x_{src})$, $(p_{tgt} := \texttt{alloca}\ x_{tgt}) \Rightarrow$ **return** $x_{src} \sim_P x_{tgt}$
| $(p_{src} := \texttt{alloca}\ x_{src})$, $(\texttt{lnop}) \Rightarrow$ **return** `true`
| $(\_ := \texttt{alloca}\ \_), \_ \mid \_, (\_ := \texttt{alloca}\ \_) \Rightarrow$ **return** `false`
// store
| $(\texttt{store}\ p_{src}\ v_{src})$, $(\texttt{store}\ p_{tgt}\ v_{tgt}) \Rightarrow$ **return** $(p_{src} \sim_P p_{tgt}) \wedge (v_{src} \sim_P v_{tgt})$
| $(\texttt{store}\ p_{src}\ v_{src})$, $(\texttt{lnop}) \Rightarrow$ **return** $\text{Priv}(p_{src}) \in P$
| $(\texttt{store}\ \_), \_ \mid \_, (\texttt{store}\ \_) \Rightarrow$ **return** `false`
// target is load
| $(v_{src} := \texttt{load}\ p_{src})$, $(v_{tgt} := \texttt{load}\ p_{tgt}) \Rightarrow$ **return** $p_{src} \sim_P p_{tgt}$
| $\_$, $(v_{tgt} := \texttt{load}\ p_{tgt}) \Rightarrow$ **return** `false`
// target is div
| $(\_ := \texttt{div}\ \_\ b_{src})$, $(\_ := \texttt{div}\ \_\ b_{tgt}) \Rightarrow$ **return** $b_{src} \sim_P b_{tgt}$
| $\_$, $(\_ := \texttt{div}\ \_\ b_{tgt}) \Rightarrow$ **return** $\text{IsNonzero}(P, b_{tgt})$
// misc.
| $\_$, $\_ \Rightarrow$ **return** `true`
end match

H Details of ERHL Proof Checker

H.1 Semantic Interpretation of Hoare triples for call instructions

We give the semantic interpretation of the Hoare triple for call instructions $I_{src} = (x_{src} := \texttt{call}\ f_{src}\ args_{src})$ and $I_{tgt} = (x_{tgt} := \texttt{call}\ f_{tgt}\ args_{tgt})$ as follows:

$$
\begin{aligned}
&[\![\{P\}\ I_{src} \sim I_{tgt}\ \{Q\}]\!] \stackrel{\text{def}}{=} \\
&\quad \forall \sigma_{src}, \sigma_{tgt}, \alpha.\ [\![P]\!]_\alpha(\sigma_{src}, \sigma_{tgt}) \wedge Instr(\sigma_{src}) = I_{src} \wedge Instr(\sigma_{tgt}) = I_{tgt} \implies \\
&\qquad\qquad (f_{src} \sim_\alpha f_{tgt}) \ \wedge\ (args_{src} \sim_\alpha args_{tgt}) \wedge \\
&\forall v_{src}, v_{tgt}, \sigma'_{src}, \sigma'_{tgt}, \alpha' \sqsupseteq \alpha.\ [\![\top]\!]_{\alpha'}(\sigma'_{src}, \sigma'_{tgt}) \wedge v_{src} \sim_{\alpha'} v_{tgt} \wedge \sigma_{src} \overset{v_{src}}{\Rightarrow} \sigma'_{src} \wedge \sigma_{tgt} \overset{v_{tgt}}{\Rightarrow} \sigma'_{tgt} \implies \\
&\qquad\qquad \exists \alpha'' \sqsupseteq \alpha'.\ [\![Q]\!]_{\alpha''}(\sigma'_{src}, \sigma'_{tgt})\ .
\end{aligned}
$$

where, $v_{src} \sim_\alpha v_{tgt}$ means v_{src} is injected into v_{tgt} via memory injection α, $\top$ is the assertion with no predicates and the full maydiff set (*i.e.,* the values of every register may differ), and $\sigma \overset{v}{\Rightarrow} \sigma'$ means σ is about to call a function, and σ' is a possible return state after the function call for the case that the callee returns v.

The semantic interpretation means that the source and the target is about to call equivalent functions with equivalent arguments, and for all future extension α' of the current memory injection α, if the return values and return states are related by α', then there exists a future extension α'' of α' for which $[\![Q]\!]$ is satisfied for the return states.

In order to prove a Hoare triple, we need to guarantee that the source and target call equivalent functions with equivalent arguments, and in turn, we can rely on the fact that the callees return equivalent states. Using this semantics interpretation of calls, we proved semantics preservation of programs using the basic approach of parametric bisimulation [21]. See /coq/proof/SimulationLocal.v:164 and /coq/proof/AdequacyLocal.v:243 for more details.

H.2 Post-Assertion Computation for Commands

CheckEquivBeh Algorithm 4 is the CheckEquivBeh() algorithm for commands (see /coq/def/Postcond.v:769). Here, $x_{src} \sim_P y_{tgt}$ means one of the followings holds: (i) $x = y$ and x is not in the maydiff set; (ii) $(x_{src} \sqsupseteq y_{src}) \in P_{src}$ and y is not in the maydiff set; or (iii) x is not in the maydiff set and $(x_{tgt} \sqsupseteq y_{tgt}) \in P_{tgt}$. Basically, this implies $x_{src} \sim_\alpha y_{tgt}$ holds for all α that is compatible with P. Also, $e_{src} \sim_P e'_{tgt}$ means e and e' are of the same expression kind, *e.g.,* they are both `add`, and for all matching operands x, x', $x_{src} \sim_P x'_{tgt}$ holds. IsNonzero() performs an analysis for proving that the value is nonzero. For simplicity, we omit the details.

In essence, if the target instruction may emit an event or invoke undefined behavior, the source instruction should be similar to that. Note that we allow source load instruction with target `lnop` instruction, which indeed occurs in the validation of `mem2reg`.

Algorithm 5 CalcPostAssnCmd(P: Assn, I_{src}, I_{tgt}: Command): Assn

1: P' := Prune(P, I_{src}, I_{tgt})
2: $(P''_{src}, P''_{tgt}, M'')$:= AddMemoryPreds(I_{src}, I_{tgt}, P')
3: P'''_{src} := AddLessdefPreds(I_{src}, P''_{src})
4: P'''_{tgt} := AddLessdefPreds(I_{tgt}, P''_{tgt})
5: **return** ReduceMaydiff($P'''_{src}, P'''_{tgt}, M''$)

Algorithm 6 Prune(P: Assn, I_{src}, I_{tgt}: Command): Assn

1: $(P_{src}, P_{tgt}, M) := P$
2: P'_{src} := PruneU(P_{src}, I_{src})
3: P'_{tgt} := PruneU(P_{tgt}, I_{tgt})
4: $M' := M \cup \{\text{Def}(I_{src}), \text{Def}(I_{tgt})\}$
5: **return** (P'_{src}, P'_{tgt}, M')

CalcPostAssn Algorithm 5 is the post-assertion computation algorithm for commands I_{src}, I_{tgt} (see /coq/def/Postcond.v:1058 for definition and /coq/proof/SoundPostcondCmd.v:309, /coq/proof/SoundPostcondCall.v:254 for the soundness of (PostAssn) for commands). At line 1, it removes predicates that no longer hold after executing the commands (Prune). Then at line 2, it adds Uniq and Priv predicates when necessary (AddMemoryPreds). At lines 3-4, it adds lessdef predicates to unary assertions for both source and target (AddLessdefPreds), and at line 5, finally tries to remove registers from the maydiff set.

Prune Algorithm 6 is the assertion pruning algorithm for commands I_{src}, I_{tgt} (see /coq/def/Postcond.v:355, /coq/def/Postcond.v:386). At lines 2-3, it removes predicates from unary assertions for both source and target (PruneU). Then at line 4, the left-hand sides of I_{src}, I_{tgt} are added to the maydiff set.

PruneU PruneU(S, I) removes the following memory predicates from a unary assertion:

- If I defines a register r, remove all predicates on r (see /coq/def/Postcond.v:373).
- If I is a store instruction $*p := v$, remove all lessdef equations on $*q$ for which we cannot prove $p \perp q$. We can prove $p \perp q$ if we have it as a predicate, or [$p \neq q$, either p or q is unique, and the other is physical] holds (see /coq/def/Postcond.v:341).
- If I is a call instruction, remove all lessdef equations on $*q$ unless Priv(q) holds (see /coq/def/Postcond.v:409).
- Remove Uniq(p) if p leaked, *i.e.,* copied to another register, used as the instruction's operand (see /coq/def/Postcond.v:347), or a function is called unless Priv(p) holds (see /coq/def/Postcond.v:411).

AddMemoryPreds AddMemoryPreds(I_{src}, I_{tgt}, P) adds the following memory predicates to P:

- If (I_{src}, I_{tgt}) are allocations (p_{src} := `alloca`$(\cdots)$), (p_{tgt} := `alloca`$(\cdots)$), add Uniq(p_{src}) to the source assertion. If $p_{src} = p_{tgt}$ then remove p_{src} from the maydiff set (see /coq/def/Postcond.v:893).
- If I_{src} is an allocation (p_{src} := `alloca`$(\cdots)$) and I_{tgt} is `lnop`, then add Uniq(p_{src}), Priv(p_{src}) to the source assertion (see /coq/def/Postcond.v:905).
- If (I_{src}, I_{tgt}) are call instructions (x_{src} := `call`$(\cdots)$), (x_{tgt} := `call`$(\cdots)$) and $x_{src} = x_{tgt}$, then remove x_{src} from the maydiff set (see /coq/def/Postcond.v:927).

AddLessdefPreds AddLessdefPreds(I, S) adds the following lessdef predicates to S:

- If I is a side-effect-free operation (x := op *args*), add ($x \sqsupseteq$ op *args*) and (op *args* $\sqsupseteq x$) (see /coq/def/Postcond.v:936). Note that load is regarded as side-effect-free.
- If I is a store instruction $*p := v$, add $*p \sqsupseteq v$ (see /coq/def/Postcond.v:939).
- If I is an allocation instruction $*p$:= `alloca`$(\cdots)$, add $*p \sqsupseteq$ `undef` (see /coq/def/Postcond.v:942).

H.3 Post-Assertion Computation for Phinodes

Phinodes do not emit any events so that any pair of phinodes behaves equivalently. Thus, for phinodes, CheckEquivBeh() checks nothing. As described in §4, CalcPostAssn() works for phinodes as follows (see /coq/def/Postcond.v:689 for the definition and /coq/proof/SoundPostcondPhinodes.v:960 for the soundness of (PostAssn) for phinodes):

- Remove old registers, and copy predicates on physical registers into those about old ones (see /coq/def/Postcond.v:673).
- Remove predicates on those registers defined in the phinodes (see /coq/def/Postcond.v:674).
- Add predicates $x \sqsupseteq \bar{y}, \bar{y} \sqsupseteq x$ for each assignment $x := y$ that is performed in the phinodes (see /coq/def/Postcond.v:680).
- Tries to reduce the maydiff set (see /coq/def/Postcond.v:686).

Misc. CheckCFG is implemented in /coq/def/Validator.v:76 and /coq/def/Validator.v:129.
CheckInit is implemented in /coq/def/Validator.v:247, and its specification is proved in /coq/proof/SimulationValid.v:1314.
CheckIncl is implemented in /coq/def/Hints.v:187, and the soundness of (Incl) is proved in /coq/proof/SoundImplies.v:279.

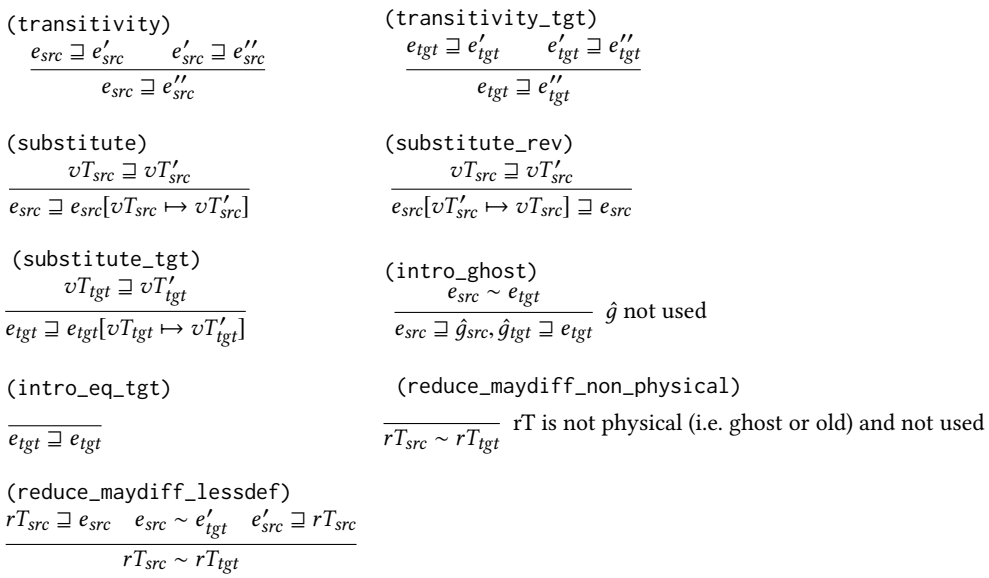

Figure 16. Formally Verified Inference Rules

I Non-Arithmetic Inference Rules

In order to support `mem2reg`, `gvn`, and `licm`, we use 9 non-arithmetic inference rules. In the Coq development, we define the 9 rules in coq/def/Infrules.v:392, and formally verified their soundness in coq/proof/SoundInfrules.v:52.

Each of the rules is based on one of the proof rules in Fig. 16. Here, $rT_{src} \sim rT'_{tgt}$ in premises means $rT_{src} \sim_P rT'_{tgt}$ for the current assertion P, Also, $e_{src} \sim e'_{tgt}$ means e and e' are of the same expression kind, *e.g.*, they are both add, and for all matching operands rT, rT', $rT_{src} \sim_P rT'_{tgt}$ holds for the current assertion P. $rT_{src} \sim rT_{tgt}$ in conclusions means you can remove rT from the maydiff set.