

# A Generic Framework for Verified Compilers Using Isabelle/HOL’s Locales

Martin Desharnais and Stefan Brunthaler

National Cyber Defence Research Institut (CODE), UniBw M, Germany  
{martin.desharnais, brunthaler}@unibw.de

## Abstract

In this paper, we present a prototype version of a generic framework for formalizing compiler transformations. Our framework leverages Isabelle/HOL’s *locales*—a module system for generic formalizations—to abstract over concrete languages and transformations. The framework thus enables us to state common definitions for language semantics, program behaviours, forward and backward simulations, and compilers. We provide generic operations, such as compiler composition, and prove general theorems, resulting in reusable proof components. By demonstrating our idea on a concrete example, we provide evidence of how locales allow reuse and, therefore, enable encapsulation of verification artefacts into modules.

## 1 Introduction

The mechanically verified formalization of software components has been the subject of much research in the last decades. Especially influential were the CompCert [3] and CakeML [2] projects, which produced realistic compilers from a (large subset of) two real-world programming languages (C99 and Standard ML) to real hardware platforms. These compilers showed both that mechanized verification is feasible and that it has a measurable effect on the dependability of the compiler [6].

We can now observe a shift in perspective, where the idea of mechanically verified software components is becoming a concrete and desirable goal. Formalization projects are increasing in number, but also in size, complexity, and lifetime. There is an analogy to be made with the emergence, in the second half of the 20th century, of software engineering to the point that the term *proof engineering* starts to be used. New and interesting questions now emerge. How to avoid repetition in definitions and proofs? Which concepts can be generalized and reused? How to separate a formalization in independent components, so that multiple people can work in parallel? What should be the interface between such components? How can tooling make proof engineers more productive? What is a good balance between proof readability and the time required to (mechanically) verify it? etc.

In the case of compiler verification, we have a very well-understood domain, with well-known terminology, that builds on decades of research and empirical experience. But as is the case for a lot of small software prototypes, small-scale formalizations constantly redefined similar abstractions and concepts. This is something we wanted to avoid when we started a small formalization, in Isabelle/HOL [4], of three small stack-based languages implementing different optimizations. Inspired by the concept of modularization in software engineering, we separated the general concepts from the language-dependent parts. We learned about, and made use of, Isabelle’s locales to devise a small generic framework<sup>1</sup> for the verification of program transformations.

## 2 Background

In this section, we start with brief overview of the operational semantics of programming languages and follow with a short introduction to Isabelle’s locales.

<sup>1</sup>Available at <https://github.com/mdesharnais/framework-VeriComp> (commit 16906564).

## 2.1 Programming Language Semantics

The operational semantics of a programming language can be defined as a transition system representing the execution of a program written in this language. A language  $L = \langle S, I, F, \rightarrow \rangle$  is defined by a set  $S$  of program states, a set  $I \subseteq S$  of initial states, a set  $F \subseteq S$  of final states, and a transition relation  $\rightarrow \subseteq S \times S$ . The execution of a program is modelled as a sequence of states  $s_1 \rightarrow s_2 \rightarrow \dots$  with  $s_1 \in I$ . An execution is called terminating if there exists a state  $s_n$  such that  $s_1 \rightarrow s_2 \rightarrow \dots \rightarrow s_n$  and  $\nexists s_{n+1}. s_n \rightarrow s_{n+1}$ , and non-terminating otherwise. A terminating execution is said to be successful if  $s_n \in F$  and to go wrong otherwise. These execution behaviours are usually called the program's behaviour and written  $s \Downarrow b$ .

The compiler from a language  $L_1$  to  $L_2$  is a partial function  $\mathcal{C} : S_1 \rightarrow S_2$ , which maps a program  $s \in I_1$  to  $\mathcal{C}(s) \in I_2$ .

Two programs  $s$  and  $c$  are said to be equivalent if they exhibit the same behaviour, i.e.  $\forall b. s \Downarrow b \iff c \Downarrow b$ . This can be established using a bisimulation [5]: the conjunction of a backward and a forward simulation. Consider a binary relation  $\approx$ , between program states, expressing that two states are to be considered equivalent for a given use case. This relation is called a simulation whenever  $\forall s s' c, s \approx c \wedge s \rightarrow s' \implies \exists c', s' \approx c' \wedge c \rightarrow c'$ . A backward simulation, thus, shows that every behaviour of the compiled program is also a behaviour of the source program, i.e. the compilation is correct (sound). A forward simulation shows that every behaviour of the source program can be achieved by the compiled program, i.e. the compilation is complete.

## 2.2 Isabelle's locales

Locales are an Isabelle construct to define parametric theories [1]. They are based on the concept of proof contexts. A theorem of the form

$$\bigwedge p_1 p_2 \dots p_n. A_1 \implies A_2 \implies \dots \implies A_m \implies C$$

has a set  $\{p_1, p_2, \dots, p_n\}$  of parameters, a set  $\{A_1, A_2, \dots, A_m\}$  of assumptions, and proves a conclusion  $C$ . Taken together, the sets of parameters and assumptions is called the proof context. Locales enable the user to define a named proof context and reuse it for multiple conclusions, thus avoiding having to repeat its components in every theorem. Consider for example a formalization of monoids.

**locale monoid =**

**fixes**

$f :: 'a \Rightarrow 'a \Rightarrow 'a$  (**infix  $\cdot$** ) **and**

$e :: 'a$

**assumes**

*associativity*:  $x \cdot (y \cdot z) = (x \cdot y) \cdot z$  **and**

*left-identity*:  $e \cdot x = x$  **and**

*right-identity*:  $x \cdot e = x$

**context monoid begin**

**primrec**  $pow :: nat \Rightarrow 'a \Rightarrow 'a$  **where**

$pow\ 0\ x = e$  |

$pow\ (Suc\ n)\ x = x \cdot pow\ n\ x$

**lemma** *pow-add*:

$pow\ (n + m)\ x = pow\ n\ x \cdot pow\ m\ x$

**proof ... qed**

**end**

The *monoid* locale consists of a sequence of parameters, introduced by the **fixes** keyword, and a sequence of assumptions, introduced by the **assumes** keyword. When working in a locale context—introduced by the **context** command or directly following a locale definition—new definitions and theorems can be derived from the locale parameters, its assumptions, and previous derived terms and theorems.

The automatically introduced *locale predicate monoid* ::  $('a \Rightarrow 'a \Rightarrow 'a) \Rightarrow 'a \Rightarrow bool$  identifies locale interpretations, i.e. parameters for which the assumptions hold. We can see that locale contexts really just are syntactic sugar for manual contexts by inspecting the theorem *monoid.pow-add* from outside the locale context.

$$monoid\ f\ e \implies pow\ f\ e\ (n + m)\ x = f\ (pow\ f\ e\ n\ x)\ (pow\ f\ e\ m\ x)$$

Locales can also be extended by more parameters and assumptions.

```
locale monoid-homomorphisms =  $M_f$ : monoid  $f e_f$  +  $M_g$ : monoid  $g e_g$ 
for  $f :: 'a \Rightarrow 'a \Rightarrow 'a$  (infix  $\cdot$ ) and  $e_f$  and  $g :: 'b \Rightarrow 'b \Rightarrow 'b$  (infix  $\diamond$ ) and  $e_g$  +
fixes  $map :: 'a \Rightarrow 'b$ 
assumes
   $map$ -distributive:  $map (x \cdot y) = map x \diamond map y$  and
   $map$ -identity:  $map e_f = e_g$ 
```

The *monoid-homomorphisms* locale extends two instances of *monoid*, fixes a projection function *map* between their underlying types, and states two assumptions on its interaction with the two monoid structures. The extended locale contexts are named, so that elements can be accessed, e.g. by writing *M<sub>f</sub>.left-identity*. The sequence of parameters required by the extended locales are introduced by the **for** keyword.

Finally, locales can be interpreted, with the **interpretation** command, by providing values for the parameters and proving that the assumptions hold.

```
interpretation monoid-nat-addition: monoid (+) (0 :: nat)
proof — Proof that the assumptions hold qed
```

Following interpretation, all derived definitions and theorems, specialized for the provided arguments, are available in the *monoid\_nat\_addition* namespace.

### 3 The Design of the Framework

The framework has three main components: some abstract definitions of languages and compilers using locales, a generic definition of program behaviour, and some composition operations over simulations and compilers.

#### 3.1 Locales

The definition of programming languages is separated into two parts: an abstract semantics and a concrete program representation.

```
locale semantics =
fixes  $step :: 'state \Rightarrow 'state \Rightarrow bool$  and  $final :: 'state \Rightarrow bool$ 
assumes  $final$ -finished:  $final s \implies finished\ step\ s$ 
```

```
locale language = semantics step final
for  $step :: 'state \Rightarrow 'state \Rightarrow bool$  and  $final :: 'state \Rightarrow bool$  +
fixes  $load :: 'prog \Rightarrow 'state$ 
```

The *semantics* locale represents the semantics as an abstract machine. It is expressed by a transition system with a transition relation *step*—usually written as an infix ( $\rightarrow$ ) arrow—and final states *final*. The *language* locale represents the concrete program representation (type variable *'prog*), which can be transformed into a program state (type variable *'state*) by the *load* function. The set of initial states of the transition system is implicitly defined by the codomain of *load*.

```
locale backward-simulation = L1: semantics step1 final1 + L2: semantics step2 final2 + well-founded ( $\sqsubset$ )
for
   $step1 :: 'state1 \Rightarrow 'state1 \Rightarrow bool$  and  $step2 :: 'state2 \Rightarrow 'state2 \Rightarrow bool$  and
   $final1 :: 'state1 \Rightarrow bool$  and  $final2 :: 'state2 \Rightarrow bool$  and
   $order :: 'index \Rightarrow 'index \Rightarrow bool$  (infix  $\sqsubset$ ) +
fixes  $match :: 'index \Rightarrow 'state1 \Rightarrow 'state2 \Rightarrow bool$ 
```

**assumes**

*match-final*:  $match\ i\ s_1\ s_2 \implies final2\ s_2 \implies final1\ s_1$  **and**  
*simulation*:  $match\ i\ s_1\ s_2 \implies s_2 \rightarrow_2 s_2' \implies$   
 $(\exists i'. s_1 \rightarrow_1^+ s_1' \wedge match\ i'\ s_1'\ s_2') \vee (\exists i'. match\ i'\ s_1\ s_2' \wedge i' \sqsubset i)$

A simulation is defined between two semantics  $L1$  and  $L2$ . A *match* predicate expresses that two states from  $L1$  and  $L2$  are equivalent. The *match* predicate is also parameterized with an ordering used to avoid stuttering.

The only two assumptions of a backward simulation are that a final state in  $L2$  will also be a final in  $L1$ , and that a step in  $L2$  will either represent a non-empty sequence of steps in  $L1$ —the  $(\rightarrow_1^+)$  relation is the transitive closure of the  $(\rightarrow_1)$  relation—or will result in an equivalent state. Stuttering is ruled out by the requirement that the index on the *match* predicate decreases with respect to the well-founded  $(\sqsubset)$  ordering.

**locale** *compiler* =

*L1*: language *step1* *final1* *load1* + *L2*: language *step2* *final2* *load2* +  
*backward-simulation* *step1* *step2* *final1* *final2* *order* *match*  
**for**  
*step1* ::  $'state1 \Rightarrow 'state1 \Rightarrow bool$  **and** *step2* ::  $'state2 \Rightarrow 'state2 \Rightarrow bool$  **and**  
*final1* ::  $'state1 \Rightarrow bool$  **and** *final2* ::  $'state2 \Rightarrow bool$  **and**  
*load1* ::  $'prog1 \Rightarrow 'state1$  **and** *load2* ::  $'prog2 \Rightarrow 'state2$  **and**  
*order* ::  $'index \Rightarrow 'index \Rightarrow bool$  **and**  
*match* ::  $'index \Rightarrow 'state1 \Rightarrow 'state2 \Rightarrow bool$  +  
**fixes** *compile* ::  $'prog1 \Rightarrow 'prog2\ option$   
**assumes** *compile-load*:  $compile\ p_1 = Some\ p_2 \implies \exists i. match\ i\ (load1\ p_1)\ (load2\ p_2)$

The *compiler* locale relates two languages,  $L1$  and  $L2$ , by a backward simulation and provides a *compile* partial function from a concrete program in  $L1$  to a concrete program in  $L2$ . The only assumption is that a successful compilation results in a program which, when loaded, is equivalent to the loaded initial program.

### 3.2 Behaviours

We define a generic datatype to encode three broad execution behaviours: successful termination (*Terminates*), non-terminating execution (*Diverges*), and going wrong (*Goes-wrong*).

**datatype**  $'state\ behaviour = Terminates\ 'state \mid Diverges \mid Goes-wrong\ 'state$

Terminating behaviours are annotated with the last state of the execution in order to compare the result of two executions with the *rel-behaviour* ::  $('a \Rightarrow 'b \Rightarrow bool) \Rightarrow 'a\ behaviour \Rightarrow 'b\ behaviour \Rightarrow bool$  relation.

$$f\ s_1\ s_2 \implies rel-behaviour\ f\ (Terminates\ s_1)\ (Terminates\ s_2)$$

$$rel-behaviour\ f\ Diverges\ Diverges$$

$$f\ s_1\ s_2 \implies rel-behaviour\ f\ (Goes-wrong\ s_1)\ (Goes-wrong\ s_2)$$

The exact meaning of the three behaviours is defined in the *semantics* locale, where a  $(\Downarrow) :: 'state \Rightarrow 'state\ behaviour \Rightarrow bool$  relation is defined to assign an execution behaviour to a program state. The  $(\rightarrow^*)$  relation is the reflexive transitive closure of the  $(\rightarrow)$  relation and  $(\rightarrow^\infty)$  is its coinductive, infinitely transitive closure. The predicate *finished* ::  $('a \Rightarrow 'a \Rightarrow bool) \Rightarrow 'a \Rightarrow bool$  identifies a state that cannot make a transition.

$$\frac{s_1 \rightarrow^* s_2 \quad \text{finished } (\rightarrow) s_2 \quad \text{final } s_2}{s_1 \Downarrow \text{Terminates } s_2} \text{state-terminates} \quad \frac{s_1 \rightarrow^\infty}{s_1 \Downarrow \text{Diverges}} \text{state-diverges}$$

$$\frac{s_1 \rightarrow^* s_2 \quad \text{finished } (\rightarrow) s_2 \quad \neg \text{final } s_2}{s_1 \Downarrow \text{Goes-wrong } s_2} \text{state-goes-wrong}$$

Even though the  $(\rightarrow)$  transition relation in the *semantics* locale need not be deterministic, if it happens to be, then the behaviour of a program becomes deterministic too.

$$\frac{\bigwedge x y z. \frac{x \rightarrow y \quad x \rightarrow z}{y = z} \quad s \Downarrow b_1 \quad s \Downarrow b_2}{b_1 = b_2}$$

The main correctness theorem states that, for any two matching programs, any not wrong behaviour of the later is also a behaviour of the former. In other words, if the compiled program does not crash, then its behaviour, whether it terminates or not, is also a valid behaviour of the source program. The predicate *is-wrong* :: '*state behaviour*  $\Rightarrow$  *bool*' identifies wrong behaviours.

$$\frac{\text{match } i \ s_1 \ s_2 \quad s_2 \Downarrow_2 \ b_2 \quad \neg \text{is-wrong } b_2}{\exists b_1 \ i'. \ s_1 \Downarrow_1 \ b_1 \wedge \text{rel-behaviour } (\text{match } i') \ b_1 \ b_2}$$

Because this theorem is proven in the context of the *backward-simulation* and, thus, only depends on its parameters and assumptions, it is independent of the concrete programming language, and need only be to be proven once. It automatically holds for all interpretations of *backward-simulation*.

As a corollary, the preservation of behaviour can be lifted to the compilation of concrete program representation.

$$\frac{\text{compile } p_1 = \text{Some } p_2 \quad \text{load2 } p_2 \Downarrow_2 \ b_2 \quad \neg \text{is-wrong } b_2}{\exists b_1 \ i. \ \text{load1 } p_1 \Downarrow_1 \ b_1 \wedge \text{rel-behaviour } (\text{match } i) \ b_1 \ b_2}$$

### 3.3 Generic Composition of Simulations and Compilers

We define the generic composition of matching functions,  $\diamond :: ('a \Rightarrow 'b \Rightarrow 'c \Rightarrow \text{bool}) \Rightarrow ('d \Rightarrow 'c \Rightarrow 'e \Rightarrow \text{bool}) \Rightarrow 'a \times 'd \Rightarrow 'b \Rightarrow 'e \Rightarrow \text{bool}$ , and orderings, *lex-prod* ::  $('a \Rightarrow 'a \Rightarrow \text{bool}) \Rightarrow ('b \Rightarrow 'b \Rightarrow \text{bool}) \Rightarrow 'a \times 'b \Rightarrow 'a \times 'b \Rightarrow \text{bool}$ , such that the composition of two backward simulations is itself a backward simulation.

$$\frac{\text{backward-simulation } (\rightarrow_1) (\rightarrow_2) \text{final}_1 \text{final}_2 (\sqsubset_1) (\approx_1) \quad \text{backward-simulation } (\rightarrow_2) (\rightarrow_3) \text{final}_2 \text{final}_3 (\sqsubset_2) (\approx_2)}{\text{backward-simulation } (\rightarrow_1) (\rightarrow_3) \text{final}_1 \text{final}_3 (\text{lex-prod } (\sqsubset_1^+) (\sqsubset_2)) ((\approx_1) \diamond (\approx_2))}$$

We define the generic  $(\Leftarrow) :: ('a \Rightarrow 'b \text{ option}) \Rightarrow ('c \Rightarrow 'a \text{ option}) \Rightarrow 'c \Rightarrow 'b \text{ option}$  composition operator on compilers, which corresponds to the monadic bind of the *option* type found in a compiler's codomain.

$$(\mathcal{C}_2 \Leftarrow \mathcal{C}_1) p \equiv \text{Option.bind } (\mathcal{C}_1 \ p) \ \mathcal{C}_2$$

Its correctness can then be generically proven for any two interpretations of the *compiler* locale.

$$\frac{\text{compiler } (\rightarrow_1) (\rightarrow_2) \text{final}_1 \text{final}_2 \text{load}_1 \text{load}_2 (\sqsubset_1) (\approx_1) \ \mathcal{C}_1 \quad \text{compiler } (\rightarrow_2) (\rightarrow_3) \text{final}_2 \text{final}_3 \text{load}_2 \text{load}_3 (\sqsubset_2) (\approx_2) \ \mathcal{C}_2}{\text{compiler } (\rightarrow_1) (\rightarrow_3) \text{final}_1 \text{final}_3 \text{load}_1 \text{load}_3 (\text{lex-prod } (\sqsubset_1^+) (\sqsubset_2)) ((\approx_1) \diamond (\approx_2)) (\mathcal{C}_2 \Leftarrow \mathcal{C}_1)}$$

## 4 An Instantiation of the Framework

The first programming languages for which we instantiated the framework are three interpreted, stack-based languages. The first one, *Std*, is a standard assembly language with push/pop and load/store instructions, conditional jumps, and unary/binary operations. The second language, *Inca* expands *Std* with inline caching, i.e. operations that are faster for specific operand types but fallback to a generic operation otherwise. The third language, *Ubx*, goes one step further by introducing operations that operate on unboxed operands.

Because we did not want to fix a concrete representation for identifiers or the dynamic environment, we define each language in a locale that abstracts over the relevant types and operations. The *Std* language is defined as follows.

```

datatype 'id instr = ...
datatype ('id, 'env) state = ...
definition final :: ('id, 'env) state  $\Rightarrow$  bool where ...
locale std =
  fixes env-empty :: 'env and env-get :: 'env  $\Rightarrow$  'id  $\Rightarrow$  'value option and env-add :: 'env  $\Rightarrow$  'id  $\Rightarrow$  'value  $\Rightarrow$  'env
begin
  inductive step :: ('id, 'env) state  $\Rightarrow$  ('id, 'env) state  $\Rightarrow$  bool where ...
  definition load :: 'id instr list  $\Rightarrow$  ('id, 'env) state where ...
  lemma final-finished: final s  $\Longrightarrow$  finished step s ...
  sublocale std-sem: semantics step final ...
  sublocale std-lang: language step final load ...
end

```

Because locales do not support the definition of new types, the *instr* and *state* datatypes had to be defined in the top-level theory. Moreover, they needed to abstract over the *'id* and *'env* types, which are fixed only inside the locale. The function *final* does not depend on any locale parameter, so it could have been defined at either place. The *step* relation and *load* function, in contrast, depend on the functions to manage the environment and thus need to be defined inside the locale. They instantiate the *instr* and *state* types using the locale's fixed *'id* and *'env*. The lemma *final-finished* can then be stated and proven. Finally, the *semantics* and *language* locales can be interpreted<sup>2</sup>, thereby proving that *Std* corresponds to our abstraction of language with a semantics. This also specializes all general results of said locales to this concrete language definition. The two other languages, *Inca* and *Ubx*, follow the same structure.

```

locale std-inca-simulation = std: std env-empty env-get env-add + inca: inca env-empty env-get env-add
for env-empty :: 'env and env-get :: 'env  $\Rightarrow$  'id  $\Rightarrow$  'val option and env-add :: 'env  $\Rightarrow$  'id  $\Rightarrow$  'val  $\Rightarrow$  'env
begin
  definition match :: nat  $\Rightarrow$  ('id, 'env) Std.state  $\Rightarrow$  ('id, 'env) Inca.state  $\Rightarrow$  bool where ...
  lemma match-final: match i s1 s2  $\Longrightarrow$  Inca.final s2  $\Longrightarrow$  Std.final s1 ...
  lemma simulation: match i s1 s2  $\Longrightarrow$  inca.step s2 s2'  $\Longrightarrow$ 
    ( $\exists i' s_1' . plus\ std.step\ s_1\ s_1' \wedge match\ i'\ s_1'\ s_2'$ )  $\vee$  ( $\exists i' . match\ i'\ s_1\ s_2' \wedge i' < i$ ) ...
  sublocale std-ubx-backward-simulation: backward-simulation std.step inca.step Std.final Inca.final (<) match ...
end

```

Proving a backward simulation between *Std* and *Inca* requires to extend the *std* and *inca* locales, define the required *match*, prove the two required *match-final* and *simulation* lemmas, and finally interpret the *backward-simulation* locale.

```

definition compile :: 'id Std.instr list  $\Rightarrow$  'id Inca.instr list option where ...
locale std-inca-compiler = sim: std-inca-simulation env-empty env-get env-add
for env-empty :: 'env and env-get :: 'env  $\Rightarrow$  'id  $\Rightarrow$  'value option and env-add :: 'env  $\Rightarrow$  'id  $\Rightarrow$  'value  $\Rightarrow$  'env

```

<sup>2</sup>The **sublocale** command is a variation of the **interpretation** command.

**begin**

**lemma** *compile-load*:  $compile\ p1 = Some\ p2 \implies sim.match\ i\ (std.load\ p1)\ (inca.load\ p2)\ \dots$

**sublocale** *std-to-inca-compiler*:

*compiler\ std.step\ inca.step\ Std.final\ Inca.final\ std.load\ inca.load\ (<)\ sim.match\ compile\ \dots*

**end**

Defining the compiler and proving its correctness is done in a similar fashion. Because the compilation function depends neither on the dynamic environment nor on the functions defined for the backward simulation, it can be defined in the top-level theory. The *std-inca-compiler* locale merely proves the *compile-load* property and can proceed to interpret the *compiler* locale.

## 5 Discussion

Using locales as a modularization tool for our generic framework turned out to be elegant at times and frustrating in other cases.

### 5.1 Strengths of the Approach

**Parameters, assumptions, and derived elements are clearly separated.** The syntax used to define a locale enables the user to clearly state the parameters and assumptions that are abstracted over. Derived elements such as function definitions and lemmas are clearly separated by being defined later in a locale context. The fact that these extensions can be done at any point following the locale's definition gives a lot of flexibility when structuring the formalization.

**It is possible to abstract over multiple types.** Locales enable fixed variables to depend on multiple type variables. This makes them more general than type classes, with which they have otherwise a lot in common. While traditional type classes permit to abstract over operations on a given abstract type, locales permit to abstract over both operations on concrete types and multiple abstract types. In fact, type classes in Isabelle/HOL are just syntactic sugar for locales with a single type variable.

**It is possible to have multiple interpretations for a given set of type.** Because a locale interpretation introduces a new namespace when specializing the derived elements, multiple instantiations are possible for a given set of types. A classical example for such a situation is a partial order over the integers. Using traditional type classes, one has to decide a canonical order that will be associated with the integer type. In order to use an alternate order, one has to define a bijection to an alternative type which instantiate the type class accordingly. As many distinct types and bijections are required as distinct instantiations are wished.

### 5.2 Weaknesses of the Approach

**Parametric types and type aliases cannot be defined in locales.** This limitation requires the user to generalize data types to abstract over any type variable fixed in the locale definition and define them outside of the locale. This was e.g. the case for the *instr* and *state* data types of the *Std*, *Inca*, and *Ubx* languages. This generalization is trivial, since a fixed type variable in a locale is akin to a type variable in a data type definition. The burden shows up when referring to the generic type in a type annotation, where it must be explicitly instantiated. Because parametric type aliases are also not supported, the instantiation has to be repeated over and over. As the number of type variables increases, type annotations become complex and hard to read.

**When extending existing locales, type annotations on fixed variables are required to name type variables.** These variables appear in the **for** section and their types are inferred from their usage in instantiating the extended locales to be extended. Type inference even succeeds in cases where some type variables must be unified between multiple locale instantiations, as is the case in the *compiler* locale. The user must nevertheless provide some type annotations in order to provide the names with which said type variables can be referred to later.

**Proving lemmas involving locale predicates have high syntactic overhead.** An example of such situation is the *compiler\_composition* lemma, where both the two hypotheses and the conclusion are locale predicates of the *compiler* locale. Proving this lemma involves accessing the *language* instance predicates accessible with expressions such as *assms(1)[THEN compiler.axioms(1)]*. The problem with this syntax is twofold: it depends on the order in which the axioms were stated and it does not scale well when the user needs to extract multiple axioms from multiple assumptions. The first problem could be solved by automatically adding lemmas using the name of the extension, e.g. *compiler.L1* would be a synonym of *compiler.axioms(1)* to refer to the first *language* instance predicate of the compiler's definition. The second could be alleviated if unnamed contexts supported extending locales.

**Unnamed contexts cannot extend existing locales.** They allow to fix variables and state assumptions that are automatically propagated to definitions and lemmas proven in their scope. This is a subset of the possibilities when defining a locale. A next step could be to also allow locale extensions. By propagating the requirement using each locale's instance predication, this would alleviate the syntactical burden of proving lemmas such as *compiler\_composition* by allowing a syntax such as *C1.L1*.

**References to a locale's fixed variables and derived definitions are syntactically different.** When extending locales, as is the case in the *backward\_simulation* locale, derived definitions of the two languages are accessible under names such as *L1.behaves* and *L2.behaves*. Fixed parameters, in contrast, are only accessible with their given name *step1* and *step2*. This lack of uniformity means that refactoring a locale, by replacing a fixed parameter by an equivalent derived definition, might require the user to not just adapt all interpretations, but also all derived definitions and theorems to use the prefixed name. Although this would not necessary be a problem in the absence of name clashes, it can be argued that a uniform naming scheme that permits the systematical use the interpretation's prefix would be preferable.

## 6 Conclusion

We presented the first version of a generic framework for formalizing compilers in Isabelle/HOL. It is based on locales to abstract over the concrete languages and program transformations, provides general definitions for program behaviours and compiler compositions, and generically proves preservation of behaviour. This framework emerged as a side product of our formalization of three stack-based languages that implementing different optimizations. It helped us to emphasize the commonalities between the different theories and to reduce some duplication. Possible future work includes extending the semantics to support traces, exploring how to extract executable programs from such formalizations, and exploring how to simplify or automate repetitive operations such as the composition of multiple compilers.

## 7 Acknowledgement

This paper is part of the project CONCORDIA, a project that has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 830927.

## References

- [1] C. Ballarín. Locales: A Module System for Mathematical Theories. *Journal of Automated Reasoning*, 52(2):123–153, 2014.
- [2] R. Kumar, M. O. Myreen, M. Norrish, and S. Owens. CakeML: A verified implementation of ML. In *Principles of Programming Languages (POPL)*, pages 179–191. ACM Press, Jan. 2014.
- [3] X. Leroy. A formally verified compiler back-end. *Journal of Automated Reasoning*, 43(4):363–446, 2009.
- [4] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL: a proof assistant for higher-order logic*, volume 2283. Springer Science & Business Media, 2002.
- [5] D. Sangiorgi. *Introduction to bisimulation and coinduction*. Cambridge University Press, 2011.
- [6] X. Yang, Y. Chen, E. Eide, and J. Regehr. Finding and understanding bugs in C compilers. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, 2011.