

# Relational Equivalence Proofs Between Imperative and MapReduce Algorithms

Bernhard Beckert, Timo Bingmann, Moritz Kiefer, Peter Sanders,  
Mattias Ulbrich, and Alexander Weigl

Institute of Theoretical Informatics  
Karlsruhe Institute of Technology, Germany

**Abstract.** *MapReduce* frameworks are widely used for the implementation of distributed algorithms. However, translating imperative algorithms into these frameworks requires significant structural changes to the algorithm. As the costs of running faulty algorithms at scale can be severe, it is highly desirable to verify the correctness of the translation, i.e., to prove that the *MapReduce* version is equivalent to the imperative original.

We present a novel approach for proving equivalence between imperative and *MapReduce* algorithms based on partitioning the equivalence proof into a sequence of equivalence proofs between intermediate programs with smaller differences. Our approach is based on the insight that two kinds of sub-proofs are required: (1) uniform transformations changing the control-flow structure that are mostly independent of the particular context in which they are applied; and (2) context-dependent transformations that are not uniform but that preserve the overall structure and can be proved correct using coupling invariants.

We demonstrate the feasibility of our approach by evaluating it on two prototypical algorithms commonly used as examples in *MapReduce* frameworks: *k*-means and PageRank. To carry out the proofs, we use the interactive theorem prover *Coq* with partial proof automation. The results show that our approach and its prototypical implementation based on *Coq* enables equivalence proofs of non-trivial algorithms and could be automated to a large degree.

## 1 Introduction

**Motivation.** Frameworks such as *MapReduce* [8], Spark [22] and Thrill [2] address the challenges arising in the implementation of distributed algorithms by providing a limited set of operations whose execution is automatically parallelized and distributed among the nodes in a cluster. However, translating an existing imperative algorithm into such a framework is a challenge in itself and the original algorithmic structure is often lost during that translation since imperative constructs do not translate directly to the provided primitives. Implementing efficient algorithms using *MapReduce* frameworks can thus require significant changes to the original algorithm.

By proving the equivalence of the original imperative algorithm and its *MapReduce* version, one can verify that no bugs have been introduced during the translation. While such proofs do not directly provide correctness guarantees for the *MapReduce* algorithm, they transfer correctness results from the imperative version to the *MapReduce* implementation. The transferred correctness properties can be formal proofs whose reach then extends to the distributed implementation, but can also be informal arguments, e.g., if the algorithm is a well-known, simple textbook reference implementation or if it has been successfully applied previously.

In this paper, we use the term “*MapReduce*” in a broader sense than implied by the two functions “map” and “reduce”. While some frameworks such as Hadoop’s MapReduce [21] module are programmed strictly by specifying these two functions, the more popular and widely used distributed frameworks provide many additional primitives for performance reasons and to make them easier to program with. Theoretically, these additional primitives can be reduced to only map and reduce operations [4], but this overly complicates the program description and is generally not used in real-world applications. In Appendix A.1, we list the few additional primitives we consider part of the extended *MapReduce* operation set.

**Contribution of this paper.** We present an interactive verification approach with which a *MapReduce* implementation of an algorithm can be proved equivalent to an imperative implementation (to the best of our knowledge this is the first framework for the purpose of such equivalence proofs, see Sect. 6). Proofs are conducted as chains/sequences of individual, smaller behavior-preserving program transformations.

We came to the important insight that two kinds of sub-proofs are required: (1) uniform transformations changing the control-flow structure that are mostly independent of the particular context in which they are applied; and (2) context-dependent transformations that are not uniform but that preserve the overall structure and can be proved correct using coupling invariants. We identified and describe a catalogue of 13 individual rules. Correctness of 10 of those rules was proven formally using the *Coq* theorem prover.

Our approach has a high potential for automation. The required interaction is designed to be as high-level as possible. The proof is guided by user-specified intermediate programs from which the individual transformations are derived. The rules are designed such that their side conditions can be proved automatically and we describe how pattern matching can be used to allow for a more flexible specification of intermediate steps.

We describe a workflow for integrating this approach with existing interactive theorem provers. We have successfully implemented the approach as a prototype within the interactive theorem prover *Coq* [19] and evaluated the feasibility of our approach by applying it to the *k*-means and PageRank algorithms. These two are prototypical algorithms commonly used as examples in *MapReduce* frameworks, because they exhibit the most common patterns found in large-scale distributed data processing applications. By showing that our approach can be applied to

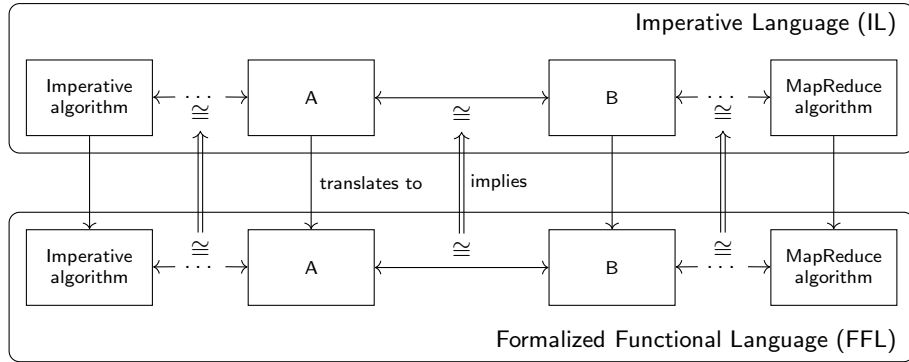


Fig. 1: Chain of equivalent programs is translated into formalized functional language

these two examples, we demonstrate that it can be extended to a much larger set of applications.

**Overview of the approach.** The main challenge in proving the equivalence of an imperative and a *MapReduce* algorithm lies in the potentially large structural difference between two such algorithms. Existing relational verification approaches (like [11,10,15,20]) exploit the fact that the two programs versions to be compared are structurally similar, which allows the verification to focus on describing and proving the similarity of the implementations rather than describing what they actually compute. To deal with the complexity arising from the large structural differences, the equivalence of imperative and *MapReduce* algorithms is not shown in one step, but as a succession of equivalence proofs for structurally closer program versions.

To this end, we require that the translation of the algorithm is broken down (by the user) into a chain of intermediate programs. For each pair of neighboring programs in this chain, the difference is comparatively small and can usually be reduced to one isolated transformation.

The imperative algorithm, the intermediate programs, as well as the *MapReduce* implementation, are given in a high-level imperative programming language (IL). IL is based on a while language and supports integers, booleans, fixed length arrays and sum and product types. It does not support recursion. Besides the imperative language constructs, IL supports *MapReduce* primitives. Given that we have stated previously that *MapReduce* programs tend to be of a more functional nature, it might seem odd at first to not use a functional language for specifying *MapReduce* algorithms. However, this is in accordance with the fact that most of the existing *MapReduce* frameworks are not implemented as separate languages but as frameworks built on top of imperative languages such as Java, Scala, or C++ [22,2]. Thereby sequential parts of *MapReduce* algorithms can still be implemented using imperative language features.

<pre> <b>Function</b> SumArrays(xs,ys) <b>begin</b>   sum ← replicate(n, 0);   <b>for</b> i ← 0 <b>to</b> n - 1 <b>do</b>       sum[i] ← xs[i] + ys[i];   <b>end</b>   <b>return</b> sum; <b>end</b> </pre>	<pre> <b>Function</b> SumArraysZipped(xs,ys) <b>begin</b>   sum ← replicate(n, 0);   zipped ← zip(xs,ys);   <b>for</b> i ← 0 <b>to</b> n - 1 <b>do</b>       sum[i] ← fst(zipped[i]) +           snd(zipped[i]);   <b>end</b>   <b>return</b> sum; <b>end</b> </pre>
---	--

Fig. 2: Two IL programs which calculate the element-wise sum of two arrays.

Each program specified in the high-level imperative language is then automatically translated into the formalized functional language (FFL) described in Sect. 2. The equivalence proofs are conducted on programs in this functional language. An overview of this process can be seen in Fig. 1. For each pair of neighboring programs in the chain, a proof obligation is generated that requires proving their equivalence. These proof obligations are then discharged independently of each other (using the workflow described in Sect. 4). Since, by construction, the semantics of IL programs is the same as that of corresponding FFL programs, the equivalence of two IL programs follows from the equivalence of their translations to FFL. Figure 2 shows two example IL programs for calculating the element-wise sum of two arrays.

The implementation of our approach based on the *Coq* theorem prover has only limited proof automation and still requires a significant amount of interactive proofs. We are convinced, however, that our approach can be extended such that it becomes highly automatised and only few user interactions or none at all are required – besides providing the intermediate programs. Further challenges include the extension of our approach to features such as references and aliasing which are commonly found in imperative languages.

**Structure of this paper.** In Sec. 2, we lay the formal groundwork for our approach by defining the programming language used for equivalence proofs and the notion of program equivalence used in this paper. Sec. 3 describes the two kinds of program transformations that we have identified and the techniques for proving equivalence using these transformations. The technical framework for equivalence proofs and the potential for automation are described in Sec. 4 and their its evaluation is in Sec. 5. In Sec. 6, we discuss work related to the ideas presented in this paper. Finally, we conclude in Sec. 7 and consider possible future work.

## 2 Formal Foundations and Program Equivalence

In this section, we briefly describe the language FFL, introduce a reduction big-step semantics for FFL and discuss the notion of equivalence for FFL programs. A full description of the syntax and semantics of FFL can be found in Appendix A.

The primary design goal of FFL is the capability to represent both imperative and *MapReduce* programs written in IL. To achieve this, we follow the work by Radoi et al. [18] and use a simply typed lambda calculus extended by the theories of sums, products, and arrays. Furthermore, the language also contains the programming primitives usually found in *MapReduce* frameworks. We also want to limit the number of primitives included in FFL while still retaining expressiveness. This simplifies proving general properties of FFL and proving the correctness of rewrite rules. We accomplish this by building upon the work of Chen et al. [5], who describes how to reduce the large number of primitives provided by *MapReduce* frameworks to a smaller core.

Two new primitives `iter` and `fold` were added to translate imperative loops directly. Compared to transforming imperative programs into a recursive form, this allows a translation closer to the original program formulation. The `fold` operator is used to translate bounded for-each iterator loops into FFL. The evaluation of the expression `fold f v0 xs` starts with the initial loop state  $v_0$  and iterates over each value of the array  $xs$  updating the loop state by applying  $f$ . General while loops are translated using the `iter` function. `iter f v0` is evaluated by repeatedly applying  $f$  to the loop state (which is initially  $v_0$ ) until  $f$  returns `unit` to indicate termination. Program terms incorporating `iter` need not evaluate to a value since the construct allows formulating non-terminating programs.

The big-step operational reduction semantics [14] of FFL is defined as a binary relation  $\Rightarrow_{bs}$ . Note that, since FFL is based on lambda calculus, programs in FFL as well as values are FFL expressions. The semantics predicate is thus a partial, functional relation on FFL-terms.

**Definition 1.** *An FFL term  $t$  evaluates to an FFL term  $v$  if  $t \Rightarrow_{bs} v$  holds. A term  $t$  is called stuck if there exists no  $v$  such that  $t \Rightarrow_{bs} v$ . Terms that evaluate to themselves are called values.*

A formal definition of the syntax of FFL is given in form of typing rules in Fig. 12; and the semantics of FFL is shown as inference rules in Fig. 14, both in Appendix A. The evaluation of a program  $t$  in an input state (i.e., for an argument tuple  $a$ ) resulting in an output state  $v$  (a result tuple) can be formalized as the reduction evaluation of the application of the program to the arguments:  $\langle t, a \rangle \Downarrow_{bs} v := \mathbf{app}(t, a) \Rightarrow_{bs} v$ .

The semantics of FFL is deterministic. This may seem odd because most *MapReduce* frameworks take considerable leeway from fully deterministic execution in the name of performance. For example, some operations may be evaluated in a non-deterministic order depending on how fast data arrives over the network leading to non-determinism if these operations are not commutative and associative. However, non-determinism in *MapReduce* algorithms is usually not desired

<pre> fold(<math>\lambda sum. \lambda i.</math>   write(<math>sum, i, xs[i] + ys[i]</math>),   replicate(<math>n, 0</math>),   range(<math>0, n</math>)) </pre>	<pre> snd(iter(<math>\lambda(i, sum).</math>   if <math>i &lt; n</math>   then inr (<math>i + 1,</math>     write(<math>sum, i, xs[i] + ys[i]</math>))   else inl unit,   (<math>0, replicate(n, 0)</math>))) </pre>
(a)	(b)

Fig. 3: Translation of function `SumArrays` (see Fig. 2) into FFL using `fold` and `iter` (where `inl` and `inr` denote the left and right injection into a sumtype).

and the problem of checking whether or not a *MapReduce* algorithm is deterministic is orthogonal to proving that it is equivalent to an imperative algorithm. We thus consider a deterministic language model to be suitable for our purposes and defer checking of determinism to other tools such as those developed by Chen et al. [6,7].

Since FFL includes the potential for run-time errors such as out-of-bound array accesses but does not include an explicit error term, the step-relation  $\Rightarrow_{bs}$  is not total. The absence of an explicit error term also has the consequence that one cannot distinguish between non-termination and runtime errors according to the definition of program equivalence in Definition 2.

The introduction of the semantics relation allows us to define a notion of program equivalence for FFL terms.

**Definition 2.** *Two well-typed FFL terms  $s$  and  $t$  are called equivalent if they (a) are of the same type  $\tau$  and (b) evaluate to the same values  $v$ . We write  $s \cong_{\tau} t$  in this case. Using  $\vdash t : \tau$  to denote that the closed FFL term  $t$  has type  $\tau$ , this definition can be formalised as follows:*

$$s \cong_{\tau} t \quad := \quad \vdash s : \tau \wedge \vdash t : \tau \wedge \forall v. (s \Rightarrow_{bs} v) \Leftrightarrow (t \Rightarrow_{bs} v) \quad (1)$$

This definition of program equivalence also enforces *mutual termination* [9], i.e., the property that equivalent programs either both terminate or both diverge. In particular, two non-terminating terms of the same type are equivalent.

Note that the generated proof obligations require proving the equivalence of functions applied to arbitrary inputs instead of requiring proving the equivalence of functions themselves

*Example 1.* Figure 3 shows two transformations of the function `SumArrays` (see Fig. 2) into FFL. In Fig. 3 (a), the loop is translated using `fold`, and in Fig. 3 (b) using the more general `iter`. In both cases, it can be observed that the local variables  $i$  and  $sum$  become  $\lambda$ -bound variables of the translation of the enclosing block, in this case the loop body.

The first translation has the initial state `replicate( $n, 0$ )`, an array of length  $n$  with all values set to 0, and it iterates over the indices in the array  $([0; 1; \dots; n -$

$$\text{fold}(\lambda acc. \lambda x. f(acc, g(x)), \quad \text{fold}(\lambda acc. \lambda y. f(acc, y),$$

$$\begin{array}{c} i, \\ xs) \end{array} \quad \rightsquigarrow \quad \begin{array}{c} i, \\ \text{map}(g, xs) \end{array}$$

**Side conditions:**  $acc \notin FV(f)$ ,  $x \notin FV(f)$ ,  $y \notin FV(f)$ ,  $x \notin FV(g)$ ,  $acc \notin FV(g)$

Fig. 4: Rewrite rule for separating a loop body into two functions  $f$  and  $g$  such that the evaluation of  $g$  is independent of all other iterations and can be computed in parallel.  $FV(g)$  is the set of free, unbound variables in the term  $g$ .

1]), updating the array  $sum$  in each iteration using the *write* function of the McCarthy theory of arrays.

The translation in Fig. 3 (b) starts from the initial loop state  $(0, \text{replicate}(n, 0))$ . In each iteration, an *if*-clause is used to check if the loop condition still evaluates to *true*. If that is the case, the index is incremented and  $sum$  is updated, otherwise the program exits the loop as indicated by `inl unit` and evaluates to the current loop state.

### 3 Program Transformations

With the reduction of imperative and *MapReduce* implementations to the common language FFL, we are able to prove equivalence between two programs by constructing a chain of single, isolated program transformations. We categorize the transformations by their dependence on the surrounding context. A *context-independent* transformation is a uniform transformation as it replaces only one isolated subterm in the program by an equivalent term. This replacement has no effects on other parts of the program and has only conditions on the replaced subterm. In contrast, *context-dependent* transformations do not replace individual terms but require many small changes throughout different parts of the programs.

For example, consider the IL programs in Fig. 2. In the left IL program, the loop iterates over two separate arrays  $xs$  and  $ys$  of the same length. In the right IL program, the loop iterates over a single array that represents the *zipped* version of  $xs$  and  $ys$ . Inspection of the FFL versions from Fig. 3 shows that this program transformation requires two changes to individual subterms: (a) the initial loop state, and (b) adaption of the read and write references.

We use two complementary techniques for proving the correctness of a transformation depending on whether it is context-independent or context-dependent: The equivalence of programs related by context-independent transformations is proven using rewrite rules (Sect. 3.1) while the equivalence of programs related by context-dependent transformation is shown using coupling predicates (Sect. 3.2).

#### 3.1 Handling Context-Independent Transformations Using Rewrite Rules

Intermediate programs are mostly linked by uniform context-independent transformations on isolated subterms. Instead of performing and proving these local

transformations manually, we can capture them into generalized rewrite rules. That equivalence is preserved when these generalized rewrite rules are applied, needs to be proven only once. By maintaining and using a collection of local transformations that have been proven correct, we can lower proof complexity and later increase the computer assistance and automation.

A rewrite rule describes a bidirectional program transformation that allows the replacement of a subterm within a program. It is composed of two patterns and a set of side conditions which are sufficient for the transformation to preserve program equivalence. A pattern is an FFL term containing metavariables.

To apply a rewrite rule on a program, we have to identify a subterm of the program that (a) matches the first pattern and (b) satisfies the side conditions. The transformed program is obtained by the instantiation of the other pattern with the matched metavariables. Since the sets of bound metavariables in the two patterns can be different, some metavariables may not be uniquely instantiated, leading to a degree of freedom in the translation. We will discuss the practical implications of this in Sect. 4.3.

While there is no hard limit on the complexity of the side conditions that can be part of rewrite rules, it is desirable to use side conditions that are simple and easy to check. This prevents the application of rewrite rules from producing auxiliary complex proofs due to complex side conditions. In our experiments we only encountered the following three different kinds of side conditions:

1. Two arrays  $xs$  and  $ys$  have the same length, i.e.,  $\text{length}(xs) \cong_{\text{int}} \text{length}(ys)$ .
2.  $t$  is not stuck.
3.  $x \notin FV(t)$  where  $FV(t)$  is the set of free variables in the term  $t$ .

Sect. 4.3 discusses how these side conditions could be discharged automatically.

To illustrate the kind of rewrite rules used in the equivalence proofs described in this paper, we present two of the most commonly used rewrite rules in detail. To demonstrate the feasibility of formal correctness proofs for rewrite rules, we have proven the correctness of most (10 out of 13 rules) of our rules in *Coq*. A full listing of all rewrite rules can be found in Appendix B.

The first rule, shown in Fig. 4, decomposes the loop body of a `fold` expression into two separate functions  $f$  and  $g$ , where  $g$  is independent of other iterations. Thus,  $g$  can be computed in parallel using a `map` operation. This rewrite rule illustrates that rewrite rules used in proofs can often also function as guidelines for parallelizing and distributing imperative algorithms.

The second rule, shown in Fig. 5, is similar to the previous rule in that it tries to separate independent parts of the loop body so that they can be executed in parallel. However, in this case, the part that is extracted is only independent of other iterations that access different indices. The `group` operation can be used to group all accesses to the same index. Using `map` one can then calculate the new values for each index in  $xs$  in parallel and update  $ys$  with those new values.



$$\begin{array}{ccc}
\text{fold}(\lambda acc. \lambda(i, x). & & \text{fold}(\lambda acc. \lambda(i, v). \text{write}(acc, i, v), \\
\text{write}(acc, i, f(i, x, acc[i])), & \rightsquigarrow & \text{map}(\lambda(i, vs). \\
ys, & & (i, \text{fold}(\lambda x'. \lambda x. f(i, x, x'), ys[i], vs)), \\
xs) & & \text{group}(xs)))
\end{array}$$

**Side conditions:**  $acc \notin FV(f), x \notin FV(f), x' \notin FV(f), i \notin FV(f), vs \notin FV(f)$

Fig. 5: Rewrite rule for grouping loop iterations which access the same index of an array.

$$\begin{array}{l}
\mathcal{C}(i_0, i'_0) \\
\wedge \quad (\forall i, i', j. \mathcal{C}(i, i') \implies \mathcal{C}(f(i, xs[j]), f'(i', xs'[j]))) \\
\implies \mathcal{C}(\text{fold}(f, i_0, xs), \text{fold}(f', i'_0, xs'))
\end{array}$$

Fig. 6: Coupling invariant rule for `fold` for a coupling predicate  $\mathcal{C}$ . Free variables are implicitly universally quantified.

### 3.2 Handling Context-Dependent Transformations Using Coupling Predicates

While context-independent transformations are nicely handled using rewrite rules, context-dependent transformations can usually not be captured by patterns and simple side conditions. Coupling predicates provide a flexible and effective solution to proving the correctness of context-dependent transformations – at the cost of requiring more user interactions than rewrite rules. The use of coupling predicates is based on the observation that analyzing two loops in lock-step and proving that a relational property, i.e., the coupling predicate, holds after each iteration is sufficient to prove that it holds after the execution of both loops. Fig. 6 shows the corresponding coupling invariant rule for `fold`. For the purpose of presentation, we ignore the distinction between syntactic terms and the values to which they evaluate. Besides this rule for `fold`, there is a similar rule for `iter`.

One compelling example for using coupling predicates is given in the beginning of this section. The presented program transformation is provable equivalent with the coupling invariant rule from Fig. 2. If these arrays are part of the accumulator in a `fold` or `iter`, capturing this transformation by a rule patterns is not possible: While the transformation of the initial accumulator value can be captured using patterns, this is not sufficient since all references to the accumulator in the loop body also need to be updated. These references can be nested arbitrarily deep inside the loop body and there can be arbitrarily many references. This makes it impossible to capture them by a single pattern which can only bind a fixed number of variables and thereby only make a fixed number of transformations. To make matters worse, it is not even sufficient to just transform the loop itself since the loops are not equivalent: the right loop evaluates to two separate arrays while the other evaluates to an array of tuples. It is thus necessary to prove the equivalence of the enclosing terms under the assumption that

the loop in one program evaluates to a tuple of two arrays `pair(xs, ys)` while the other loop evaluates to `zip(xs, ys)`. This assumption can then be proven correct using the coupling predicate stating that this holds after each iteration.

Another commonly found transformation is the removal of unused elements from a tuple representing the loop accumulator. As it was the case for the previous transformation, the loops themselves are not equivalent and it is necessary to prove enclosing terms equivalent using the assumption that the values present in both loop accumulators are equivalent. As before, this assumption can be proven correct using a coupling predicate which states that this holds after each iteration.

## 4 Transformation Application Strategy

Splitting the translation into a chain of intermediate programs and translating these into FFL leaves us with the problem of proving neighboring programs equivalent. In order to reduce the amount of user interaction required to conduct these basic equivalence proofs, we define an iterative heuristic search strategy to identify the locations within the programs on which the program transformations described in Sect. 3 will be applied. Alg. 1 depicts this search strategy as pseudocode. First, we use the structural difference operation (`Diff`, see Sect. 4.1) to identify subterms  $P'$  and  $Q'$  whose equivalence implies the equivalence of the full programs  $P$  and  $Q$ . Second, we start an iterative bottom-up process in which we try to prove the equivalence of the subterms  $P'$ ,  $Q'$  and their enclosing terms (`ProveEquivalent`), until we reached the top level programs  $P$  and  $Q$ . During the bottom-up process, the subterms  $P'$  and  $Q'$  may be found to be equivalent only in some cases but not in others. But that is fine as long as we are able to prove that the cases in which they are non-equivalent are not relevant in the context in which  $P'$  and  $Q'$  occur. Thus, we extract the premises under which  $P'$  and  $Q'$  are equivalent, and bubble them up to the equivalence proof for the parent terms (`AddMissingPremises`, `Widen`, see Sec. 4.2) If we arrive at the top-level terms and cannot prove those equivalent, the proof fails.

### 4.1 Using Congruence Rules to Simplify Proofs

While the difference between neighboring programs in the chain – which are more closely related – tends to be small, the size of these programs can still be large. This complicates interactive proofs for the user, and can also slow down automated proofs. To reduce the complexity, we prove the equivalence of subterms and then use congruence rules to derive the equivalence of the full programs. A concrete example of a congruence rule is shown in Fig. 7a.

We have found that a simple structural comparison (`Diff` in Alg. 1) is well suited for finding smaller subterms whose equivalence implies the equivalence of the full programs. `Diff` computes the smallest two subterms such that replacing them by placeholders results in identical terms. An example of `Diff` can be seen in Fig. 7b.

```

input : Two FFL terms  $P$  and  $Q$ 
output :  $true$  if  $P$  and  $Q$  could be proven equivalent
Premises  $\leftarrow \{\}$ ;
 $(P', Q') \leftarrow \text{Diff}(P, Q)$ ;
repeat
  equivalent?  $\leftarrow \text{ProveEquivalent}(P', Q', \text{Premises})$ ;
  if  $\text{equivalent?}$  then
    return  $true$ ;
  else
    Premises  $\leftarrow \text{AddMissingPremises}(\text{Premises})$ ;
     $(P', Q') \leftarrow \text{Widen}(P', Q')$ ;
  end
until  $P' = P$  and  $Q' = Q$ ;
return  $false$ ;

```

Alg 1. Strategy for individual equivalence proofs between a pair of FFL programs.

$$\frac{xs \cong_{[\alpha]} ys \quad i \cong_{\text{Int}} j}{\text{read}(xs, i) \cong_{\alpha} \text{read}(ys, i)}$$

Fig. 7a: Congruence rule for `read`

$$\begin{aligned} & \text{Diff}(\text{fold}(\lambda(x, y). x + y, 0, xs), \\ & \quad \text{fold}(\lambda(x, y). y + x, 0, xs)) \\ & = (\lambda(x, y). x + y, \lambda(x, y). y + x) \end{aligned}$$

Fig. 7b: Example of applying `Diff`

## 4.2 Missing Premises and Widening

During the iterative bottom-up process in Alg. 1,  $P'$  and  $Q'$  may turn out to be non-equivalent in some cases. The strategy then tries to extract required contextual conditions (premises) that are sufficient to ensure equivalence of  $P'$  and  $Q'$  (`AddMissingPremises`). In the next step, we try to prove the equivalence of enclosing terms (`Widen`), which contain  $P'$  and  $Q'$  as subterms. Additionally, in the widening-step, we take care of the generated premises. These have either to be shown to always hold in the context of `Widen`( $P', Q'$ ) or in the context of further widening.

These two steps – premise extraction and widening – are commonly required to prove the equivalence of loop bodies. The example in Fig. 8 illustrates this. Applying `Diff` instantiates  $P'$  and  $Q'$  with the two loop bodies, as they are

<pre> sum <math>\leftarrow</math> 0; <b>for</b> <math>i \leftarrow</math> 0 <b>to</b> <math>n - 1</math> <b>do</b>     sum <math>\leftarrow</math> sum + xs[<math>i</math>];     xs <math>\leftarrow</math> F'(xs, ys); <b>end</b> </pre>	<pre> sum <math>\leftarrow</math> 0; <b>for</b> <math>i \leftarrow</math> 0 <b>to</b> <math>n - 1</math> <b>do</b>     zipped <math>\leftarrow</math> zip(xs, ys);     sum <math>\leftarrow</math> sum + fst(zipped[<math>i</math>]);     xs <math>\leftarrow</math> F(zipped); <b>end</b> </pre>
---	---

Fig. 8: Two potentially equivalent IL programs operating on two separate arrays (left) and the result of applying `zip` to these arrays (right). `xs`, `ys` are arrays of length  $n$ . `F` and `F'` return arrays of the length of their input.

the topmost non-equal subterms. A coupling invariant implying that the two loops are started in equivalent is not sufficient to ensure equivalent loop states after execution since `zip` is only defined for arrays of the same length. Thus, the coupling invariant needs to include the premise that `xs` and `ys` are of equal length.

In some cases, additional premises sufficient for proving equivalence can be found by working backward from missing assumption in failed proofs. In the example above, proving that the program states are equivalent at the end of each loop iteration assuming that they are equivalent at the beginning will fail due to the missing premise that `xs` and `ys` have the same length. We thus add this premise and try to prove the loop bodies equivalent using that premise. If that is successful, we widen the context to enclosing terms. In the outer context, we attempt to prove that the additional premises are satisfied and derive the equivalence of the full loops based on proved coupling invariant.

### 4.3 Potential for Automation of Proofs using Rewrite Rules

Since equivalence proofs using rewrite rules are particularly common but also quite repetitive, this section is devoted to their potential for proof automation. A graphical overview of the individual steps can be found in Fig. 9.

1. We perform an approximate matching procedure to generate candidate programs which match the patterns in the rewrite rule.
2. We attempt to prove that these candidates are equivalent to the input programs or otherwise we reject them.
3. We prove that the side conditions hold for these candidates.

By the correctness of the rewrite rule, the candidates are equivalent.

**4.3.1 Matching of Rewrite Rules** While automatic rewriting systems have been used in the related context of automatically translating imperative algorithms to *MapReduce* algorithms [18], the specific ways in which rewrite rules are used in our approach brings new challenges as well as simplifications.

The challenge lies in the fact that the intermediate programs often do not match the patterns found in rewrite rules directly. There are two typical solutions: normal forms and generalization of patterns. Both are not applicable here. First, there is no suitable normal form of FFL programs. Additionally, both programs are defined by the user, so we cannot assume a specific program structure. Second, the formulation of generalized rewrite rules for matching the large variety of user-defined programs is difficult to obtain and also their correctness proofs are harder to obtain.

The benefit of the programs  $A, A'$  (resp.  $B, B'$ ) being provided by the user is that this can reduce ambiguities. In particular, the schematic variables in the two patterns usually overlap to a large degree, but not fully. The matching of the program  $A$  against the corresponding pattern can lead to unassigned metavariables, which we need to instantiate with correct choices to prove the

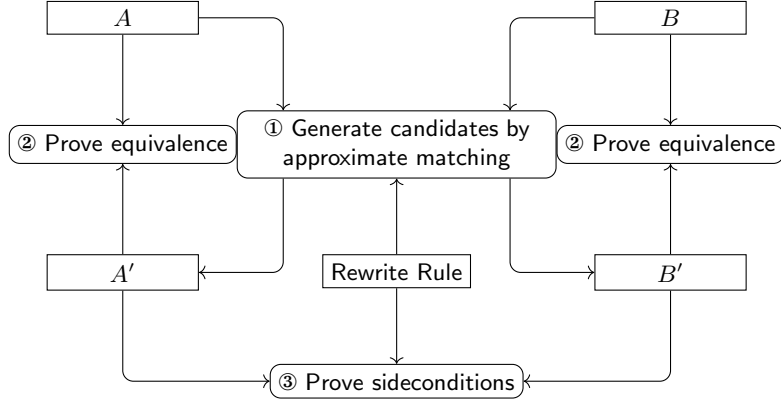


Fig. 9: Workflow for equivalence proofs using rewrite rules. The user has to provide the programs  $A, A', B, B'$  and also the rewrite rule. The equivalence proofs ③ are computer-aided in *Coq*.

equivalence. Now, we have the benefit, that the target program  $A'$  is also defined by the user. So, we can obtain missing assignments by matching the other pattern against the other program.

To find the intermediate programs which match the patterns in rewrite rules, an approximate match procedure is used to find assignments for schematic variables in patterns. The approximate matching is an extension of the classical pattern with the background knowledge and heuristic of easy-to-prove differences. Applying these assignments to patterns yields candidates for the intermediate program. Once two candidates that match the patterns in a rewrite rule have been identified, it is necessary to prove that (a) the candidates are equivalent to the programs used as the input of the approximate matching procedure, and (b) the side conditions hold and the equivalence of the candidates follows thereby from the correctness of the rewrite rule..

While we have only implemented rudimentary partial automation of the equivalence proof construction, analyzing the *Coq* proofs produced in our experiments has shown that these proofs can be reduced to the correctness of a small number of simple transformations. Proving the correctness of these transformations automatically is feasible and could drastically reduce the need for user interaction.

During our evaluation, one of the most prevalent transformations is *call-by-name beta-reduction* or *lambda abstraction* depending on the direction of the transformation for proving the equivalence (② in Fig. 9). Call-by-name beta-reduction refers to the *beta-reduction* found in programming languages with lazy semantics, which contrary to *call-by-value beta-reduction* does not evaluate the argument before applying substitution. Since we are working in a call-by-value setting, call-by-name beta-reduction does not always produce an equivalent program. However, the resulting program is equivalent if, for each case where the

argument would have been evaluated in the original program, all occurrences in the new program will also be evaluated.

Most other transformations are special cases of constant-folding, e.g., reducing expressions such as  $\text{fst}(\text{pair}(a, b))$  to  $a$ . Constant-folding does not produce an equivalent program in general if the terms that are being folded are inside the body of a lambda. A sufficient criterion for the resulting program to be equivalent is that the terms being folded are always evaluated.

**4.3.2 Proving Side Conditions** In Sect. 3.1 we listed the three different kinds of side conditions used in our rewrite rules. The first of those,  $x \notin FV(t)$ , is purely syntactical and can easily be checked automatically. While proving that a term is not stuck can be difficult in general, in our experiments this could usually be reduced to the term being a value, which again is a syntactical condition. The third kind of side condition states that two arrays have the same length. This can usually be proven recursively by reduction to operations that produce arrays of a specific length, e.g.,

$$\forall n, a, b. \text{length}(\text{replicate}(n, a)) = \text{length}(\text{replicate}(n, b)) ,$$

or to length-preserving operations such as `map`. Note that it can be necessary to strengthen loop invariants to carry this fact through a loop as explained in Sect. 4.2.

## 5 Evaluation and Case Study

To demonstrate the feasibility of our approach, we have created a toolchain. The user specifies a sequence of intermediate programs in a simple imperative language. These programs are then automatically translated into a formalization of the previously described functional programming language FFL in *Coq*. In addition to generating proof obligations, our toolchain reduces these obligations using the mentioned structural comparison `Diff`, and it applies congruence rules to reconstruct an equivalence proof of the full programs.

Using this toolchain, we have proven the equivalence of imperative and *MapReduces* implementations of the *PageRank* algorithm [3] and the *k-means* [16] algorithm in *Coq*. Fig. 10 shows the imperative and the *MapReduce* implementation of *PageRank* that we have used in our experiments.

While we have created the imperative implementations of the two algorithms ourselves, the *MapReduce* versions are very close to the implementations accompanying the Thrill [2] framework. This reinforces our claim that FFL is capable of representing *MapReduce* algorithms and is thereby suitable for this approach. In total, the formalization of FFL, the rewrite rules, and proofs of various properties, encompasses about 8000 lines of *Coq* code. The equivalence proofs of *PageRank* and *k-means* each require about 3700 lines of *Coq* proofs. That includes the automatically generated translation of the chain of equivalent programs (for *k-means* this chain consists of 9 programs while for *PageRank* it consists of 6 programs),

```

Function PageRank(links, numLinks, n)
begin
  ranks ←
    Replicate(numLinks,  $\frac{1}{\text{numLinks}}$ );
  for i = 1 to n do
    ranks' ← Replicate(numLinks, 0);
    for p = 0 to numLinks - 1 do
      contrib ←  $\frac{\text{ranks}[p]}{\text{Length}(\text{links}[p])}$ ;
      foreach q ← links[p] do
        ranks'[q] ←
          ranks'[q] + contrib;
      end
    end
    for p = 0 to numLinks - 1 do
      ranks[p] ←
        Dampen(ranks'[p], numLinks);
    end
  end
return ranks;
end

Function PageRank(links, numLinks, n)
begin
  ranks ←
    Replicate(numLinks,  $\frac{1}{\text{numLinks}}$ );
  for i = 1 to n do
    outRanks ← Zip(links, ranks);
    contribs ←
      FlatMap(
         $\lambda(\text{ls}, r).$ 
        Map( $\lambda l.(l, \frac{r}{\text{Length}(l)})$ , ls),
        outRanks);
    rankUpdates ←
      Reduce(+, 0, contribs);
    ranks' ← Replicate(numLinks, 0);
    foreach (l, r) ← rankUpdates do
      ranks'[l] ← r;
    end
    ranks ←
      Map( $\lambda r.$  Dampen(r, numLinks),
        ranks');
  end
return ranks;
end

```

Fig. 10: Imperative (left) and *MapReduce* (right) versions of the *PageRank* algorithm (the function `Replicate(n,v)` creates an array of length `n` with all elements set to `v`; and `Dampen` is an arbitrary function).

which take up large parts of these proofs. The proofs rely on the rewrite rules which we have formalized in *Coq* as well as coupling predicates.

## 6 Related Work

A common approach to relational verification and program equivalence is the use of product programs [1]. Product programs combine the states of two programs and interleave their behavior in a single program. *RVT* [11] proves the equivalence of C programs by combining them in a product program. By assuming that the program states are equal after each loop iteration, *RVT* avoids the need for user-specified or inferred loop invariants and coupling predicates.

Hawblitzel et al. [13] use a similar technique for handling recursive function calls. Felsing et al. [10] demonstrate that coupling predicates for proving the equivalence of two programs can often be inferred automatically. While the structure of imperative and *MapReduce* algorithms tends to be quite different, splitting the translation into intermediate steps yields programs which are often structurally similar. We have found that in this case, techniques such as coupling predicates arise naturally and are useful for selected parts of an equivalence proof.

Radoi et al. [18] describe an automatic translation of imperative algorithms to *MapReduce* algorithms based on rewrite rules. While the rewrite rules are very similar to the ones used in our approach, we complement rewrite rules by coupling predicates. Furthermore we are able to prove equivalence for algorithms

for which the automatic translation from Radoi et al. is not capable of producing efficient *MapReduce* algorithms. The objective of verification imposes different constraints than the automated translation – in particular both programs are provided by the user, so there is less flexibility needed in the formulation of rewrite rules.

Chen et al. [5] and Radoi et al. [18] describe languages and sequential semantics for *MapReduce* algorithms. Chen et al. describe an executable sequential specification in the Haskell programming language focusing on capturing non-determinism correctly. Radoi et al. use a language based on a lambda calculus as the common representation for the previously described translation from imperative to *MapReduce* algorithms. While this language closely resembles the language used in our approach, it lacks support for representing some imperative constructs such as arbitrary *while*-loops.

Grossman et al. [12] verify the equivalence of a restricted subset of Spark programs by reducing the problem of checking program equivalence to the validity of formulas in a decidable fragment of first-order logic. While this approach is fully automatic, it limits programs to Presburger arithmetic and requires that they are synchronized in some way.

To the best of our knowledge, we are the first to propose a framework for proving equivalence of *MapReduce* and imperative programs.

## 7 Conclusion

We have presented a new approach for proving the equivalence of imperative and *MapReduce* algorithms. This approach relies on splitting the transformation into a chain of intermediate programs. The individual equivalence proofs are then categorized in context-independent and context-dependent transformations. Equivalence proofs for context-independent transformations are handled using rewrite rules, while equivalence proofs for context-dependent transformations are based on coupling predicates. We have demonstrated the feasibility of end-to-end equivalence proofs using this approach by applying it two well-known non-trivial algorithms.

While we have hinted at the potential for automating this approach, implementing automation is left as future work. In particular, it would be interesting to explore whether existing tools for relational verification using coupling predicates can be used or if new tools are necessary. Further future work includes extending the approach presented here to support the full expressiveness provided by languages which are used to implement imperative and *MapReduce* algorithms.

## References

1. Barthe, G., Crespo, J.M., Kunz, C.: Relational Verification Using Product Programs, pp. 200–214. Springer Berlin Heidelberg, Berlin, Heidelberg (2011), [http://dx.doi.org/10.1007/978-3-642-21437-0\\_17](http://dx.doi.org/10.1007/978-3-642-21437-0_17)



2. Bingmann, T., Axtmann, M., Jöbstl, E., Lamm, S., Nguyen, H.C., Noe, A., Schlag, S., Stumpp, M., Sturm, T., Sanders, P.: Thrill: High-performance algorithmic distributed batch data processing with C++. In: IEEE International Conference on Big Data. pp. 172–183. IEEE (Dec 2016), <https://doi.org/10.1109/BigData.2016.7840603>, preprint arXiv:1608.05634
3. Brin, S., Page, L.: The anatomy of a large-scale hypertextual web search engine. *Comput. Netw. ISDN Syst.* 30(1-7), 107–117 (Apr 1998), [http://dx.doi.org/10.1016/S0169-7552\(98\)00110-X](http://dx.doi.org/10.1016/S0169-7552(98)00110-X)
4. Chambers, C., Raniwala, A., Perry, F., Adams, S., Henry, R.R., Bradshaw, R., Weizenbaum, N.: FlumeJava: Easy, efficient data-parallel pipelines. *ACM SIGPLAN Notices* 45(6), 363–375 (2010)
5. Chen, Y.F., Hong, C.D., Lengál, O., Mu, S.C., Sinha, N., Wang, B.Y.: An executable sequential specification for spark aggregation (2017)
6. Chen, Y.F., Hong, C.D., Sinha, N., Wang, B.Y.: Commutativity of Reducers, pp. 131–146. Springer (2015), [http://dx.doi.org/10.1007/978-3-662-46681-0\\_9](http://dx.doi.org/10.1007/978-3-662-46681-0_9)
7. Chen, Y., Song, L., Wu, Z.: The commutativity problem of the mapreduce framework: A transducer-based approach. *CoRR* abs/1605.01497 (2016), <http://arxiv.org/abs/1605.01497>
8. Dean, J., Ghemawat, S.: MapReduce: Simplified data processing on large clusters. *Commun. ACM* 51(1), 107–113 (Jan 2008), <http://doi.acm.org/10.1145/1327452.1327492>
9. Elenbogen, D., Katz, S., Strichman, O.: Proving mutual termination. *Form. Methods Syst. Des.* 47(2), 204–229 (Oct 2015), <http://dx.doi.org/10.1007/s10703-015-0234-3>
10. Felsing, D., Grebing, S., Klebanov, V., Rümmer, P., Ulbrich, M.: Automating regression verification. In: Proceedings of the 29th ACM/IEEE International Conference on Automated Software Engineering. pp. 349–360. ASE '14, ACM, New York, NY, USA (2014), <http://doi.acm.org/10.1145/2642937.2642987>
11. Godlin, B., Strichman, O.: Regression verification. In: Proceedings of the 46th Annual Design Automation Conference. pp. 466–471. DAC '09, ACM, New York, NY, USA (2009), <http://doi.acm.org/10.1145/1629911.1630034>
12. Grossman, S., Cohen, S., Itzhaky, S., Rinetzky, N., Sagiv, M.: Verifying Equivalence of Spark Programs, pp. 282–300. Springer International Publishing, Cham (2017), [https://doi.org/10.1007/978-3-319-63390-9\\_15](https://doi.org/10.1007/978-3-319-63390-9_15)
13. Hawblitzel, C., Kawaguchi, M., Lahiri, S., Rebêlo, H.: Mutual summaries: Unifying program comparison techniques. In: Informal proceedings of BOOGIE 2011 workshop (2011), <https://www.microsoft.com/en-us/research/publication/mutual-summaries-unifying-program-compari>
14. Kahn, G.: Natural semantics. *STACS 87* pp. 22–39 (1987)
15. Lahiri, S.K., Hawblitzel, C., Kawaguchi, M., Rebêlo, H.: Symdiff: A language-agnostic semantic diff tool for imperative programs. In: Proceedings of the 24th International Conference on Computer Aided Verification. pp. 712–717. CAV'12, Springer-Verlag, Berlin, Heidelberg (2012), [http://dx.doi.org/10.1007/978-3-642-31424-7\\_54](http://dx.doi.org/10.1007/978-3-642-31424-7_54)
16. Lloyd, S.: Least squares quantization in PCM. *IEEE Transactions on Information Theory* 28(2), 129–137 (1982), <https://doi.org/10.1109/TIT.1982.1056489>
17. McCarthy, J.: A basis for a mathematical theory of computation1). In: Braffort, P., Hirschberg, D. (eds.) *Computer Programming and Formal Systems*, Studies in Logic and the Foundations of Mathematics, vol. 35, pp. 33 – 70. Elsevier (1963), <http://www.sciencedirect.com/science/article/pii/S0049237X08720184>

18. Radoi, C., Fink, S.J., Rabbah, R., Sridharan, M.: Translating imperative code to mapreduce. SIGPLAN Not. 49(10), 909–927 (Oct 2014), <http://doi.acm.org/10.1145/2714064.2660228>
19. development team, T.C.: The Coq proof assistant reference manual. LogiCal Project (2004), <http://coq.inria.fr>, version 8.6
20. Verdoolaege, S., Janssens, G., Bruynooghe, M.: Equivalence checking of static affine programs using widening to handle recurrences. ACM Trans. Program. Lang. Syst. 34(3), 11:1–11:35 (Nov 2012), <http://doi.acm.org/10.1145/2362389.2362390>
21. White, T.: Hadoop: The definitive guide. O’Reilly Media, Inc. (2012)
22. Zaharia, M., Chowdhury, M., Franklin, M.J., Shenker, S., Stoica, I.: Spark: Cluster computing with working sets. In: Proceedings of the 2Nd USENIX Conference on Hot Topics in Cloud Computing. pp. 10–10. HotCloud’10, USENIX Association, Berkeley, CA, USA (2010), <http://dl.acm.org/citation.cfm?id=1863103.1863113>

## A Syntax and Semantics of FFL

### A.1 Syntax

The programming language FFL is used in our approach to represent the imperative input program, the target *MapReduce* implementations, and intermediate programs in which both programming paradigms appear in combination. FFL is designed to accommodate both functional and imperative concepts. It bases on

1. *simply typed lambda calculus* with the following theories:
2. *direct sums* and *products*,
3. (mathematical) *integers*,
4. McCarthy's theory of *arrays* [17],
5. *iterating* primitives to represent imperative loops,
6. and primitives found in *MapReduce* frameworks.

Fig. 11 shows an exhaustive list of the constructors of FFL. These can be classified according to the six categories mentioned above.

#### Simply typed lambda calculus

- variables  $x$ , function abstraction  $\lambda x. e$  and application  $\text{app}(e_1, e_2)$
- if  $e_1$  then  $e_t$  else  $e_f$
- Boolean literals  $\text{bool}(b)$  for  $b \in \{\text{true}, \text{false}\}$

#### Direct sums and products

- introduction  $\text{pair}(e_1, e_2)$  and elimination forms  $\text{fst}(e)$ ,  $\text{snd}(e)$  of binary product types
- introduction  $\text{inl}(e)$ ,  $\text{inr}(e)$ , and elimination forms

$$\text{case } e \text{ of } \text{inl}(l) \mapsto e_l; \text{inr}(r) \mapsto e_r$$

of binary sum types

$b := \text{true} \mid \text{false}$			
$i := \dots \mid -2 \mid -1 \mid 0 \mid 1 \mid 2 \mid \dots$			
$e := x$			
$\text{app}(e, e)$	$\lambda x. e$		
$\text{int}(i)$	$\text{bool}(b)$	$\text{list}[e, \dots, e]$	
$\text{add}(e, e)$	$\text{sub}(e, e)$	$\text{mul}(e, e)$	
$\text{gt}(e, e)$	$\text{lt}(e, e)$		
$\text{unit}$	$\text{pair}(e, e)$		
$\text{fst}(e)$	$\text{snd}(e)$		
$\text{inl}(e)$	$\text{inr}(e)$	$\text{case } e \text{ of } \text{inl}(l) \mapsto e; \text{inr}(r) \mapsto e$	
$\text{iter}(e, e)$	$\text{fold}(e, e, e)$	$\text{if } e \text{ then } e \text{ else } e$	
$\text{read}(e, e)$	$\text{write}(e, e, e)$	$\text{readAtKey}(e, e)$	$\text{writeAtKey}(e, e, e)$
$\text{replicate}(e, e)$	$\text{range}(e, e)$	$\text{length}(e)$	
$\text{map}(e, e)$	$\text{group}(e)$	$\text{zip}(e, e)$	$\text{concat}(e)$

Fig. 11: Grammar for terms  $e$  of FFL.

- unit, the single inhabitant of Unit

**Theory of integers**

- integer literals  $\text{int}(i)$  for  $i \in \mathbb{Z}$
- comparison of integers  $\text{gt}(e_1, e_2)$ ,  $\text{lt}(e_1, e_2)$
- binary operations on integers  $\text{add}(e_1, e_2)$ ,  $\text{sub}(e_1, e_2)$ ,  $\text{mul}(e_1, e_2)$

**Theory of arrays**

- array literals  $\text{list}[e_1, \dots, e_n]$
- array read  $\text{read}(e_1, e_2)$
- array write  $\text{write}(e_1, e_2, e_3)$
- $\text{replicate}(e_1, e_2)$  which produces an array of the specified length  $e_1$  containing the same element  $e_2$  at each index,
- $\text{range}(e_1, e_2)$  which produces an array containing all elements in the range from  $e_1$  to  $e_2$  and
- $\text{length}(e)$  which returns the length of array  $e$ .

**Iteration functions**

- $\text{iter}(e_1, e_2)$  and
- $\text{fold}(e_1, e_2, e_3)$

**MapReduce primitives**

- $\text{map}(f, xs)$  which applies a function  $f$  to all elements in the array  $xs$
- $\text{zip}(xs, ys)$  which combines two arrays  $xs, ys$  of the same length into an array containing pairs of elements of  $xs$  and  $ys$
- $\text{concat}(xss)$  which flattens an array of arrays by concatenating the arrays
- operations which operate on arrays of key-value pairs:  $\text{readAtKey}(xs, k)$  returns the value associated with a key
- $\text{writeAtKey}(xs, k, v)$  sets the value associated with a key, and  $\text{group}(xs)$  produces an array that associates each key in the original array with all values associated with it
- $\text{readAtKey}$  and  $\text{writeAtKey}$  operate on the first pair with a matching key; for  $\text{group}$  the order of the keys in the resulting array matches the order of their first occurrences in the input and the order of the values associated with each key matches their order in the input array

Fig. 12 lists the typing rules for FFL. The set of expressions derivable using the given rules is the set of well-typed programs in FFL.  $\Gamma$  denotes the context (type assignment for free variables) under which the rules apply. The following type constructors are used in the figure.

Type	Terms of this type
Int/Bool	integers/booleans
$\alpha \times \beta$	pairs of terms of type $\alpha$ and $\beta$
$\alpha + \beta$	elements of type $\alpha$ or $\beta$ (direct sum)
$\alpha \rightarrow \beta$	functions from terms of type $\alpha$ to terms of type $\beta$
$[\alpha]$	arrays of terms of type $\alpha$
Unit	unit

$$\begin{array}{c}
\frac{}{\Gamma, x : \tau \vdash x : \tau} \\
\frac{\Gamma, x : \alpha \vdash e : \beta}{\Gamma \vdash \lambda x. e : \alpha \rightarrow \beta} \\
\frac{b \in \{\text{true}, \text{false}\}}{\Gamma \vdash \text{bool}(b) : \text{Bool}} \\
\frac{\Gamma \vdash a : \text{Int} \quad \Gamma \vdash b : \text{Int}}{\Gamma \vdash \text{add}(a, b) : \text{Int}} \\
\frac{\Gamma \vdash a : \text{Int} \quad \Gamma \vdash b : \text{Int}}{\Gamma \vdash \text{mul}(a, b) : \text{Int}} \\
\frac{\Gamma \vdash a : \text{Int} \quad \Gamma \vdash b : \text{Int}}{\Gamma \vdash \text{gt}(a, b) : \text{Bool}} \\
\frac{\Gamma \vdash b : \text{Bool} \quad \Gamma \vdash e_t : \tau \quad \Gamma \vdash e_f : \tau}{\Gamma \vdash \text{if } b \text{ then } e_t \text{ else } e_f : \tau} \\
\frac{\Gamma \vdash ts : [\tau]}{\Gamma \vdash \text{length}(ts) : \text{Int}} \\
\frac{\Gamma \vdash p : \alpha \times \beta}{\Gamma \vdash \text{fst}(p) : \alpha} \\
\frac{\Gamma \vdash a : \alpha}{\Gamma \vdash \text{inl}(a) : \alpha + \beta} \\
\frac{\Gamma \vdash e : \alpha + \beta \quad \Gamma, l : \alpha \vdash a : \tau \quad \Gamma, r : \beta \vdash b : \tau}{\Gamma \vdash \text{case } e \text{ of } \text{inl}(l) \mapsto a; \text{inr}(r) \mapsto b : \tau} \\
\frac{\Gamma \vdash f : \alpha \times \beta \rightarrow \alpha \quad \Gamma \vdash e : \alpha \quad \Gamma \vdash ts : [\beta]}{\Gamma \vdash \text{fold}(f, e, ts) : \alpha} \\
\frac{\Gamma \vdash f : \alpha \rightarrow \text{Unit} + \alpha \quad \Gamma \vdash e : \alpha}{\Gamma \vdash \text{iter}(f, e) : \alpha} \\
\frac{\Gamma \vdash ts : [\tau] \quad \Gamma \vdash i : \text{Int}}{\Gamma \vdash \text{read}(ts, i) : \tau} \\
\frac{\Gamma \vdash ts : [\alpha \times \beta] \quad \Gamma \vdash k : \alpha}{\Gamma \vdash \text{readAtKey}(ts, k) : \text{unit} + \beta} \\
\frac{\Gamma \vdash n : \text{Int} \quad \Gamma \vdash t : \alpha}{\Gamma \vdash \text{replicate}(n, t) : [\alpha]} \\
\frac{\Gamma \vdash ts : [\alpha \times \beta]}{\Gamma \vdash \text{group}(ts) : [\alpha \times [\beta]]} \\
\frac{\Gamma \vdash tss : [[\alpha]]}{\Gamma \vdash \text{concat}(tss) : [\alpha]} \\
\frac{\Gamma \vdash f : \alpha \rightarrow \beta \quad \Gamma \vdash t : \alpha}{\Gamma \vdash \text{app}(f, t) : \beta} \\
\frac{\Gamma \vdash t_1 : \tau \quad \dots \quad \Gamma \vdash t_n : \tau}{\Gamma \vdash \text{list}[t_1, \dots, t_n] : [\tau]} \\
\frac{i \in \mathbb{Z}}{\Gamma \vdash \text{int}(i) : \text{Int}} \\
\frac{\Gamma \vdash a : \text{Int} \quad \Gamma \vdash b : \text{Int}}{\Gamma \vdash \text{sub}(a, b) : \text{Int}} \\
\frac{\Gamma \vdash a : \text{Int} \quad \Gamma \vdash b : \text{Int}}{\Gamma \vdash \text{lt}(a, b) : \text{Bool}} \\
\frac{}{\Gamma \vdash \text{unit} : \text{Unit}} \\
\frac{\Gamma \vdash a : \alpha \quad \Gamma \vdash b : \beta}{\Gamma \vdash \text{pair}(a, b) : \alpha \times \beta} \\
\frac{\Gamma \vdash p : \alpha \times \beta}{\Gamma \vdash \text{snd}(p) : \beta} \\
\frac{\Gamma \vdash b : \beta}{\Gamma \vdash \text{inr}(b) : \alpha + \beta} \\
\frac{\Gamma \vdash ts : [\tau] \quad \Gamma \vdash i : \text{Int} \quad \Gamma \vdash v : \tau}{\Gamma \vdash \text{write}(ts, i, v) : [\tau]} \\
\frac{\Gamma \vdash ts : [\alpha \times \beta] \quad \Gamma \vdash k : \alpha \quad \Gamma \vdash t : \beta}{\Gamma \vdash \text{writeAtKey}(ts, k, t) : [\alpha \times \beta]} \\
\frac{\Gamma \vdash l : \text{Int} \quad \Gamma \vdash r : \text{Int} \quad \Gamma \vdash s : \text{Int}}{\Gamma \vdash \text{range}(l, r) : [\text{Int}]} \\
\frac{\Gamma \vdash f : \alpha \rightarrow \beta \quad \Gamma \vdash ts : [\alpha]}{\Gamma \vdash \text{map}(f, ts) : [\beta]} \\
\frac{\Gamma \vdash as : [\alpha] \quad \Gamma \vdash bs : [\beta]}{\Gamma \vdash \text{zip}(as, bs) : [\alpha \times \beta]}
\end{array}$$

Fig. 12: Typing rules for FFL.

## A.2 Semantics

The semantics of FFL is given as a reduction big-step semantics formalized as a binary predicate  $\Rightarrow_{bs}$  on well-typed and closed FFL terms. The values (terms which evaluate to themselves) in FFL are:  $\lambda$ -abstractions, literals, and sums and products of values. Fig. 14 lists the evaluation rules for most constructs and Fig. 13 shows those for the key-value operations in FFL.

$$\begin{array}{c}
 \frac{xs \Rightarrow_{bs} [(k_I, v_I), \dots, (k_j, v_j), \dots, (k_n, v_n)] \quad \forall 1 \leq i < j, k_i \neq k \quad k_j = k}{\text{readAtKey}(xs, k) \Rightarrow_{bs} \text{inr}(v)} \\
 \\
 \frac{xs \Rightarrow_{bs} [(k_I, v_I), \dots, (k_n, v_n)] \quad \forall 1 \leq i \leq n, k_i \neq k}{\text{readAtKey}(xs, k) \Rightarrow_{bs} \text{inl}(\text{unit})} \\
 \\
 \frac{xs \Rightarrow_{bs} [(k_I, v_I), \dots, (k_j, v_j), \dots, (k_n, v_n)] \quad \forall 1 \leq i < j, k_i \neq k \quad k_j = k}{\text{writeAtKey}(xs, k, v) \Rightarrow_{bs} [(k_I, v_I), \dots, (k_j, v), \dots, (k_n, v_n)]} \\
 \\
 \frac{xs \Rightarrow_{bs} [(k_I, v_I), \dots, (k_n, v_n)] \quad \forall 1 \leq i \leq n, k_i \neq k}{\text{writeAtKey}(xs, k, v) \Rightarrow_{bs} [(k_I, v_I), \dots, (k_n, v_n), (k, v)]} \\
 \\
 \frac{xs \Rightarrow_{bs} []}{\text{group}(xs) \Rightarrow_{bs} []} \\
 \\
 \frac{xs \Rightarrow_{bs} [(k_I, v_I), \dots, (k_n, v_n)] \quad \text{group}([(k_1, v_1), \dots, (k_{n-1}, v_{n-1})]) \Rightarrow_{bs} [(k'_1, vs_1), \dots, (k'_m, vs_m)] \quad \forall 1 \leq i \leq m, k_n \neq k'_i}{\text{group}(xs) \Rightarrow_{bs} [(k'_I, vs_I), \dots, (k'_m, vs_m), (k_n, [v_n])]} \\
 \\
 \frac{xs \Rightarrow_{bs} [(k_1, v_1), \dots, (k_n, v_n)] \quad \text{group}([(k_1, v_1), \dots, (k_n, v_n)]) \Rightarrow_{bs} [(k'_1, vs_1), \dots, (k'_j, [v'_1, \dots, v'_n]), \dots, (k'_n, vs_n)] \quad \forall 1 \leq i < j, k'_n \neq k_i \quad k_j = k}{\text{group}(xs) \Rightarrow_{bs} [(k'_I, vs_I), \dots, (k, [v'_1, \dots, v'_n, v_n]), \dots, (k_n, [v_n])]}
 \end{array}$$

Fig.13: Semantics of key-value operations in FFL (for readability, we use the shorthands  $[t_1, \dots, t_n]$  for list literals  $\text{list}[t_1, \dots, t_n]$ , and  $(t_1, t_2)$  for pairs  $\text{pair}(t_1, t_2)$ ).

$$\begin{array}{c}
\frac{f \Rightarrow_{bs} \lambda x. b \quad t \Rightarrow_{bs} t' \quad [t'/x]b \Rightarrow_{bs} v}{\text{app}(f, t) \Rightarrow_{bs} v} \\
\frac{}{\lambda x. e \Rightarrow_{bs} \lambda x. e} \\
\frac{}{\text{bool}(b) \Rightarrow_{bs} \text{bool}(b)} \\
\frac{t_1 \Rightarrow_{bs} t'_1 \quad \dots \quad t_n \Rightarrow_{bs} t'_n}{\text{list}[t_1, \dots, t_n] \Rightarrow_{bs} \text{list}[t'_1, \dots, t'_n]} \\
\frac{b \Rightarrow_{bs} \text{bool}(\text{true}) \quad e_t \Rightarrow_{bs} e_{t'}}{\text{if } b \text{ then } e_t \text{ else } e_f \Rightarrow_{bs} e_{t'}} \\
\frac{a \Rightarrow_{bs} \text{int}(i) \quad b \Rightarrow_{bs} \text{int}(j)}{\text{add}(a, b) \Rightarrow_{bs} \text{int}(i + j)} \\
\frac{a \Rightarrow_{bs} \text{int}(i) \quad b \Rightarrow_{bs} \text{int}(j)}{\text{mul}(a, b) \Rightarrow_{bs} \text{int}(i \cdot j)} \\
\frac{a \Rightarrow_{bs} \text{int}(i) \quad b \Rightarrow_{bs} \text{int}(j)}{\text{gt}(a, b) \Rightarrow_{bs} \text{bool}(i > j)} \\
\frac{}{\text{unit} \Rightarrow_{bs} \text{unit}} \\
\frac{p \Rightarrow_{bs} \text{pair}(a, b)}{\text{fst}(p) \Rightarrow_{bs} a} \\
\frac{a \Rightarrow_{bs} a'}{\text{inl}(a) \Rightarrow_{bs} \text{inl}(a')} \\
\frac{e \Rightarrow_{bs} \text{inl}(v) \quad [v/l]a \Rightarrow_{bs} a'}{\text{case } e \text{ of } \text{inl}(l) \mapsto a; \text{inr}(r) \mapsto b \Rightarrow_{bs} a'} \\
\frac{i \Rightarrow_{bs} \text{int}(i') \quad ts \Rightarrow_{bs} \text{list}[t_1, \dots, t_{i'}, \dots, t_n]}{\text{read}(ts, i) \Rightarrow_{bs} t_{i'}} \\
\frac{i \Rightarrow_{bs} \text{int}(i') \quad ts \Rightarrow_{bs} \text{list}[t_1, \dots, t_{i'}, \dots, t_n] \quad t \Rightarrow_{bs} t'}{\text{write}(ts, i, t) \Rightarrow_{bs} \text{list}[t_1, \dots, t', \dots, t_n]} \\
\frac{acc \Rightarrow_{bs} acc' \quad \text{app}(f, acc') \Rightarrow_{bs} \text{inl}(\text{unit})}{\text{iter}(f, acc) \Rightarrow_{bs} acc'} \\
\frac{f \Rightarrow_{bs} \lambda f. b \quad acc \Rightarrow_{bs} acc' \quad ts \Rightarrow_{bs} \text{list}[]}{\text{fold}(f, acc, ts) \Rightarrow_{bs} acc'} \\
\frac{ts \Rightarrow_{bs} \text{list}[t_1, t_2, \dots, t_n] \quad \text{app}(f, \text{pair}(acc, t_1)) \Rightarrow_{bs} acc' \quad \text{fold}(f, acc', \text{list}[t_2, \dots, t_n]) \Rightarrow_{bs} v}{\text{fold}(f, acc, ts) \Rightarrow_{bs} v} \\
\frac{n \Rightarrow_{bs} \text{int}(i) \quad t \Rightarrow_{bs} v}{\text{replicate}(n, v) \Rightarrow_{bs} \text{list}[\underbrace{v, \dots, v}_i]} \\
\frac{l \Rightarrow_{bs} \text{int}(l') \quad h \Rightarrow_{bs} \text{int}(h')}{\text{range}(l, h) \Rightarrow_{bs} \text{list}[\text{int}(l'), \text{int}(l' + 1), \dots, \text{int}(h' - 1)]} \\
\frac{ts \Rightarrow_{bs} \text{list}[t_1, \dots, t_n] \quad \text{app}(f, t_1) \Rightarrow_{bs} t'_1 \quad \dots \quad \text{app}(f, t_n) \Rightarrow_{bs} t'_n}{\text{map}(f, ts) \Rightarrow_{bs} \text{list}[t'_1, \dots, t'_n]} \\
\frac{as \Rightarrow_{bs} \text{list}[a_1, \dots, a_n] \quad bs \Rightarrow_{bs} \text{list}[b_1, \dots, b_n]}{\text{zip}(as, bs) \Rightarrow_{bs} \text{list}[\text{pair}(a_1, b_1), \dots, \text{pair}(a_n, b_n)]} \\
\frac{tss \Rightarrow_{bs} \text{list}[\text{list}[t_{1,1}, \dots, t_{1,n}], \dots, \text{list}[t_{m,1}, \dots, t_{m,n'}]]}{\text{concat}(tss) \Rightarrow_{bs} \text{list}[t_{1,1}, \dots, t_{1,n}, \dots, t_{m,1}, \dots, t_{m,n'}]}
\end{array}$$

Fig. 14: Big-step semantics of FFL.

## B Rewrite Rules

Below we list all rewrite rules that we have identified. Rule 1 has been explained previously and deals with extracting parts of the loop body that are independent of other iterations into a `map` expression which allows computing them in parallel. Rules 2 and 3 group accesses to the same index of an array or the same key in an array of key-value pairs. Accesses to different indices/keys can then be computed in parallel. Rule 4 fuses two consecutive applications of `map` into a single application. Rule 5 deals with loops that update an array that they also read from by writing to a separate array instead. Rule 6 flattens nested `fold` expressions over an array of arrays into a single `fold` expression over the concatenated arrays. Rules 7 and 8 transform specific kinds of `iter` to `fold` and `fold` to `map`. Rules 9 and 10 transform `fold` and `map` expressions over an index range to `fold` and `map` expressions that operate directly on the values stored at these indices. Rule 11 commutes writing back updates to an array with applying a function to all result values. Finally, Rules 12 and 13 commute array reads with `zip` and `map`.

### 1. Extract independent part of loop body to map.

$$\begin{array}{ccc} \text{fold}(\lambda acc. \lambda x. f(acc, g(x)), & & \text{fold}(\lambda acc. \lambda y. f(acc, y), \\ acc, & \rightsquigarrow & acc, \\ xs) & & \text{map}(g, xs)) \end{array}$$

**Side conditions:**  $acc \notin FV(f)$ ,  $x \notin FV(f)$ ,  $y \notin FV(f)$ ,  
 $x \notin FV(g)$ ,  $acc \notin FV(g)$ ,  $g$  is not stuck

### 2. Group accesses to the same index of an array.

$$\begin{array}{ccc} \text{fold}(\lambda acc. \lambda(i, x). & & \text{fold}(\lambda acc. \lambda(i, v). \text{write}(acc, i, v), \\ \text{write}(acc, i, f(i, x, acc[i])), & \rightsquigarrow & ys, \\ ys, & & \text{map}(\lambda(i, vs). \\ xs) & & (i, \text{fold}(\lambda x'. \lambda x. f(i, x, x'), \\ & & \text{read}(ys, i), \\ & & vs)), \\ & & \text{group}(xs))) \end{array}$$

**Side conditions:**

$acc \notin FV(f)$ ,  $x \notin FV(f)$ ,  $x' \notin FV(f)$ ,  $i \notin FV(f)$ ,  $vs \notin FV(f)$

### 3. Group accesses to the same key.

$$\begin{array}{ccc} \text{fold}(\lambda acc. \lambda(k, v). & & \text{fold}(\lambda acc. \lambda(k, v). \\ \text{writeAtKey}(acc, & & \text{writeAtKey}(acc, k, v), \\ k, & & m, \\ f(k, & \rightsquigarrow & \text{map}(\lambda(k, vs). \\ v, & & (k, \text{fold}(\lambda v'. \lambda v. f(k, v, v'), \\ \text{readAtKey}(acc, k))), & & \text{readAtKey}(m, k), \\ m, & & vs)), \\ xs) & & \text{group}(xs))) \end{array}$$

**Side conditions:**

$acc \notin FV(f)$ ,  $v \notin FV(f)$ ,  $v' \notin FV(f)$ ,  $k \notin FV(f)$ ,  $vs \notin FV(f)$



**4. Fuse consecutive calls to map into a single call of map.**

$$\text{map}(f, \text{map}(g, xs)) \iff \text{map}(\lambda x. \text{app}(f, \text{app}(g, x)), xs)$$

**Side conditions:**  $x \notin FV(f), x \notin FV(g), f$  is not stuck,  $g$  is not stuck

**5. Separate arrays that are read from and written to.**

$$\begin{array}{ccc} \text{fold}(\lambda(xs', i). & & \text{fold}(\lambda(ys', i). \\ \text{write}(xs', & & \text{write}(ys', \\ i, & & i, \\ \text{app}(f, & \iff & \text{app}(f, \\ (i, \text{read}(xs', i))), & & (i, \text{read}(xs, i))), \\ xs, & & ys \\ \text{range}(0, \text{length}(xs))) & & \text{range}(0, \text{length}(xs))) \end{array}$$

**Side conditions:**

$\text{length}(xs) = \text{length}(ys), xs' \notin FV(f), i \notin FV(f), ys' \notin FV(xs), i \notin FV(xs)$

**6. Flatten fold over array of arrays.**

$$\text{fold}(\lambda(acc', xs). \text{fold}(f, acc', xs), acc, xss) \iff \text{fold}(f, acc, \text{concat}(xss))$$

**Side conditions:**  $f$  is not stuck,  $acc \notin FV(f), xs \notin FV(f)$

**7. Transform iter to fold.**

$$\begin{array}{ccc} \text{fst}(\text{iter}(\lambda(acc, i). \text{if } i < \text{max} & & \iff \text{fold}(f, \\ \text{then } \text{inr}(\text{app}(f, (acc, i)), & & \text{acc}_0, \\ i + 1) & & \text{range}(\text{min}, \text{max})) \\ \text{else } \text{inl}(\text{unit}), & & \\ (acc_0, \text{min})) \end{array}$$

**Side conditions:**

$f$  is not stuck,  $acc \notin FV(f), i \notin FV(f), i \notin FV(\text{max}), acc \notin FV(\text{max})$

**8. Transform fold to map.**

$$\begin{array}{ccc} \text{fold}(\lambda(ys', i). \text{twrite}(ys', & & \\ i, & & \\ \text{app}(f, \text{read}(xs, i))), & \iff & \text{map}(f, xs) \\ ys, & & \\ \text{range}(0, \text{length}(xs))) & & \end{array}$$

**Side conditions:**

$ys' \notin FV(f), i \notin FV(f), ys' \notin FV(xs), i \notin FV(xs), f$  is not stuck

**9. fold over the values in an array instead of over the index range.**

$$\begin{array}{ccc} \text{fold}(\lambda(acc, i). \text{app}(f, (acc, \text{read}(xs, i))), & & \\ i, & \iff & \text{fold}(f, i, xs) \\ \text{range}(i, \text{length}(xs))) & & \end{array}$$

**Side conditions:**

$i \notin FV(xs), acc \notin FV(xs), i \notin FV(f), acc \notin FV(f), f$  is not stuck

**10. map over the values in an array instead of over the index range.**

$$\begin{array}{ccc} \text{map}(\lambda i. \text{app}(f, \text{read}(xs, i)), & & \\ \text{range}(i, \text{length}(xs))) & \iff & \text{map}(f, xs) \end{array}$$

**Side conditions:**  $i \notin FV(xs), i \notin FV(f), f$  is not stuck

**11. Commute writing back updates to an array and applying map to the result.**

$$\text{map}(f, \text{fold}(\lambda(xs', (i, x)). \text{write}(xs', i, x), xs, ys)) \rightsquigarrow \text{fold}(\lambda(xs', (i, x)). \text{write}(xs', i, x), \text{map}(f, xs), \text{map}(\lambda(i, x). (i, \text{app}(f, x)), ys))$$

**Side conditions:**  $i \notin FV(f), x \notin FV(f)$

**12. Commute read and zip.**

a)  $\text{fst}(\text{read}(\text{zip}(xs, ys), i)) \rightsquigarrow \text{read}(xs, i)$

**Side conditions:**  $\text{length}(xs) = \text{length}(ys), ys$  is not stuck

b)  $\text{snd}(\text{read}(\text{zip}(xs, ys), i)) \rightsquigarrow \text{read}(ys, i)$

**Side conditions:**  $\text{length}(xs) = \text{length}(ys), xs$  is not stuck

**13. Commute read and map.**

a)  $\text{read}(\text{map}(f, xs), i) \rightsquigarrow \text{app}(f, \text{read}(xs, i))$

b)  $\text{readAtKey}(\text{map}(\lambda(k, v). (k, \text{app}(f, v)), xs), k) \rightsquigarrow \text{app}(f, \text{readAtKey}(xs, k))$