Tilting at windmills with Coq: formal verification of a compilation algorithm for parallel moves

Laurence Rideau, Bernard P. Serpette, Xavier Leroy

► To cite this version:

Laurence Rideau, Bernard P. Serpette, Xavier Leroy. Tilting at windmills with Coq: formal verification of a compilation algorithm for parallel moves. *Journal of Automated Reasoning*, 2008, 40 (4), pp.307-326. 10.1007/s10817-007-9096-8 . inria-00289709

HAL Id: inria-00289709

<https://hal.inria.fr/inria-00289709>

Submitted on 23 Jun 2008

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Tilting at windmills with Coq: formal verification of a compilation algorithm for parallel moves

Laurence Rideau · Bernard Paul Serpette ·
Xavier Leroy

the date of receipt and acceptance should be inserted later

Abstract This article describes the formal verification of a compilation algorithm that transforms parallel moves (parallel assignments between variables) into a semantically-equivalent sequence of elementary moves. Two different specifications of the algorithm are given: an inductive specification and a functional one, each with its correctness proofs. A functional program can then be extracted and integrated in the Compcert verified compiler.

Keywords Parallel move · Parallel assignment · Compilation · Compiler correctness · The Coq proof assistant

1 Introduction

Parallel assignment is a programming construct found in some programming languages (Algol 68, Common Lisp) and in some compiler intermediate languages (the RTL intermediate language of the GNU Compiler Collection). A parallel assignment is written $(x_1, \dots, x_n) := (e_1, \dots, e_n)$, where the x_i are variables and the e_i are expressions possibly involving the variables x_i . The semantics of this parallel assignment is to evaluate the expressions $e_1, \dots, e_n$, then assign their respective values to the variables $x_1, \dots, x_n$. Since the left-hand side variables can occur in the right-hand side expressions, the effect of a parallel assignment is, in general, different from that of the sequence of elementary assignments $x_1 := e_1; \dots; x_n := e_n$. Welch [15] gives examples of uses of parallel assignments.

Compiling a parallel assignment instruction amounts to finding a *serialisation* – a sequence of single assignments $x'_1 = e'_1; \dots; x'_m = e'_m$ – whose effect on the variables x_i and on other program variables is the same as that of the source-level parallel assignment. Since parallel assignment includes permutation of variables $(x_1, x_2) =$

L. Rideau and B. P. Serpette
INRIA Sophia-Antipolis Méditerranée, B.P. 93, 06902 Sophia-Antipolis, France
E-mail: Laurence.Rideau@inria.fr, Bernard.Serpette@inria.fr

X. Leroy
INRIA Paris-Rocquencourt, B.P. 105, 78153 Le Chesnay, France
E-mail: Xavier.Leroy@inria.fr

(x_2, x_1) as a special case, the generated sequence of elementary assignments cannot, in general, operate in-place over the x_i : additional storage is necessary under the form of temporary variables. A trivial compilation algorithm, outlined by Welch [15], uses n temporaries $t_1, \dots, t_n$, not present in the original program:

$$t_1 := e_1; \dots; t_n := e_n; x_1 := t_1; \dots; x_n := t_n$$

As shown by the example $(x_1, x_2) = (x_2, x_1)$, it is possible to find more efficient sequences than that use fewer than n temporaries.

Finding such sequences is difficult in general: Sethi [14] shows that compiling a parallel assignment over n variables using at most K temporaries, where K is a constant independent of n , is an NP-hard problem. However, there exists a family of parallel assignments for which serialisation is computationally easy: parallel assignments of the form $(x_1, \dots, x_n) := (y_1, \dots, y_n)$, where the right-hand sides y_i are restricted to be variables (either the x_i variables or other variables). We call these assignments *parallel moves*. They occur naturally in compilers when enforcing calling conventions [1]. May [13] shows another example of use of parallel moves in the context of machine code translation. It is a folklore result that parallel moves can be serialised using at most one temporary register, and in linear time [13].

The purpose of this paper is to formalize a compilation algorithm for parallel moves and mechanically verify its semantic correctness, using the Coq proof assistant [7,4]. This work is part of a larger effort, the CompCert project [10,6,5], which aims at mechanically verifying the correctness of an optimizing compiler for a subset of the C programming language. Every pass of the CompCert compiler is specified in Coq, then proved to be semantics-preserving: the observable behavior of the generated code is identical to that of the source code.

The parallel move compilation algorithm is used in the pass that enforces calling conventions for functions. Consider a source-level function call $f(e_1, \dots, e_n)$. Earlier compiler passes produce intermediate code that computes the values of the arguments $e_1, \dots, e_n$ and deposits them in locations $l_1, \dots, l_n$. (Locations are either processor registers or stack slots.) The function calling conventions dictate that the separately-compiled function f expects its parameters in conventional locations $l'_1, \dots, l'_n$ determined by the number and types of the parameters. The register allocation phase is instructed to prefer the locations $l'_1, \dots, l'_n$ for the targets of $e_1, \dots, e_n$, but such preferences cannot always be honored. Therefore, the pass that enforces calling conventions must, before the call, insert elementary move instructions that perform the parallel move $(l'_1, \dots, l'_n) := (l_1, \dots, l_n)$. Only two hardware registers (one integer register, one floating-point register) are available at this point to serve as temporaries. Therefore, the naive compilation algorithm for parallel moves will not do, and we had to implement and prove correct the space-efficient algorithm.

The formal specification and correctness proof of this compilation algorithm is challenging. Indeed, this was one of the most difficult parts of the whole CompCert development. While short and conceptually clear, the algorithm we started with (described section 3) is very imperative in nature. The underlying graph-like data structure, called *windmills* in this paper, is unusual. Finally, the algorithm involves non-structural recursion. We show how to tackle these difficulties by progressive refinement from a high-level, nondeterministic, relational specification of semantics-preserving transformations over windmills.

The remainder of this paper is organized as follows. Section 2 defines the windmill structure. Section 3 shows the folklore, imperative serialisation algorithm that we

used as a starting point. Two inductive specifications of the algorithm follow, a non-deterministic one (section 4), and a deterministic one (section 5), with the proofs that they are correct and consistent. Section 6 derives a functional implementation from these specifications and proves consistency with the inductive specifications, termination, and correctness. Section 7 proves additional syntactic properties of the result of the compilation function. In section 8, the correctness result of section 6 is extended to the case where variables can partially overlap. Concluding remarks are presented in section 9.

The complete, commented source for the Coq development presented here is available at <http://gallium.inria.fr/~xleroy/parallel-move/>.

2 Definitions and notations

An elementary move is written $(s \mapsto d)$, where the source register s and the destination register d range over a given set $\mathcal{R}$ of registers. Parallel moves as well as sequences of elementary moves are written as lists of moves $(s_1 \mapsto d_1) \cdots (s_n \mapsto d_n)$. The $\cdot$ operator denotes the concatenation of two lists. We overload it to also denote prepending a move in front of a list, $(s \mapsto d) \cdot l$, and appending a move at the end of a list, $l \cdot (s \mapsto d)$. The empty list is written $\emptyset$.

We assume given a set $\mathcal{T} \subseteq \mathcal{R}$ of temporary registers: registers that are not mentioned in the initial parallel move problem and that the compiler can use to break cycles. We also assume given a function $T : \mathcal{R} \rightarrow \mathcal{T}$ that associates to any register r a temporary register $T(r)$ appropriate for saving the value of r when breaking a cycle involving r . In the simplest case, only one temporary `tmp` is available and T is the constant function $T(r) = \text{tmp}$. In more realistic cases, architectural or typing constraints may demand the use of several temporaries, for instance one integer temporary to move integer or pointer values, and one floating-point temporary to move floating-point values. In this example, $T(r)$ selects the appropriate temporary as a function of the type of r .

A parallel move $(s_1 \mapsto d_1) \cdots (s_n \mapsto d_n)$ is well defined only if the destination registers are pairwise distinct: $d_i \neq d_j$ if $i \neq j$. (If a register appeared twice as a destination, its final value after the parallel move would not be uniquely defined.) We call such parallel moves *windmills*.

Definition 1 (Windmills) A parallel move μ is a windmill if, for all $l_1, s_i, d_i, l_2, s_j, d_j$ and l_3 ,

$$\mu = l_1 \cdot (s_i \mapsto d_i) \cdot l_2 \cdot (s_j \mapsto d_j) \cdot l_3 \Rightarrow d_i \neq d_j$$

The name “windmill” comes from the graphical shape of the corresponding transfer relation. We can view a parallel move as a transfer relation. Each move $(s_i \mapsto d_i)$ corresponds to an edge in the graph of this relation. A parallel move is a windmill if every register has at most one predecessor for the transfer relation. Although this property is similar to the specification of forests (set of disjoint trees), it allows cycles such as $(r_1 \mapsto r_2) \cdot (r_2 \mapsto r_1)$. This is unlike a tree, where, by definition, there exists a unique element – the root – that has no predecessor.

The graph of the relation corresponding to a windmill is composed of cycles – the *axes* – whose elements are the roots of trees – the *blades*. Figure 1 figure shows a set of 4 windmills: the general case, the special case of a tree, a simple cycle with four registers, and the special case of a self-loop.

Fig. 1 Examples of windmills

3 An imperative algorithm

There are two special cases of parallel move problems where serialization is straightforward. In the case of a simple cycle, i.e. an axle without any blade, the transfer relation is:

$$(r_1 \mapsto r_2) \cdot (r_2 \mapsto r_3) \cdots (r_{n-1} \mapsto r_n) \cdot (r_n \mapsto r_1)$$

Serialization is done using a single temporary register $t = T(r_1)$:

$$t := r_1; r_1 := r_n; r_n := r_{n-1}; \dots; r_3 := r_2; r_2 := t.$$

The other easy case corresponds to a transfer relation that is a tree, such as for instance

$$(r_1 \mapsto r_2) \cdot (r_1 \mapsto r_3) \cdot (r_2 \mapsto r_4)$$

In this case, serialization corresponds to enumerating the edges in a bottom-up topological order:

$$r_4 := r_2; r_2 := r_1; r_3 := r_1.$$

C. May [13] describes an algorithm that generalizes these two special cases. This algorithm follows the topology of the transfer relation. First, remove one by one all the edges that have no successors, emitting the corresponding sequential assignments in the same order. (In windmill terminology, this causes the blades to disappear little by little.) Eventually, all that remains are simple, disjoint cycles (windmill axles) which can be serialized using one temporary as described above.

We now consider a variant of this algorithm that processes blades and axles simultaneously, in a single pass. The algorithm is given in figure 2 in Caml syntax [11]. It can be read as pseudo-code knowing that `a.(i)` refers to the *i*-th element of array `a`.

It takes as arguments two arrays `src` and `dst` containing respectively the source registers $s_1, \dots, s_n$ and the target registers $d_1, \dots, d_n$, as well as the temporary-generating function `tmp`. The elementary moves produced by serialization are successively printed using the C-like `printf` function.

The algorithm implements a kind of depth-first graph traversal using the `move_one` function that takes an edge $(src_i \mapsto dst_i)$ of the transfer relation as an argument. In line 9 and 10, all successors of dst_i are handled. Line 17 handles the case of an already analysed edge. In the case of line 12, it is a simple recursive call. In the case of line 14, a cycle is discovered and a temporary register is used to break it.

```

1  type status = To_move | Being_moved | Moved
2
3  let parallel_move src dst tmp =
4    let n = Array.length src in
5    let status = Array.make n To_move in
6    let rec move_one i =
7      if src.(i) ≠ dst.(i) then begin
8        status.(i) ← Being_moved;
9        for j = 0 to n - 1 do
10         if src.(j) = dst.(i) then
11           match status.(j) with
12             | To_move →
13               move_one j
14             | Being_moved →
15               printf "%s␣:=␣%s;\n" (tmp src.(j)) src.(j);
16               src.(j) ← tmp src.(j)
17             | Moved →
18               ()
19         done;
20         printf "%s␣:=␣%s;\n" dst.(i) src.(i);
21         status.(i) ← Moved
22       end in
23     for i = 0 to n - 1 do
24       if status.(i) = To_move then move_one i
25     done

```

Fig. 2 A one-pass, imperative algorithm to compile parallel moves

When exiting this loop on line 19, all the nodes that can be reached by dst_i have been analysed, the edge $(src_i \mapsto dst_i)$ is serialized and marked as analyzed.

When a cycle is discovered, the stack of active calls to `move_one` correspond to the edges $((r_n \mapsto r_1), \dots, (r_2 \mapsto r_3), (r_1 \mapsto r_2))$ and, at line 14, the edge $(r_1 \mapsto r_2)$ is analysed again. The side effect of line 16 thus necessarily acts on the edge that started the recursion of the `move_one` function, i.e. the edge at the bottom of the stack. The main loop, lines 23 to 25, ensures that all edges are analyzed at least once.

4 Nondeterministic specification

The algorithm in figure 2 does not lend itself easily to program proof: first, it is written imperatively, making it quite removed from a mathematical specification; second, it commits to a particular strategy for solving the problem, obscuring the essential invariants for the correctness proof. In this section, we develop an abstract specification of the steps of the algorithm, and show that these steps preserve semantics. This specification is not deterministic: it does not constrain which step must be taken in a given state.

In the algorithm of figure 2, notice that every edge of the transfer relation begins in the state `To_move` (line 5), then takes the state `Being_moved` (line 8), and finally reaches the state `Moved` (line 21). Rather than associating a status to each edge, we will represent the current state of the compilation as three disjoint lists of edges, each list containing edges having the same status. The state is therefore a triple (μ, σ, τ) of lists of edges: μ is the *to-move* list, σ the *being-moved* list and τ the *moved* list.

Each step of the algorithm will either extract an element from the *to-move* list μ and add it to the *being-moved* list σ , or remove an element of the latter and add it to the *moved* list τ . The *to-move* and *being-moved* lists are used as work lists; the algorithm will stop when both are empty.

The *being-moved* list σ is used as a stack, simulating the recursive calls to the `move_one` function in the imperative algorithm. In addition to pushing and popping moves from the beginning of σ , the last element of σ can also be modified when a cycle is discovered.

The last component of the state, the *moved* list τ , is used to accumulate the sequence of elementary moves that the imperative algorithm emits using the `printf` function. Moves are successively added to the front of the list τ . Therefore, the elementary moves listed in τ are to be executed from right to left.

4.1 Inference rules

The inference rules below define a rewriting relation $\triangleright$ between states (μ, σ, τ) . Each rule describes a step that the compilation algorithm is allowed to take. A run of compilation is viewed as a sequence of rewrites from the initial state $(\mu, \emptyset, \emptyset)$ to a final state $(\emptyset, \emptyset, \tau)$, where μ is the parallel move problem we set out to compile, and `reverse`(τ) is the sequence of elementary moves generated by this compilation run.

$$\frac{}{(\mu_1 \cdot (r \mapsto r) \cdot \mu_2, \sigma, \tau) \triangleright (\mu_1 \cdot \mu_2, \sigma, \tau)} \text{ [Nop]}$$

The first rule deals with the case where the transfer relation has an edge $(r \mapsto r)$ whose source and destination are identical. This case is specially handled at line 7 of the imperative algorithm. Note that in the imperative algorithm, this edge is not annotated as `Moved` and keeps the status `To_move` until the end of the algorithm. Nevertheless this edge cannot result from a recursive call of the function `move_one` because, in this case, the caller would correspond to an edge $(s \mapsto r)$ and that would violate the fact that the transfer relation is a windmill: indeed, there would exist two different edges with the same destination. It is therefore valid to move the test of line 7 up to the level of the main loop at line 24.

$$\frac{}{(\mu_1 \cdot (s \mapsto d) \cdot \mu_2, \emptyset, \tau) \triangleright (\mu_1 \cdot \mu_2, (s \mapsto d), \tau)} \text{ [Start]}$$

The [Start] rule corresponds to the first call to the `move_one` function at line 24. We know that in this case no edge has the status `Being_Moved` and this rule will thus be the only one to handle the case where the *being-moved* list is empty.

$$\frac{}{(\mu_1 \cdot (d \mapsto r) \cdot \mu_2, (s \mapsto d) \cdot \sigma, \tau) \triangleright (\mu_1 \cdot \mu_2, (d \mapsto r) \cdot (s \mapsto d) \cdot \sigma, \tau)} \text{ [Push]}$$

The [Push] rule corresponds to the recursive call at line 13. If the edge being analysed is $(s \mapsto d)$ (i.e. the top of the *being-moved* list) and if there exists a successor $(d \mapsto r)$ in the *to-move* list, the latter is transferred to the top of the *being-moved* stack.

$$\frac{}{(\mu, \sigma \cdot (s \mapsto d), \tau) \triangleright (\mu, \sigma \cdot (T(s) \mapsto d), (s \mapsto T(s)) \cdot \tau)} \text{ [Loop]}$$

The [Loop] rule corresponds to the special case at lines 15 and 16 where a cycle is discovered. This rule is of more general use. Indeed, it is not essential to check the presence of a cycle before inserting a transfer using a temporary register. Thus, lines 15 and 16 can also be moved up to the level of the main loop (line 24). The fact that one uses the temporary register only in the case of a cycle must be regarded as an optimization. The crucial issue is to make sure that a transfer through a temporary register must be used whenever a cycle arises. This is ensured by the next rule.

$$\frac{\text{NoRead}(\mu, d_n) \wedge d_n \neq s_0}{(\mu, (s_n \mapsto d_n) \cdot \sigma \cdot (s_0 \mapsto d_0), \tau) \triangleright (\mu, \sigma \cdot (s_0 \mapsto d_0), (s_n \mapsto d_n) \cdot \tau)} \quad [\text{Pop}]$$

The [Pop] rule corresponds to returning from a recursive call to the function `move_one`. The first premise of this rule, `NoRead`(μ, d_n), is formally defined as

$$\text{NoRead}(\mu, d_n) \stackrel{\text{def}}{=} \forall \mu_1, s, \mu_2, \quad \mu = \mu_1 \cdot (s \mapsto d) \cdot \mu_2 \Rightarrow s \neq d_n.$$

This premise checks that the edge $(s_n \mapsto d_n)$ under study no longer has its successor d_n in the *to-move* list. It therefore prevents the [Pop] rule from being used if the value of d_n is still needed. The second premise, $d_n \neq s_0$, makes sure that the [Pop] rule cannot be used if d_n participates in a cycle, forcing the use of the [Loop] rule instead in the case of a cycle.

$$\frac{\text{NoRead}(\mu, d_n)}{(\mu, (s \mapsto d), \tau) \triangleright (\mu, \emptyset, (s \mapsto d) \cdot \tau)} \quad [\text{Last}]$$

The [Last] rule corresponds to returning from the main loop at line 24. It is a special case of the [Pop] rule. As in the special case of the [Nop] rule, observe that a *to-move* list of the form $((r \mapsto r))$ can either be eliminated by the [Nop] rule, or rewritten to itself in two steps ([Start] then [Last]), or rewritten to $((T(r) \mapsto r) \cdot (r \mapsto T(r)))$ using a temporary register (rules [Start] then [Loop] then [Last]). These three results are equally valid.

4.2 Well-formedness invariant

In preparation for proving the semantic correctness of the rewriting rules, we first define a well-formedness property that acts as a crucial invariant in this proof.

Definition 2 (Invariant) A triple (μ, σ, τ) is well-formed, written $\vdash (\mu, \sigma, \tau)$, if and only if:

1. $\mu \cdot \sigma$ is a windmill, in the sense of definition 1.
2. The μ list does not contain temporary registers: for all μ_1, s, d, μ_2 ,

$$\mu = \mu_1 \cdot (s \mapsto d) \cdot \mu_2 \Rightarrow (s \notin \mathcal{T} \wedge d \notin \mathcal{T}).$$

3. The σ list can only use a temporary register as the source of its last edge: for all σ_1, s_0, d_0 ,

$$\sigma = \sigma_1 \cdot (s_0 \mapsto d_0) \Rightarrow d_0 \notin \mathcal{T}$$

and for all σ_2, s, d, σ_3 ,

$$\sigma_1 = \sigma_2 \cdot (s \mapsto d) \cdot \sigma_3 \wedge \sigma_3 \neq \emptyset \Rightarrow (s \notin \mathcal{T} \wedge d \notin \mathcal{T}).$$

4. The σ list is a path: $\sigma = (r_{n-1} \mapsto r_n) \cdots (r_2 \mapsto r_3) \cdot (r_1 \mapsto r_2)$ for some registers $r_1, \dots, r_n$.

Notice that if μ is a windmill and does not use any temporary register, then the initial triple $(\mu, \emptyset, \emptyset)$ is well-formed.

Lemma 1 (The invariant is preserved) *The rewriting rules transform a well-formed triple into a well-formed triple. If $S_1 \triangleright S_2$ and $\vdash S_1$, then $\vdash S_2$.*

Proof By case analysis on the rule that concludes $S_1 \triangleright S_2$. For $S_1 = (\mu, \sigma, \tau)$, only the [Push] rule adds an element to σ , but the edge thus added extends the path already present in σ .

4.3 Dynamic semantics for assignments

To state semantic preservation for the rewriting relation, we need to give dynamic semantics to the parallel moves μ and the sequential moves σ and τ occurring in intermediate states. Let $\mathcal{V}$ be the set of run-time values for the language. The run-time state of a program is represented as an environment $\rho : \mathcal{R} \rightarrow \mathcal{V}$ mapping registers to their current values. Assigning value v to register r transforms the environment ρ into the environment $\rho[r \leftarrow v]$ defined by

$$\rho[r \leftarrow v] = \begin{cases} r \mapsto v \\ r' \mapsto \rho(r') & \text{if } r' \neq r. \end{cases}$$

(We assume that assigning to a register r preserves the values of all other registers r' , or in other words that r does not overlap with any other register. We will revisit this hypothesis in section 8.)

The run-time effect of a parallel move μ is to transform the current environment ρ in the environment $\llbracket \mu \rrbracket_{//}(\rho)$ defined as:

$$\llbracket (s_n \mapsto d_n) \cdots (s_0 \mapsto d_0) \rrbracket_{//}(\rho) = \rho[d_0 \leftarrow \rho(s_0)] \cdots [d_n \leftarrow \rho(s_n)]$$

In this definition parentheses are omitted, as we require that μ is a windmill. The following lemma, showing that the order of moves within μ does not matter, is intensively used in our proofs:

Lemma 2 (Independence) *If $\mu = \mu_1 \cdot (s \mapsto d) \cdot \mu_2$ is a windmill, then*

$$\llbracket \mu \rrbracket_{//}(\rho) = \llbracket (s \mapsto d) \cdot \mu_1 \cdot \mu_2 \rrbracket_{//}(\rho)$$

Proof By induction on μ_1 . The base case $\mu_1 = \emptyset$ is trivial. For the inductive step, we need to prove that the first two elements of μ_1 can be exchanged: $(\rho[d_0 \leftarrow \rho(s_0)])[d_1 \leftarrow \rho(s_1)] = (\rho[d_1 \leftarrow \rho(s_1)])[d_0 \leftarrow \rho(s_0)]$. The windmill structure of μ_1 implies that $d_0 \neq d_1$, and therefore this equality is verified.

The run-time effect of a sequence of elementary moves τ , performed from left to right, is defined by:

$$\begin{aligned} \llbracket \emptyset \rrbracket_{\rightarrow}(\rho) &= \rho \\ \llbracket (s \mapsto d) \cdot \tau \rrbracket_{\rightarrow}(\rho) &= \llbracket \tau \rrbracket_{\rightarrow}(\rho[d \leftarrow \rho(s)]) \end{aligned}$$

The sequences of elementary moves τ appearing in the states of the rewriting relation are actually to be interpreted from right to left, or equivalently to be reversed before execution. It is therefore useful to define directly the semantics of a sequence of elementary moves τ performed from right to left:

$$\begin{aligned} \llbracket \emptyset \rrbracket \leftarrow (\rho) &= \rho \\ \llbracket (s \mapsto d) \cdot \tau \rrbracket \leftarrow (\rho) &= \rho' [d \leftarrow \rho'(s)] \text{ where } \rho' = \llbracket \tau \rrbracket \leftarrow (\rho) \end{aligned}$$

Lemma 3 (Reversed sequential execution) *For all sequences τ of elementary moves, $\llbracket \tau \rrbracket \leftarrow (\rho) = \llbracket \text{reverse}(\tau) \rrbracket \rightarrow (\rho)$.*

Proof Easy induction over τ .

Definition 3 (Equivalent environments) We say that two environments ρ_1 and ρ_2 are equivalent, and we write $\rho_1 \equiv \rho_2$, if and only if all non-temporary registers have the same values in both environments:

$$\rho_1 \equiv \rho_2 \stackrel{\text{def}}{=} \forall r \notin \mathcal{T}, \rho_1(r) = \rho_2(r).$$

Definition 4 (Semantics of a triple) The dynamic semantics of the triple (μ, σ, τ) corresponds to first executing τ sequentially from right to left, then executing $\mu \cdot \sigma$ in parallel:

$$\llbracket (\mu, \sigma, \tau) \rrbracket (\rho) = \llbracket \mu \cdot \sigma \rrbracket // (\llbracket \tau \rrbracket \leftarrow (\rho)).$$

Notice that for initial states, we have $\llbracket (\mu, \emptyset, \emptyset) \rrbracket (\rho) = \llbracket \mu \rrbracket // (\rho)$, while for final states, we have $\llbracket (\emptyset, \emptyset, \tau) \rrbracket (\rho) = \llbracket \tau \rrbracket \leftarrow (\rho) = \llbracket \text{reverse}(\tau) \rrbracket \rightarrow (\rho)$.

4.4 Correctness

Lemma 4 (One-step semantic preservation) *The rewriting rules preserve the semantics of well-formed triples: if $S_1 \triangleright S_2$ and $\vdash S_1$, then $\llbracket S_1 \rrbracket (\rho) \equiv \llbracket S_2 \rrbracket (\rho)$ for all environments ρ .*

Proof By case analysis on the rule used to derive $S_1 \triangleright S_2$. We write $S_1 = (\mu, \sigma, \tau)$.

- Rule [Nop]: use the fact that $\rho[r \leftarrow \rho(r)] = \rho$.
- Rules [Start] and [Push]: use the order independence lemma 2.
- Rule [Loop]: exploit the fact that μ and σ do not use the temporary register introduced by the rule.
- Rule [Pop]: the expected result follows from

$$\llbracket (s_n \mapsto d_n) \cdot \mu \cdot \sigma \rrbracket // (\llbracket \tau \rrbracket \leftarrow (\rho)) = \llbracket \mu \cdot \sigma \rrbracket // (\llbracket (s_n \mapsto d_n) \cdot \tau \rrbracket \leftarrow (\rho))$$

This equality holds provided that register d_n does not appear as a source in $\mu \cdot \sigma$. For μ , this is ensured by the first premise. For σ , this follows by induction from the fact that σ is a path. The base case comes from the premise $d_n \neq s_0$. The inductive case is provided by the windmill structure of σ .

- Rule [Last]: special case of rule [Pop].

This lemma is easily extended to $\triangleright^*$, the reflexive transitive closure of $\triangleright$:

Lemma 5 (Semantic preservation) *Several steps of rewriting preserve the semantics of well-formed triples: if $S_1 \triangleright^* S_2$ and $\vdash S_1$, then $\llbracket S_1 \rrbracket(\rho) \equiv \llbracket S_2 \rrbracket(\rho)$ for all environments ρ .*

Proof By induction on the number of steps.

As a corollary, we obtain the main semantic preservation theorem of this section.

Theorem 1 (Correctness of $\triangleright^*$) *If μ is a windmill containing no temporary register, and if the triple $(\mu, \emptyset, \emptyset)$ can be rewritten into $(\emptyset, \emptyset, \tau)$ then the parallel execution of μ is equivalent to the sequential execution of $\text{reverse}(\tau)$:*

$$(\mu, \emptyset, \emptyset) \triangleright^* (\emptyset, \emptyset, \tau) \Rightarrow \llbracket \mu \rrbracket_{//}(\rho) \equiv \llbracket \text{reverse}(\tau) \rrbracket \rightarrow (\rho).$$

5 Deterministic specification

The next step towards an effective algorithm consist in *determinising* the inductive rules of section 4 by ensuring that at most one rule matches a given state, and by specifying which edge to extract out of the *to-move* list in the rules [Nop], [Push] and [Start].

$$\frac{}{((r \mapsto r) \cdot \mu, \emptyset, \tau) \hookrightarrow (\mu, \emptyset, \tau)} \text{ [Nop']}$$

$$\frac{s \neq d}{((s \mapsto d) \cdot \mu, \emptyset, \tau) \hookrightarrow (\mu, (s \mapsto d), \tau)} \text{ [Start']}$$

The first two rules treat the case of an empty *being-moved* list and consider only the first element on the left of the *to-move* list μ .

$$\frac{\text{NoRead}(\mu_1, d)}{(\mu_1 \cdot (d \mapsto r) \cdot \mu_2, (s \mapsto d) \cdot \sigma, \tau) \hookrightarrow (\mu_1 \cdot \mu_2, (d \mapsto r) \cdot (s \mapsto d) \cdot \sigma, \tau)} \text{ [Push']}$$

The [Push'] rule treats the case in which the analysed edge has a successor in the μ list. It forces the algorithm to consider the leftmost successor found in μ .

$$\frac{\text{NoRead}(\mu, r_0)}{(\mu, (s \mapsto r_0) \cdot \sigma \cdot (r_0 \mapsto d), \tau) \hookrightarrow (\mu, \sigma \cdot (T(r_0) \mapsto d), (s \mapsto r_0) \cdot (r_0 \mapsto T(r_0)) \cdot \tau)} \text{ [LoopPop]}$$

$$\frac{\text{NoRead}(\mu, d_n) \wedge d_n \neq s_0}{(\mu, (s_n \mapsto d_n) \cdot \sigma \cdot (s_0 \mapsto d_0), \tau) \hookrightarrow (\mu, \sigma \cdot (s_0 \mapsto d_0), (s_n \mapsto d_n) \cdot \tau)} \text{ [Pop']}$$

These two rules treat the case where the analysed edge has no successor in the *to-move* list, and the *being-moved* list contains at least two elements. The [LoopPop] rule is the [Pop] rule in the case of a cycle. This rule, unlike the corresponding nondeterministic rules, pops the analysed edge. This ensures that every rule moves at least one edge. The termination proof in section 6 uses this fact.

$$\frac{\text{NoRead}(\mu, d)}{(\mu, (s \mapsto d), \tau) \hookrightarrow (\mu, \emptyset, (s \mapsto d) \cdot \tau)} \quad [\text{Last}']$$

The $[\text{Last}']$ rule treats the final case where the *being-moved* list contains only one element.

Since these deterministic rules are instances or combinations of the nondeterministic ones, the following two lemmas are trivial.

Lemma 6 (Inclusion $\hookrightarrow \subseteq \triangleright^*$) *If $S_1 \hookrightarrow S_2$, then $S_1 \triangleright^* S_2$.*

Proof By case analysis. The $[\text{LoopPop}]$ case uses the transitivity of $\triangleright^*$.

Lemma 7 (Inclusion $\hookrightarrow^* \subseteq \triangleright^*$) *If $S_1 \hookrightarrow^* S_2$, then $S_1 \triangleright^* S_2$.*

Proof By induction on the number of steps.

Notice that the deterministic rules implement a strategy that is slightly different from that of the imperative algorithm. Indeed, during the imperative algorithm, when a temporary register is used to break a cycle, the analysed edge can still have a successor. This is not the case for the $[\text{LoopPop}]$ rule.

Semantic preservation for the deterministic rules follows immediately from theorem 1 and lemma 7. However, to obtain an algorithm, we still need to prove that the rules are indeed deterministic and that a normal form (i.e. the final state $(\emptyset, \emptyset, \tau)$) always exists. To do so, we now give a functional presentation of the deterministic rules.

6 The functional algorithm

From the deterministic rules in section 5, it is easy to build a function **stepf** that performs one step of rewriting. Compiling a parallel move problem, then, corresponds to iterating **stepf** until a final state is reached:

```
pmov S =
  match S with
  | (∅, ∅, _) ⇒ S
  | _ ⇒ pmov (stepf S)
end
```

However, such a function definition is not accepted by Coq, because it is not structurally recursive: the recursive call is performed on **stepf**(S) and not on a structural subterm of S . To define the **pmov** function, we will use the Coq command **Function**, which provides a simple mechanism to define a recursive function by well-founded recursion [2,3]. To this end, we first define the **stepf** function that executes one computation step, then a measure function which will be used to justify termination, and finally the theorem establishing that the measure decreases at each recursive call.

```

Function stepf (st: state) : state :=
  match st with
  | State  $\emptyset \emptyset \_ \Rightarrow$  st                                (* final state *)
  | State ((s, d) · t1)  $\emptyset$  1  $\Rightarrow$ 
    if reg_eq s d
    then State t1  $\emptyset$  1                                (* s = d; rule [Nop] *)
    else State t1 ((s, d) ·  $\emptyset$ ) 1                    (* s  $\neq$  d; rule [Start] *)
  | State t ((s, d) · b) 1  $\Rightarrow$ 
    match split_move t d with
    | Some (t1, r, t2)  $\Rightarrow$                                (* t = t1 · (d  $\mapsto$  r) · t2 *)
      State (t1 · t2)                                     (* rule [Push] *)
      ((d, r) · (s, d) · b)
      1
    | None  $\Rightarrow$ 
      match b with
      |  $\emptyset \Rightarrow$  State t  $\emptyset$  ((s, d) · 1)            (* rule [Last] *)
      | _  $\Rightarrow$ 
        if is_last_source d b                             (* check if b = _ · (d  $\mapsto$  _) *)
        then State t                                       (* rule [LoopPop] *)
        (replace_last_source (temp d) b)
        ((s, d) · (d, temp d) · 1)
        else State t b ((s, d) · 1)                       (* rule [Pop] *)
      end
    end
  end.

```

Fig. 3 The one-step rewriting function

6.1 The one-step rewriting function

We define the one-step rewriting function **stepf** in a way that closely follows the deterministic rules of section 5. It takes a state S as argument and returns a state S' such that $S \hookrightarrow S'$. The function, written in Coq syntax, is defined by case analysis in figure 3.

The first case (line 3) does not correspond to any of the inductive rules: it detects a final state $(\emptyset, \emptyset, \tau)$ and returns it unchanged. The other cases correspond to the deterministic rules indicated in comments.

The **reg_eq** function decides whether two registers are equal or not. The **split_move** function takes as arguments a windmill t and a register d . If this register is the source of an edge from list t , the function returns this edge $(d, _)$ and the two remaining sub-lists. This breaks the t list into $t_1 \cdot (d, r) \cdot t_2$ like the [Push] rule does.

Finally, the **replace_last_source** function takes as argument a list b of edges and a register r , and replaces the source of the last edge of b with r : $hb \cdot (s \mapsto d)$ becomes $hb \cdot (t \mapsto d)$, in accordance with the [LoopPop] rule.

Since the definition of the **stepf** function is close to the deterministic specification, the following theorem is easily proved:

Lemma 8 (Compatibility of stepf and $\hookrightarrow$) $S \hookrightarrow \text{stepf}(S)$ if S is not a final state (i.e. not of the form $(\emptyset, \emptyset, \tau)$).

Proof By case analysis on the shape of S and application of the corresponding rules for the $\hookrightarrow$ relation.

6.2 The general recursive function

The general recursive function consists in iterating the **stepf** function. To ensure termination of this function, we need a nonnegative integer measure over states that decreases at each call to the **stepf** function, except in the final case $(\emptyset, \emptyset, \tau)$. Examination of the deterministic rewriting rules and of the lengths of the *to-move* and *being-moved* lists reveals that:

- either an edge is removed from the *to-move* list and the *being-moved* list is unchanged (rule [Nop]);
- or an edge is transferred from the *to-move* list to the *being-moved* list (rules [Start] and [Push]);
- or the *to-move* list remains the same while an edge is transferred from the *being-moved* list to the *moved* list (rules [LoopPop], [Pop], and [Last]).

Therefore, either one of the *to-move* or *being-moved* lists loses one element, or one element is transferred from the former to the latter. Decrease is achieved by giving more weight to the *to-move* edges, as the following **meas** function does:

Definition 5 (The triple measure) $\text{meas}(\mu, \sigma, \tau) = 2 \times \text{length}(\mu) + \text{length}(\sigma)$.

Lemma 9 (Measure decreases with $\hookrightarrow$) *If $S_1 \hookrightarrow S_2$, then $\text{meas}(S_2) < \text{meas}(S_1)$.*

Proof By case analysis on the rules of the $\hookrightarrow$ relation and computation of the corresponding measures.

Combining this lemma with lemma 8 (compatibility of **stepf** with $\hookrightarrow$), we obtain:

Lemma 10 (Measure decreases with **stepf)** *If S is not a final state, i.e. not of the form $(\emptyset, \emptyset, \tau)$, then $\text{meas}(\text{stepf}(S)) < \text{meas}(S)$.*

We can then define the iterate of **stepf** using the **Function** command of Coq version 8:

```
Function pmov (S: State) {measure meas st} : state :=
  if final_state S then S else pmov (stepf S).
```

where **final_state** is a boolean-valued function returning **true** if its argument is of the form $(\emptyset, \emptyset, \tau)$ and **false** otherwise. Coq produces a proof obligation requiring the user to show that the recursive call to **pmov** is actually decreasing with respect to the **meas** function. This obligation is trivially proved by lemma 10. Coq then generates and automatically proves the following functional induction principle that enables us to reason over the **pmov** function:

$$\begin{aligned}
& (\forall S, \text{final_state}(S) = \text{true} \Rightarrow P(S, S)) \\
& \wedge (\forall S, \text{final_state}(S) = \text{false} \\
& \quad \wedge P(\text{stepf}(S), \text{pmov}(\text{stepf}(S))) \\
& \quad \Rightarrow P(S, \text{pmov}(\text{stepf}(S))) \\
& \Rightarrow (\forall S, P(S, \text{pmov}(S)))
\end{aligned}$$

Using this induction principle and lemma 8, we obtain the correctness of **pmov** with respect to the deterministic specification.

Lemma 11 *For all initial states S , $\text{pmov}(S)$ is a final state $(\emptyset, \emptyset, \tau)$ such that $S \hookrightarrow^* (\emptyset, \emptyset, \tau)$.*

6.3 The compilation function

To finish this development, we define the main compilation function, taking a parallel move problem as input and returning an equivalent sequence of elementary moves.

Definition `parmove (mu: moves) : moves :=`
`match pmov (State mu  $\emptyset$   $\emptyset$ ) with`
`| State _ _ tau  $\Rightarrow$  reverse tau`
`end.`

The semantic correctness of `parmove` follows from the previous results.

Theorem 2 *Let μ be a parallel move problem. If μ is a windmill and does not mention temporary registers, then $\llbracket \text{parmove}(\mu) \rrbracket \rightarrow (\rho) \equiv \llbracket \mu \rrbracket_{//}(\rho)$ for all environments ρ .*

Proof By definition of `parmove` and lemma 11, we have `parmove`(μ) = `reverse`(τ) and $(\mu, \emptyset, \emptyset) \hookrightarrow^* (\emptyset, \emptyset, \tau)$. By lemma 7, this entails $(\mu, \emptyset, \emptyset) \vdash^* (\emptyset, \emptyset, \tau)$. The result follows from theorem 1.

Here is an alternate formulation of this theorem that can be more convenient to use.

Theorem 3 *Let $\mu = (s_1 \mapsto d_1) \cdots (s_n \mapsto d_n)$ be a parallel move problem. Assume that the d_i are pairwise distinct, and that $s_i \notin \mathcal{T}$ and $d_i \notin \mathcal{T}$ for all i . Let $\tau = \text{parmove}(\mu)$. For all initial environments ρ , the environment $\rho' = \llbracket \tau \rrbracket \rightarrow (\rho)$ after executing τ sequentially is such that*

1. $\rho'(d_i) = \rho(s_i)$ for all $i = 1, \dots, n$;
2. $\rho'(r) = \rho(r)$ for all registers $r \notin \{d_1, \dots, d_n\} \cup \mathcal{T}$.

Proof By theorem 2, we know that $\rho' \equiv \llbracket \mu \rrbracket_{//}(\rho)$. For (1), since $d_i \notin \mathcal{T}$, we have $\rho'(d_i) = \llbracket \mu \rrbracket_{//}(\rho)(d_i) = \rho(s_i)$ by lemma 2. For (2), we have $\rho'(r) = \llbracket \mu \rrbracket_{//}(\rho)(r)$, and an easy induction on μ shows that $\llbracket \mu \rrbracket_{//}(\rho)(r) = \rho(r)$ for all registers r that are not destinations of μ .

Finally, we used the Coq extraction mechanism [12] to automatically generate executable Caml code from the definition of the `parmove` function, thus obtaining a verified implementation of the parallel move compilation algorithm that can be integrated in the CompCert compiler back-end.

7 Syntactic properties of the compilation algorithm

In this section, we show a useful syntactic property of the sequences of moves generated by the `parmove` compilation function: the registers involved in the generated moves are not arbitrary, but either appear in the initial parallel move problem, or are results of the temporary-generating function T .

More formally, let $\mu_0 = (s_1 \mapsto d_1) \cdots (s_n \mapsto d_n)$ be the initial parallel move problem. We define a property $P(s \mapsto d)$ of elementary moves $(s \mapsto d)$ by the following inference rules:

$$\frac{(s \mapsto d) \in \mu_0}{P(s \mapsto d)} \quad \frac{P(s \mapsto d)}{P(T(s) \mapsto d)} \quad \frac{P(s \mapsto d)}{P(s \mapsto T(s))}$$

We extend this property pointwise to lists l of moves and to states:

$$P(l) \stackrel{\text{def}}{=} \forall (s \mapsto d) \in l, P(s \mapsto d) \quad P(\mu, \sigma, \tau) \stackrel{\text{def}}{=} P(\mu) \wedge P(\sigma) \wedge P(\tau)$$

Obviously, the property P holds for μ_0 and therefore for the initial state $(\mu_0, \emptyset, \emptyset)$. Moreover, the property is preserved by every rewriting step.

Lemma 12 *If $S_1 \triangleright S_2$ and $P(S_1)$ hold, then $P(S_2)$ holds.*

Proof By examination of the nondeterministic rewriting rules. The result is obvious for all rules except [Loop], since these rules do not generate any new moves. The [Loop] rule replaces a move $(s \mapsto d)$ that satisfies P with the two moves $(T(s) \mapsto d)$ and $(s \mapsto T(s))$. The new moves satisfy P by application of the second and third inference rules defining P .

It follows that the result of $\text{parmove}(\mu_0)$ satisfies property P .

Lemma 13 *For all moves $(s \mapsto d) \in \text{parmove}(\mu_0)$, the property $P(s \mapsto d)$ holds.*

Proof Follows from lemmas 7, 11 and 12.

As a first use of this lemma, we can show that for every move in the generated sequence $\text{parmove}(\mu_0)$, the source of this move is either a source of μ_0 or a temporary, and similarly for the destination.

Lemma 14 *For all moves $(s \mapsto d) \in \text{parmove}(\mu_0)$, we have $s \in \{s_1, \dots, s_n\} \cup \mathcal{T}$ and $d \in \{d_1, \dots, d_n\} \cup \mathcal{T}$.*

Proof We show that $P(s \mapsto d)$ implies $s \in \{s_1, \dots, s_n\} \cup \mathcal{T}$ and $d \in \{d_1, \dots, d_n\} \cup \mathcal{T}$ by induction on a derivation of $P(s \mapsto d)$, then conclude with lemma 13.

Another application of lemma 13 is to show that the compilation of parallel moves preserves register classes. Most processors divide their register set in several classes, e.g. a class of scalar registers and a class of floating-point registers. Instructions are provided to perform a register-to-register move within the same class, but moves between two register of different classes are not supported. Assume that every register r has an associated class $\Gamma(r)$. We say that a list of moves l respects register classes if $\Gamma(s) = \Gamma(d)$ for all $(s \mapsto d) \in l$.

Lemma 15 (Register class preservation) *Assume that temporaries are generated in a class-preserving manner: $\Gamma(T(s)) = \Gamma(s)$ for all registers s . If μ_0 respects register classes, then $\text{parmove}(\mu_0)$ respects register classes as well.*

Proof We show that $P(s \mapsto d)$ implies $\Gamma(s) = \Gamma(d)$ by induction on a derivation of $P(s \mapsto d)$, then conclude using lemma 13.

8 Extension to overlapping registers

So far, we have assumed that two distinct registers never overlap: assigning one does not change the value of the other. This is not always the case in practice. For instance, some processor architectures expose hardware registers that share some of their bits. In the IA32 architecture, for example, the register **AL** refers to the low 8 bits of the 32-bit register **EAX**. Assigning to **AL** therefore modifies the value of **EAX**, and conversely.

Overlap also occurs naturally when the “registers” manipulated by the parallel move compilation algorithm include memory locations, such as variables that have been spilled to memory during register allocation, or stack locations used for parameter passing. For example, writing a 32-bit integer at offset δ in the stack also changes the values of the stack locations at offsets $\delta + 1$, $\delta + 2$ and $\delta + 3$.

It is not straightforward to generalize the parallel move compilation algorithm to the case where source and destination registers can overlap arbitrarily. We will not attempt to do so in this section, but set out to show a weaker, but still useful result: the unmodified parallel move algorithm produces correct sequences of elementary assignments even if registers can in general overlap, provided destinations and sources do not overlap.

8.1 Formalising overlap

In the non-overlapping case, two registers r_1 and r_2 are either identical $r_1 = r_2$ or different $r_1 \neq r_2$. When registers can overlap, we have three cases: r_1 and r_2 are either identical ($r_1 = r_2$) or completely disjoint (written $r_1 \perp r_2$) or different but partially overlapping (written $r_1 \bowtie r_2$).

We assume given a disjointness relation $\perp$ over $\mathcal{R} \times \mathcal{R}$, which must be symmetric and such that $r_1 \perp r_2 \Rightarrow r_1 \neq r_2$. We define partial overlap $r_1 \bowtie r_2$ as $(r_1 \neq r_2) \wedge \neg(r_1 \perp r_2)$.

In the non-overlapping case, the semantics of an assignment of value v to register r is captured by the update operation $\rho[r \leftarrow v]$, characterized by the “good variable” property:

$$\rho[r \leftarrow v] = \begin{cases} r \mapsto v \\ r' \mapsto \rho(r') \end{cases} \quad \text{if } r' \neq r.$$

To account for the possibility of overlap, we define the weak update operation $\rho[r \Leftarrow v]$ by the following “weak good variable” property:

$$\rho[r \Leftarrow v] = \begin{cases} r \mapsto v \\ r' \mapsto \rho(r') \end{cases} \quad \text{if } r' \perp r \\ \begin{cases} r' \mapsto \text{undefined} \end{cases} \quad \text{if } r' \bowtie r.$$

As suggested by the fatter arrow $\Leftarrow$, weak update sets the target register r to the specified value v , but causes “collateral damage” on registers r' that partially overlap with r : their values after the update are undefined. Only registers r' that are disjoint from r keep their old values.

Using weak update instead of update, the semantics of a sequence τ of elementary moves becomes

$$\begin{aligned} \llbracket \emptyset \rrbracket \Rightarrow (\rho) &= \rho \\ \llbracket (s \mapsto d) \cdot \tau \rrbracket \Rightarrow (\rho) &= \llbracket \tau \rrbracket \Rightarrow (\rho[d \Leftarrow \rho(s)]) \end{aligned}$$

8.2 Effect of overlap on the parallel move algorithm

In this section, we consider a parallel move problem $\mu = (s_1 \mapsto d_1) \cdots (s_n \mapsto d_n)$ and assume that the sources $s_1, \dots, s_n$ and the destinations $d_1, \dots, d_n$ satisfy the following hypotheses:

1. No temporaries: $s_i \perp t$ and $d_i \perp t$ for all i, j and $t \in \mathcal{T}$.
2. Destinations are pairwise disjoint: $d_i \perp d_j$ if $i \neq j$.
3. Sources and destinations do not partially overlap: $s_i \neq d_j \Rightarrow s_i \perp d_j$ for all i, j .
4. Distinct temporaries do not partially overlap: $t_1 \neq t_2 \Rightarrow t_1 \perp t_2$ for all $t_1, t_2 \in \mathcal{T}$.

As we show below, these hypotheses ensure that assigning a destination or a temporary r preserves the values of all sources, destinations and temporaries other than r .

These hypotheses are easily satisfied in our application scenario within the Compert compiler. Hardware registers never partially overlap in the target architecture (the PowerPC processor). Properties of the register allocator ensure that the only possibility for partial overlap is between stack locations used as destinations for parameter passing, but such overlap is avoided by the calling conventions used.

Since $r_1 \perp r_2 \Rightarrow r_1 \neq r_2$, hypothesis (1) ensures that no temporary occurs in sources and destinations, and hypothesis (2) ensures that μ is a windmill. Therefore, the initial state $(\mu, \emptyset, \emptyset)$ is well-formed.

We now set out to prove a correctness result for the sequence of elementary moves $\tau = \text{parmove}(\mu)$ produced by the parallel move compilation function. Namely, we wish to show that the final values of the destinations d_i are the initial values of the sources s_i , and that all registers disjoint from the destinations and from the temporaries keep their initial values. To this end, we define the following relation $\cong$ between environments:

$$\rho_1 \cong \rho_2 \stackrel{\text{def}}{=} (\forall r \notin \mathcal{D}, \rho_1(r) = \rho_2(r))$$

where $\mathcal{D}$ is the set of registers that partially overlap with one of the destinations or one of the temporaries:

$$\mathcal{D} \stackrel{\text{def}}{=} \{r \mid \exists i, r \bowtie d_i\} \cup \{r \mid \exists t \in \mathcal{T}, r \bowtie t\}.$$

In other words, $\rho_1 \cong \rho_2$ holds if ρ_1 and ρ_2 assign the same values to registers, except perhaps those registers that could be set to an undefined value as a side-effect of assigning one of the destinations or one of the temporaries.

Lemma 16 *$r \notin \mathcal{D}$ if r is one of the destinations d_i , or one of the sources s_i , or a temporary $t \in \mathcal{T}$.*

Proof Follows from hypotheses (1) to (4). ✎

Using the relation $\cong$, we can relate the effect of executing moves using normal, overlap-unaware update and weak, overlap-aware update.

Lemma 17 *Consider an elementary move $(s \mapsto d)$ where $s \in \{s_1, \dots, s_n\} \cup \mathcal{T}$ and $d \in \{d_1, \dots, d_n\} \cup \mathcal{T}$. If $\rho_1 \cong \rho_2$, then $\rho_1[d \leftarrow \rho_1(s)] \cong \rho_2[d \leftarrow \rho_2(s)]$.*

Proof We need to show that

$$\rho_1[d \leftarrow \rho_1(s)](r) = \rho_2[d \leftarrow \rho_2(s)](r) \quad (*)$$

for all registers $r \notin \mathcal{D}$. By definition of $\mathcal{D}$, it must be the case that either $r = d$ or $r \perp d$, otherwise r would partially overlap d , which is a destination or a temporary, contradicting $r \notin \mathcal{D}$.

In the first case $r = d$, the left-hand side of $(*)$ is equal to $\rho_1(s)$ and the right-hand side to $\rho_2(s)$. We do have $\rho_1(s) = \rho_2(s)$ by hypothesis $\rho_1 \cong \rho_2$ and the fact that $s \notin \mathcal{D}$ by lemma 16.

In the second case $r \perp d$, using the good variable property and the fact that $r \neq d$, the left-hand side of $(*)$ is equal to $\rho_1(d)$. Using the weak good variable property, the right-hand side is $\rho_2(d)$. The equality $\rho_1(d) = \rho_2(d)$ follows from hypothesis $\rho_1 \cong \rho_2$.

The previous lemma extends to sequences of elementary moves, performed with normal updates on one side and with weak updates on the other side.

Lemma 18 *Let τ be a sequence of moves such that for all $(s \mapsto d) \in \tau$, we have $s \in \{s_1, \dots, s_n\} \cup \mathcal{T}$ and $d \in \{d_1, \dots, d_n\} \cup \mathcal{T}$. If $\rho_1 \cong \rho_2$, then $\llbracket \tau \rrbracket \rightarrow (\rho_1) = \llbracket \tau \rrbracket \rightarrow (\rho_2)$.*

Proof By structural induction on τ , using lemma 17.

Theorem 4 *Let $\mu = (s_1 \mapsto d_1) \cdots (s_n \mapsto d_n)$ be a parallel move problem that satisfies hypotheses (1) to (4). Let $\tau = \text{parmove}(\mu)$. For all environments ρ , writing $\rho' = \llbracket \tau \rrbracket \rightarrow (\rho)$, we have*

1. $\rho'(d_i) = \rho(s_i)$ for all $i = 1, \dots, n$;
2. $\rho'(r) = \rho(r)$ for all registers r disjoint from the destinations d_i and from the temporaries $\mathcal{T}$.

Proof Define $\rho_1 = \llbracket \tau \rrbracket \rightarrow (\rho)$. By theorem 3, we have $\rho_1(d_i) = \rho(s_i)$ for all i , and $\rho_1(r) = \rho(r)$ for all $r \notin \{d_1, \dots, d_n\} \cup \mathcal{T}$. By lemma 14, every move $(s \mapsto d) \in \tau$ is such that s is either one of the sources s_i or a temporary, and d is either one of the destinations d_i or a temporary. Since $\rho \cong \rho$ holds trivially, lemma 18 applies and shows that $\rho_1 \cong \rho'$.

Let d_i be one of the destinations. We have $\rho'(d_i) = \rho_1(d_i)$ since $d_i \notin \mathcal{D}$ by lemma 16. Moreover, $\rho_1(d_i) = \rho(s_i)$. The expected result follows.

Let r be a register disjoint from the destinations d_i and from the temporaries $\mathcal{T}$. By definition of $\mathcal{D}$, $r \notin \mathcal{D}$, therefore $\rho'(r) = \rho_1(r)$. Moreover, $\rho_1(r) = \rho(r)$ since $r \notin \{d_1, \dots, d_n\} \cup \mathcal{T}$. The expected result follows.

9 Conclusions

Using the Coq proof assistant, we have proved the correctness and termination of a compilation algorithm that serialises parallel moves. The main difficulty of this development was to find an appropriate set of atomic transitions between intermediate states of the algorithm, and prove that they preserve semantics; once this is done, the proof of the algorithm proper follows easily by refinements. The Coq development is relatively small: 580 lines of specifications and 610 lines of proof scripts.

The approach followed in this article enabled us to slightly improve the imperative algorithm from section 3 used as a starting point. In particular, the functional algorithm looks only for cycles involving the edge that started the recursion of the `move_one` function (i.e. the last element of the *being-moved* list, whereas the imperative algorithm looks for cycles anywhere in this list. The correctness proof for the $\triangleright$ relation implies that looking for the cycle from the bottom of the *being-moved* stack is sufficient: in the case for the [Pop] rule, we proved that all elements of the *being-moved* list except the last one cannot have the destination d_n as source.

While specified in terms of parallel moves, our results can probably be extended to parallel assignments $(x_1, \dots, x_n) := (e_1, \dots, e_n)$ where every expression e_i mentions at most one of the variables x_j .

Although efficient in practice, neither the initial imperative algorithm nor the functional algorithm proved in this paper are optimal in the number of elementary moves produced. Consider for instance the parallel move $(r_3, r_2, r_1) := (r_1, r_1, r_2)$. The algorithms generate a sequence of 4 moves: $r_3 := r_1$; $t := r_2$; $r_2 := r_1$; $r_1 := t$, where t is a temporary. However, the effect can be achieved in 3 moves, using the destination register r_3 to break the cycle: $r_3 := r_1$; $r_1 := r_2$; $r_2 := r_3$. Preliminary investigations suggest that the nondeterministic specification of section 4 can be extended with one additional rule to support this use of a destination to break cycles. However, the functional algorithm would become much more complex and require backtracking, possibly leading to exponential complexity.

A direction for further work is to prove the correctness of an imperative, array-based formulation of the parallel move compilation algorithm, similar to the algorithm given in section 3. We are considering using the Why tool [8, 9] to conduct this proof. It raises several difficulties. The two nested loops and the number of arrays needed (`src`, `dst` and `status`) imply large invariants. Moreover the proof obligations generated by Why are huge and hard to work with. Furthermore, due to the side effects over arrays, delicate non-aliasing lemmas must be proved.

References

1. Appel, A.W.: Compiling with continuations. Cambridge University Press (1992)
2. Balaa, A., Bertot, Y.: Fonctions récurrentes générales par itération en théorie des types. In: Journées Francophones des Langages Applicatifs 2002, pp. 27–42. INRIA (2002)
3. Barthe, G., Forest, J., Pichardie, D., Rusu, V.: Defining and reasoning about recursive functions: a practical tool for the Coq proof assistant. In: Proc. 8th Int. Symp. on Functional and Logic Programming (FLOPS'06), *Lecture Notes in Computer Science*, vol. 3945, pp. 114–129. Springer (2006)
4. Bertot, Y., Castéran, P.: Interactive Theorem Proving and Program Development – Coq'Art: The Calculus of Inductive Constructions. EATCS Texts in Theoretical Computer Science. Springer (2004)
5. Bertot, Y., Grégoire, B., Leroy, X.: A structured approach to proving compiler optimizations based on dataflow analysis. In: Types for Proofs and Programs, Workshop TYPES 2004, *Lecture Notes in Computer Science*, vol. 3839, pp. 66–81. Springer (2006)
6. Blazy, S., Dargaye, Z., Leroy, X.: Formal verification of a C compiler front-end. In: FM 2006: Int. Symp. on Formal Methods, *Lecture Notes in Computer Science*, vol. 4085, pp. 460–475. Springer (2006)
7. Coq development team: The Coq proof assistant. Software and documentation available at <http://coq.inria.fr/> (1989–2007)
8. Filliâtre, J.C.: Verification of non-functional programs using interpretations in type theory. *Journal of Functional Programming* **13**(4), 709–745 (2003)
9. Filliâtre, J.C.: The Why software verification tool. Software and documentation available at <http://why.lri.fr/> (2003–2007)

10. Leroy, X.: Formal certification of a compiler back-end, or: programming a compiler with a proof assistant. In: 33rd symposium Principles of Programming Languages, pp. 42–54. ACM Press (2006)
11. Leroy, X., Doligez, D., Garrigue, J., Vouillon, J.: The Objective Caml system. Software and documentation available at <http://caml.inria.fr/> (1996–2007)
12. Letouzey, P.: A new extraction for Coq. In: Types for Proofs and Programs, Workshop TYPES 2002, *Lecture Notes in Computer Science*, vol. 2646, pp. 200–219. Springer (2003)
13. May, C.: The parallel assignment problem redefined. *IEEE Transactions on Software Engineering* **15**(6), 821–824 (1989)
14. Sethi, R.: A note on implementing parallel assignment instructions. *Information Processing Letters* **2**(4), 91–95 (1973)
15. Welch, P.H.: Parallel assignment revisited. *Software Practice and Experience* **13**(12), 1175–1180 (1983)