Certified Undecidability of Intuitionistic Linear Logic via Binary Stack Machines and Minsky Machines

Yannick Forster, Dominique Larchey-Wendling

► To cite this version:

Yannick Forster, Dominique Larchey-Wendling. Certified Undecidability of Intuitionistic Linear Logic via Binary Stack Machines and Minsky Machines. The 8th ACM SIGPLAN International Conference on Certified Programs and Proofs, CPP 2019, Jan 2019, Cascais, Portugal. pp.104-117, 10.1145/3293880.3294096 . hal-02333390

HAL Id: hal-02333390

<https://hal.science/hal-02333390>

Submitted on 9 Nov 2020

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.

Certified Undecidability of Intuitionistic Linear Logic via Binary Stack Machines and Minsky Machines

Yannick Forster
Saarland University
Saarbrücken, Germany
forster@ps.uni-saarland.de

Dominique Larchey-Wendling
Université de Lorraine, CNRS, LORIA
Vandœuvre-lès-Nancy, France
dominique.larchey-wendling@loria.fr

Abstract

We formally prove the undecidability of entailment in intuitionistic linear logic in Coq. We reduce the Post correspondence problem (PCP) via binary stack machines and Minsky machines to intuitionistic linear logic. The reductions rely on several technically involved formalisations, amongst them a binary stack machine simulator for PCP, a verified low-level compiler for instruction-based languages and a soundness proof for intuitionistic linear logic with respect to trivial phase semantics. We exploit the computability of all functions definable in constructive type theory and thus do not have to rely on a concrete model of computation, enabling the reduction proofs to focus on correctness properties.

Keywords Undecidability, many-one reduction, binary stack machines, Minsky machines, intuitionistic linear logic, low-level compiler, constructive type theory, Coq

1 Introduction

Undecidability of advanced problems is usually established using computable (many-one) reductions from a problem already known to be undecidable. Such reductions rely on many subtle details and could thus be a prime example for the use of interactive theorem provers to assist ongoing research. However, formalisations of undecidability proofs are not common in the literature. There are two main obstacles in our eyes: (1) proofs on paper mostly omit the invariants needed for the verification of the reduction; (2) they omit the computability proof, which amounts to the formal verification of a program in the chosen model of computation.

Constructive type theory, as implemented in the proof assistant Coq [36], provides a particularly convenient setting for decidability and undecidability proofs. Since every function definable in constructive type theory is computable, one can use a synthetic approach to computability. This approach makes an explicit model of computation and explicit computability proofs unnecessary, enabling the proofs to focus on the invariants needed for the reduction.

In this paper we contribute a formalised chain of reductions, starting at the Post correspondence problem (PCP), via binary PCP (BPCP), BPCP with indices (iBPCP), binary stack

machines (BSM), Minsky machines (MM) and (elementary) intuitionistic linear logic (both eILL and ILL):

$$\text{PCP} \leq \text{BPCP} \leq \text{iBPCP} \leq \text{BSM} \leq \text{MM} \leq \text{eILL} \leq \text{ILL}$$

Combined with the reduction $\text{Halt} \leq \text{PCP}$ of [11], this yields a fully formalised reduction from the halting problem of Turing machines to entailment in ILL.

PCP is a problem over a finite set of cards, each having a string at the top and bottom. Such a PCP instance is solvable if there is a non-empty finite sequence of those cards with an equal upper and lower string. We follow [11] in giving a definition of PCP with countably many symbols. We reduce PCP to BPCP, where symbols are restricted to Booleans and to iBPCP, where the sequence of cards is represented by a sequence of their indices (natural numbers).

We describe a general framework for the semantics of instruction-based machine models and explain how it can be used for compositional reasoning. We then define binary stack machines as an instance of such instruction-based machines. BSMs have a fixed number of binary stacks and programs consist of consecutively indexed PUSH and POP instructions, the latter allowing conditional jumps based on the obtained Boolean or if the stack is empty. A BSM program terminates if it jumps to a non-existing index. We verify an iBPCP simulator — which is a BSM program terminating if and only if a given iBPCP instance is solvable — by enumerating all possible sequences of cards. The program size is linear in the size of the instance and always has more than 80 instructions. It consists of several independent sub-programs, making the verification crucially rely on compositional reasoning.

We define Minsky machines [32] with a fixed number of natural-number registers as another instance of instruction-based machines. MM-programs consist of INC and DEC operations, the latter allowing conditional jumps if the register is zero. We describe and verify a general compiler for instruction-based languages and use it to compile BSM programs to MM programs. The main challenge is to rearrange jumps using a linker. For the reduction we have to prove correctness of the compiler for both terminating and non-terminating programs, a notable difference to other verified compilers (e.g. [6, 28]).

We formalise intuitionistic linear logic with a sequent calculus S_{ILL} . We introduce the elementary fragment of ILL [26,

[27] together with a goal directed sequent calculus G_{eLL} which is equivalent to S_{iLL} on this fragment: both define the same entailment predicate. We give a denotational semantics for iLL using trivial phase semantics based on commutative monoids and prove soundness of S_{iLL} for this semantics. Finally, we show that entailment in G_{eLL} is complete for trivial phase semantics, yielding a reduction of eLL to iLL . Our arguments slightly generalise those of the proof in [27] to formalise the reduction chain $MM \leq eLL \leq iLL$.

We conclude the paper by summarising our efforts to build a library of undecidable problems in Coq [14] and by giving an overview of related work. The library should contain both entry points for plugging new undecidability results and tools to mechanise the corresponding reductions. We discuss some ongoing additions to our library.

The axiom-free Coq formalisation of all the results in this paper is available online and the main lemmas and theorems in the pdf version of the paper are hyper-linked with the html version of the Coq source code:

<https://uds-psl.github.io/ill-undecidability>

2 Type Theory and Formal Undecidability

2.1 Definitions

We write $\mathbb{P}$ to denote the type of propositions. The basic inductive data structures we use are *natural numbers* ($n : \mathbb{N} ::= 0 \mid S n$), *Booleans* ($b : \mathbb{B} ::= 0 \mid 1$), the *option type* ($\odot X ::= (x : X) \mid \star$) and *lists* ($l : \mathbb{L} X ::= [] \mid x :: l$). We write $l_1 \# l_2$ for the *concatenation* of two lists, $|l|$ for the length of a list, $\bar{l}$ for the reversal of a list, and $[f s \mid s \in l]$ for a map over a list. We define $l[i] : \odot X$ as the i -th element of l . Moreover, we write $s \in l$ if s is a member of l , and $l_1 \subseteq l_2$ if every member of l_1 is a member of l_2 .

We will overload the notation for lists and use it for vectors as well. Given a type X , we write the type of vectors of length n as X^n . We write the finite datatype of n elements as $\alpha, \beta : \mathbb{F}_n ::= 0 \mid \dots \mid n - 1$. If $v : X^n$, $x : X$ and $\alpha : \mathbb{F}_n$, we define $v[\alpha]$ as the α -th component of v and $v[\alpha := x]$ as the vector v as the α -th component updated by x .

2.2 Undecidability in Coq

A *decision problem* consists of a type X and a unary predicate $p : X \rightarrow \mathbb{P}$ on X . A problem (X, p) is *decidable* if there is a function $f : X \rightarrow \mathbb{B}$ s.t. $\forall x. p x \leftrightarrow f x = 1$, or equivalently, if there is a (dependent) *decider* of type $\forall x : X, \{p x\} + \{\neg p x\}$.

Being able to define the concept of decidability without referring to a model of computation is a feature of constructive type theory: due to the commitment of its designers to effective methods, all definable functions are computable. There has been a lot of work on proving problems to be decidable using this notion, see e.g. [5, 8, 9, 22, 25, 31, 34].

Note that since it is consistent in constructive type theory to assume that all problems are “decidable,” i.e. by the axiom $\forall P : \mathbb{P}, \{P\} + \{\neg P\}$, the logical negation of decidability of

a problem p can never be actually proven.¹ However, in practice, the undecidability of a problem is rarely established directly anyway. Most proofs are by reduction from another undecidable problem, and ultimately most problems are thus proven undecidable by reduction from an initial undecidable problem. In most textbook presentations, it is the halting problem (Halt) which is seen as a *seed of undecidability*.

The idea to exploit reductions to prove the undecidability of problems is well-known. It was first used to formalise undecidability proofs in a proof assistant without referring to an explicit model of computation in [11]. They define (many-one) reductions from a problem (X, p) to (Y, q) as a (computable) function $f : X \rightarrow Y$ such that $\forall x. p x \leftrightarrow q (f x)$. We write $p \leq q$ and say that p *reduces to* q if a reduction of (X, p) to (Y, q) exists. In the case of axiom-free Coq, all such reductions are inherently computable and thus, computability can be dropped from the definition. This coincides with usual practice in non-formalised papers where the computability of reductions is left out most of the time.

Actually, many-one reductions are only an instance of the more general notion of Turing-reductions (see e.g. [33]): (X, p) is Turing-reducible to (Y, q) , written $p \leq_T q$, if p is decidable given that q is decidable.

Lemma 2.1. *If $p \leq q$ then $p \leq_T q$.*

The structure of undecidability proofs can be comprehensively described using an inductive predicate $\text{undec } p$:

$$\frac{}{\text{undec Halt}} \quad \frac{p \leq_T q \quad \text{undec } q}{\text{undec } p}$$

The first rule establishes the halting problem for Turing machines as a seed of undecidability. By the second rule, undecidability can be transferred via (Turing) reductions.

Lemma 2.2. *$\text{undec } (\lambda x. \neg p x)$ implies $\text{undec } p$.*

Proof. If p is decidable then so is $\lambda x. \neg p x$. □

Lemma 2.3. *$\text{undec } p$ if and only if $\text{Halt} \leq_T p$.*

A problem that would be decidable and undecidable at the same time would entail an obviously unwanted statement:

Lemma 2.4. *If p is decidable and $\text{undec } p$ then Halt is decidable.*

We rely on the assumptions that Halt in [11] is a faithful implementation of the halting problem and that functions in Coq are indeed computable. If the latter would not be the case, the earlier mentioned publications on decidability proofs would all lose their grounds.

The choice of Halt as the seed of undecidability is arbitrary: most problems that are universally accepted as undecidable

¹This illustrates that adding axioms to Coq could render the assumption that all functions are computable invalid, thus we safely commit ourselves to axiom-free Coq.

would work. Since the halting problem Halt for a concrete Turing machine model has already been reduced to PCP in Coq [11], we can safely use PCP as an alternate seed of undecidability. This also explains why the reductions in this paper start at PCP. Remark that we use this inductive definition of undec only for illustrative purposes. For our proofs, we simply rely on many-one reductions.

3 Post Correspondence Problem

All definitions concerning strings will be parametrised over a type X of symbols. We will later use natural numbers as symbols for PCP and Booleans as symbols for BPCP. A *string* is a list of symbols. We write X^* if the symbol type is X . We call strings over Booleans *bitstrings*. The letters x, y, z, u , and v range over strings, and the letters a, b, c range over symbols. We write xy for $x \mathbin{\dot{+}} y$ and ax for $a :: x$. We use ϵ to denote the empty string.

A *card* x/y is a pair (x, y) of two strings. When we call x/y a card we see x as the upper and y as the lower string of the card. A *card set* (CS) is finite set of cards represented by a list (where the order and multiplicity do not matter). A *stack* is a sequence of cards also represented by a list (but the order and multiplicity matter). The letters A, B, C range over stacks and the letter R ranges over CS. We use the letter I to range over a list of natural numbers.

The *upper trace* A^1 and the *lower trace* A^2 of a stack A are strings defined by the following equations:

$$\square^{1,2} := \epsilon \quad (x/y :: A)^1 := x(A^1) \quad (x/y :: A)^2 := y(A^2)$$

Note that A^1 is the concatenation of the upper strings of the cards in A , and that A^2 is the concatenation of the lower strings of the cards in A . We say that a stack A *matches* if $A^1 = A^2$. A *match* is a matching stack.

We define the predicate for the *Post correspondence problem* over elements of type X as symbols as:

$$\text{PCP}(R : \text{CS}_X) := \exists A \subseteq R. A \neq \square \wedge A^1 = A^2$$

Note that $\text{PCP}(R)$ holds iff there exists a non-empty match $A \subseteq R$. If we do not specify the type X , we mean the special case $\text{PCP}(R : \text{CS}_{\mathbb{N}})$ where the symbols type is $\mathbb{N}$, used as the standard one in [11]. Here we also define the *binary Post correspondence problem* fixing Booleans as symbols: $\text{BPCP}(R : \text{CS}_{\mathbb{B}}) := \text{PCP}(R : \text{CS}_{\mathbb{B}})$.

Finally, we define a version of BPCP using explicit indices, which is easier to treat for low-level machines. We define functions I_R^k for lists of indices I and $k \in \{1, 2\}$.

$$\begin{aligned} I_R^k &:= \epsilon & (i :: I)_R^k &:= A_R^k \quad (\text{for } R[i] \text{ undefined}) \\ (i :: I)_R^1 &:= x(A_R^1) & (i :: I)_R^2 &:= y(A_R^2) \quad (\text{for } R[i] = x/y) \end{aligned}$$

and define binary PCP with indices:

$$\text{iBPCP}(R : \text{CS}_{\mathbb{B}}) := \exists I \neq \square, (\forall i \in I, i < |R|) \wedge I_R^1 = I_R^2$$

Note the equivalence when we reverse lists of indices:

Lemma 3.1. $\text{iBPCP } R \text{ iff } \exists I \neq \square, (\forall i \in I, i < |R|) \wedge \bar{I}_R^1 = \bar{I}_R^2$.

3.1 PCP reduces to BPCP

We define a function $f : \mathbb{N}^* \rightarrow \mathbb{B}^*$ by $f \epsilon := \epsilon$ and $f(nx) := 1^n 0 (f x)$ and we extend f to cards and lists of cards by pointwise application.

Lemma 3.2. *The following hold for $k \in \{1, 2\}$:*

1. $(f A)^k = f(A^k)$ and 2. If $A \subseteq C$ then $f A \subseteq f C$.

We define the inverse function $g : \mathbb{B}^* \rightarrow \mathbb{N} \rightarrow \mathbb{N}^*$ with an auxiliary argument by $g \epsilon n := \epsilon$ and $g(1x)n := g x (1 + n)$ and $g(0x)n := n(g x 0)$. We write $g x$ for $g x 0$ and again extend it pointwise to cards and stacks.

Lemma 3.3. *If $k \in \{1, 2\}$ and $A \subseteq f B$ then*

1. $(g A)^k = g(A^k)$ and 2. $g A \subseteq B$.

Theorem 3.4. $\text{PCP} \leq \text{BPCP}$.

Proof. We prove that $\text{PCP } R \leftrightarrow \text{BPCP}(f R)$.

For the direction from left to right let $\square \neq A \subseteq R$ with $A^1 = A^2$ be given. Then $f A \subseteq f R$ by Lemma 3.2 (2), $f A \neq \square$ is trivial and $(f A)^1 = (f A)^2$ follows from Lemma 3.2 (1).

For the direction from right to left let $\square \neq B \subseteq f R$ with $B^1 = B^2$ be given. $g B \subseteq R$ follows by Lemma 3.3 (2), $g B \neq \square$ is trivial and $(g B)^1 = (g B)^2$ by Lemma 3.3 (1). $\square$

3.2 BPCP reduces to iBPCP

We define

$$\begin{aligned} f R \square &:= \square \\ f R(x/y :: A) &:= n :: f R A & (\text{if } x/y \in R \text{ at pos. } n) \\ f R(x/y :: A) &:= f R A & (\text{if } x/y \notin R) \end{aligned}$$

and

$$\begin{aligned} g R \square &:= \square \\ g R(i :: I) &:= (x/y)(g R I) & (\text{if } R[i] = x/y) \\ g R(i :: I) &:= g R I & (\text{if } R[i] \text{ is undefined}) \end{aligned}$$

Lemma 3.5. *The following hold for $k \in \{1, 2\}$:*

1. If $B \subseteq R$ then $(f R B)_R^k = B^k$
2. If $(\forall i \in I, i < |R|)$ then $(g R I)^k = I_R^k$.

Proof. (1) is by induction on B , (2) by induction on I . $\square$

Theorem 3.6. $\text{BPCP} \leq \text{iBPCP}$.

4 Code

We define code over arbitrary instruction sets to reuse later for both binary stack machines and Minsky machines. Since both instruction sets contain absolute jump instructions, the positions of instructions are represented by natural numbers.

4.1 Syntax

We model the type of code type-theoretically by defining it over a type of instructions $\mathbb{I}$ as $\mathbb{N} \times \mathbb{L}\mathbb{I}$. We use the letter ρ for individual instructions and P for lists of instructions. A code $(i, [\rho_0; \rho_1; \dots; \rho_{n-1}])$ can be seen as a labelled program:

$$i : \rho_0; \quad i + 1 : \rho_1; \quad \dots \quad i + n - 1 : \rho_{n-1};$$

where $i, i + 1, \dots, i + n - 1$ represent *program counter (PC)* values which identify the positions of individual instructions. The reason we use any number i instead of just 0 as the first label is to be able to reason about the whole code as well as about segments in the code, segments which we call *subcode*. We define the *subcode relation* denoted $(i, P) <_{sc} (j, Q)$ by

$$(i, P) <_{sc} (j, Q) := \exists l, r, Q = l \# P \# r \wedge i = j + |l|$$

meaning that P is the segment of Q starting at label i . Many of the upcoming semantic relations below defined will be “somehow preserved” by the subcode relation and this is an essential ingredient of *compositional reasoning over subcode*; see Section 4.3 for more precise preservation properties.

The start label of a code (i, P) is $\text{start}(i, P) := i$. The end of code for (i, P) is the first label above the code block, i.e. $\text{end}(i, P) := i + |P|$. We define when a PC value k is in(side) or else out(side) of the code block (i, P) as

$$\begin{aligned} \text{in } k (i, P) &:= \text{start}(i, P) \leq k < \text{end}(i, P) \\ \text{out } k (i, P) &:= k < \text{start}(i, P) \vee \text{end}(i, P) \leq k. \end{aligned}$$

4.2 Semantics

The semantics of all our machines is defined by providing a type of states $\mathbb{S} := \mathbb{N} \times \mathbb{C}$ where a state $(i, v) : \mathbb{S}$ is composed of a PC value $i : \mathbb{N}$ and a *configuration* $v : \mathbb{C}$, and a step relation denoted $\rho \parallel (i_1, v_1) > (i_2, v_2)$ of type $\mathbb{I} \rightarrow \mathbb{S} \rightarrow \mathbb{S} \rightarrow \mathbb{P}$ which describes the effect of instruction ρ , transforming the state (i_1, v_1) into the state (i_2, v_2) .

We inductively extend the step relation to a *small-step semantics* of type $\mathbb{N} \times \mathbb{L}\mathbb{I} \rightarrow \mathbb{S} \rightarrow \mathbb{N} \rightarrow \mathbb{S} \rightarrow \mathbb{P}$ denoted by $(i, P) \parallel (i_1, v_1) >^n (i_2, v_2)$, and meaning that the program (i, P) , starting from state (i_1, v_1) , reaches state (i_2, v_2) after the execution of n individual instructions:

$$\frac{\begin{array}{c} (i, P) \parallel (i_1, v_1) >^0 (i_1, v_1) \\ \hline \begin{array}{cc} P = P_1 \# \rho :: P_2 & i_1 = i + |P_1| \\ \rho \parallel (i_1, v_1) > (i_2, v_2) & (i, P) \parallel (i_2, v_2) >^n (i_3, v_3) \end{array} \\ \hline (i, P) \parallel (i_1, v_1) >^{n+1} (i_3, v_3) \end{array}}$$

The first two premises of the second rule express that instruction ρ occurs at position i_1 in the program (i, P) , or equivalently $(i_1, [\rho]) <_{sc} (i, P)$. Notice that whenever j satisfies $\text{out } j (i, P)$ then $(i, P) \parallel (j, v) >^n s$ implies $n = 0$ and $s = (j, v)$, i.e. the computation is blocked outside of (i, P) .

To shorten notation, we will often use s to represent states like (i, v) and $\mathcal{P}$ to represent code like (i, P) . We define the

transitive closure $>^+$ and reflexive-transitive closure $>^*$ of the step relation as follows:

$$\begin{aligned} \mathcal{P} \parallel s >^+ s' &:= \exists n > 0, \mathcal{P} \parallel s >^n s' \\ \mathcal{P} \parallel s >^* s' &:= \exists n, \mathcal{P} \parallel s >^n s' \end{aligned}$$

We call the step relation *deterministic* if for any ρ and s, s_1 and s_2 , if $\rho \parallel s > s_1$ and $\rho \parallel s > s_2$ then $s_1 = s_2$. In this paper, we only consider deterministic step relations. In this case, the small-step semantics is also deterministic, i.e. for any $\mathcal{P}, s$ and n , there is at most one s' s.t. $\mathcal{P} \parallel s >^n s'$. However, there are generally many possible s' s.t. $\mathcal{P} \parallel s >^* s'$. Thus neither $>^*$ nor $>^+$ are a well-suited big-step semantics because they are not deterministic. By further assuming termination, we obtain the unicity of the end state.

Hence we define the (deterministic) big-step semantics as the *terminates on* predicate, and in addition, we also define the *terminates on in p steps* predicate, the *terminates on with progress* predicate and the *termination* predicate as:

$$\begin{aligned} \mathcal{P} \parallel s \rightsquigarrow (j, w) &:= \mathcal{P} \parallel s >^* (j, w) \wedge \text{out } j \mathcal{P} \\ \mathcal{P} \parallel s \rightsquigarrow^n (j, w) &:= \mathcal{P} \parallel s >^n (j, w) \wedge \text{out } j \mathcal{P} \\ \mathcal{P} \parallel s \rightsquigarrow^+ (j, w) &:= \mathcal{P} \parallel s >^+ (j, w) \wedge \text{out } j \mathcal{P} \\ \mathcal{P} \parallel s \downarrow &:= \exists s', \mathcal{P} \parallel s \rightsquigarrow s' \end{aligned}$$

Finally note that although $\text{out } j \mathcal{P}$ is a sufficient condition for j to be a state where the program cannot progress any more, this condition is not necessary. Indeed, in certain models of computation, some instructions might block a program, either unconditionally like a HALT instruction, or conditionally when a POP instruction is not allowed to run on an empty stack. To avoid having to distinguish between halting when outside or halting on a blocking instruction, we only manipulate machine models where the step relation is *total*, i.e. the property $\forall \rho \exists s s', \rho \parallel s > s'$ holds. In that case, being outside of the code is the only way to block a computation.

4.3 Compositional Reasoning over Subcode

Compositional reasoning is the ability to derive properties of the whole from the properties of its components. It is an essential tool in the certification of big developments such as low-level computer code.

In our setting, compositional reasoning translates into tools which allow the transfer of properties of subcode to larger code. For instance, the following property is a simple and obvious example of compositional reasoning:

Lemma 4.1. *Given $\mathcal{P} <_{sc} Q$ we have for any n, s, s' :*

$$\mathcal{P} \parallel s >^n s' \rightarrow Q \parallel s >^n s'$$

This reflects the fact that a computation running inside some subcode $\mathcal{P}$ of Q also runs inside Q itself. A more critical tool of compositional reasoning is the following result:

Lemma 4.2. *Given $\mathcal{P} <_{sc} Q$ we have for any s_1, s_3 :*

$$Q \parallel s_1 \rightsquigarrow s_3 \rightarrow \exists s_2, \mathcal{P} \parallel s_1 \rightsquigarrow s_2 \wedge Q \parallel s_2 \rightsquigarrow s_3$$

Any terminating computation in code Q can be decomposed into a prefix terminating computation in subcode $\mathcal{P}$ followed by a suffix terminating computation in Q .

Lemma 4.3. *If we assume the step relation $>$ to be deterministic, we get:*

$$\begin{aligned} \mathcal{P} <_{sc} Q \wedge \mathcal{P} \parallel s_1 >^+ s_2 \wedge Q \parallel s_1 \rightsquigarrow^p s_3 \\ \rightarrow \exists q < p, Q \parallel s_2 \rightsquigarrow^q s_3 \end{aligned}$$

This last result allows the analysis of terminating computations in Q by induction on their length, using the behaviour of its subcodes such as $\mathcal{P}$, as soon as they make the computation progress.

Such tools are our major examples of compositional reasoning: they are used to derive e.g. soundness and completeness properties of computations in code Q from the soundness properties of progressing computations in its subcodes.

5 Binary Stack Machines

We define *Binary Stack Machines (BSM)* as a particular instance of (low-level) code. BSM have a fixed number $n : \mathbb{N}$ of Boolean stacks modelled by lists of Booleans. We fix a number n of stacks as a parameter for all further definitions. BSM configurations are vectors v of n Boolean stacks, i.e. $\mathbb{C} := (\mathbb{L}\mathbb{B})^n$. We define BSM instructions as follows:

$$\mathbb{I}_n ::= \text{POP } (\alpha : \mathbb{F}_n) (p : \mathbb{N}) (q : \mathbb{N}) \mid \text{PUSH } (\alpha : \mathbb{F}_n) (b : \mathbb{B}).$$

5.1 Semantics for BSM

The rules defining the step relation for BSM instructions are:

$$\begin{aligned} \text{PUSH } \alpha \ b \parallel (i, v) > (1 + i, v[\alpha := b :: l]) & \quad \text{if } v[\alpha] = l \\ \text{POP } \alpha \ p \ q \parallel (i, v) > (q, v) & \quad \text{if } v[\alpha] = \square \\ \text{POP } \alpha \ p \ q \parallel (i, v) > (p, v[\alpha := l]) & \quad \text{if } v[\alpha] = 0 :: l \\ \text{POP } \alpha \ p \ q \parallel (i, v) > (1 + i, v[\alpha := l]) & \quad \text{if } v[\alpha] = 1 :: l \end{aligned}$$

The $\text{PUSH } \alpha \ b$ instruction pushes $b : \mathbb{B}$ to the stack α . The $\text{POP } \alpha \ p \ q$ instruction (1) jumps to PC value q if the stack α is empty, or (2) jumps to p if the top element of the stack α is 0 and removes it or (3) increases the PC by one if the top element of the stack α is 1 and removes it.

The above defined step relation for BSM is obviously deterministic and total: individual instructions always execute and depending on the input state, lead to a unique output state. Finally, we define the halting problem for binary stack machines as a predicate over a dependent quadruplet:

$$\text{BSM}(n : \mathbb{N}, i : \mathbb{N}, P : \mathbb{L}\mathbb{I}_n, v : (\mathbb{L}\mathbb{B})^n) := (i, P) \parallel (i, v) \downarrow.$$

Note that the execution starts with the first instruction of the program and that we do not care about its final state as long as it terminates (on a PC value outside of its code).

5.2 Overview of the Reduction from iBPCP to BSM

We first describe the overall idea of the reduction and then explain some parts in detail. Recall that an iBPCP instance is

given as a list $R : \mathbb{L}(\mathbb{L}\mathbb{B} \times \mathbb{L}\mathbb{B})$ of binary cards. The problem is to find a non-empty list of indices of cards of R such that the upper and lower trace are equal. We solve this problem using a simulator denoted pcp_bsm_R implemented as a BSM-program in which the iBPCP instance R is hard-coded.

So let us fix the list R . The program pcp_bsm_R will solve the instance generated by R . A card c_i in $R = [c_0; \dots; c_i; \dots]$ is represented by its index $i < |R|$. Finding a match is decomposed into three problems: (1) enumerate all the possible non-empty lists of card indices below $|R|$; (2) given a non-empty list of card indices, decode it into an upper trace and a lower trace and (3) compare those two traces for equality.

The simulator pcp_bsm_R has 4 binary stacks s, u, l and a : stack s contains the encoding of the current list of cards indices, u and l contain the upper and lower traces, and a is a spare stack for internal computations:

- (1) pcp_bsm_R solves part (1) by generating in s all the possible non-empty finite binary sequences. We use a simple encoding of $\mathbb{L}\mathbb{N}$ in $\mathbb{L}\mathbb{B}$: e.g. the list $[5; 3; 0; 2]$ of card indices is to be recognised as the sequence $000001 \cdot 0001 \cdot 1 \cdot 001$ (the $\cdot$ only helps at reading). Any binary sequence which is not a valid encoding of cards – be it terminating with a 0 or be it containing indices larger than $|R|$ – is going to be rejected as invalid in part (2) and pcp_bsm_R will branch again at (1) to generate the next possible binary sequence in s .
- (2) Given an encoding of a non-empty sequence of cards indices in s , pcp_bsm_R computes the upper (resp. lower) trace of that sequence in u (resp. l). To decode a card index at the top of s , it uses a dispatcher that analyses the head sequence of s of the shape $0^i 1 \dots$. If i exceeds $|R|$ or s ends up empty before a 1 is encountered, then the dispatcher rejects the whole sequence s as an invalid encoding of card indices and jumps back to part (1). Otherwise, it decodes the cards indices one after the other and, when encountering the index i of card c_i , defers the execution into a hard coded version of card $c_i = (x_i/y_i)$ that pushes the upper part of x_i into u and the lower part y_i into l , i.e. $u \rightsquigarrow x_i \uplus u$ and $l \rightsquigarrow y_i \uplus l$ and then continues the analysis of the remaining part of s . Note that after the decoding of the sequences $[i_1; \dots; i_k]$ of indices, we get $x_{i_k} \uplus \dots \uplus x_{i_1}$ in u and $y_{i_k} \uplus \dots \uplus y_{i_1}$ in l , i.e. the cards are stacked in reversed order, but we proved that this is no problem in Lemma 3.1.
- (3) Given an upper/lower trace in u/l , pcp_bsm_R solves part (3) by a 1-to-1 comparison between u and l . If u differs from l , pcp_bsm_R reverts back to generating the next binary sequence by branching at (1). If u equals l then pcp_bsm_R jumps to a terminating PC value outside of its code.

Notice that pcp_bsm_R never stops if there is no solution to the iBPCP instance R . On the other hand, since (1) enumerates any possible card list (including some possibly invalid encodings which are rejected as such), if there is a solution to R , then pcp_bsm_R will stop when (3) detects such a solution where $u = l$. Of course, these parts (1)-(3) correspond to specific BSM code for which correctness and completeness results are established. Here we only explain certain sub-programs of the pcp_bsm_R program.

5.3 Details of the Reduction from iBPCP to BSM

To solve part (1), we enumerate all index lists using a function $\text{next} : \mathbb{L}\mathbb{B} \rightarrow \mathbb{L}\mathbb{B}$ such that the sequence of its iterations starting at $[0]$ visits all non-empty binary lists:

Lemma 5.1. *There is a function $\text{next} : \mathbb{L}\mathbb{B} \rightarrow \mathbb{L}\mathbb{B}$ such that*

$$\forall A : \mathbb{L}\mathbb{B}, A \neq [] \rightarrow \exists n : \mathbb{N}, \text{next}^n [0] = A.$$

Then we show that next is BSM-computable:

Lemma 5.2. *For any $n, i : \mathbb{N}$ and $x \neq y : \mathbb{F}_n$ there is a program $\text{increment} : \mathbb{L}\mathbb{I}_n$ with 15 instructions satisfying*

$$\forall v, (i, \text{increment}) \parallel (i, v) >^* (15 + i, v[x := \text{next}(v[x])]).$$

Notice the absence of y in the specification of increment witnessing its use as a spare register: its content might be locally modified during the execution of increment but it is reverted back to its initial value afterwards. Using a simple loop over increment allows us to solve part (1) of pcp_bsm_R .

To solve part (2) we implement programs $\text{tile}_{\alpha/\beta}$ pushing the strings α and β to the stacks x and y :

Lemma 5.3. *For all $n : \mathbb{N}, x \neq y : \mathbb{F}_n$ and $\alpha, \beta : \mathbb{L}\mathbb{B}$ there is a program $\text{tile}_{\alpha/\beta} : \mathbb{L}\mathbb{I}_n$ with $|\text{tile}_{\alpha/\beta}| = |\alpha| + |\beta|$ satisfying the following specification:*

$$\begin{aligned} \forall v i w, w = v[x := \alpha + v[x], y := \beta + v[y]] \\ \rightarrow (i, \text{tile}_{\alpha/\beta}) \parallel (i, v) >^* (|\text{tile}_{\alpha/\beta}| + i, w). \end{aligned}$$

The cards decoder consists in gluing $\text{tile}_{c_1}, \text{tile}_{c_2}, \dots, \text{tile}_{c_k}$ where $R = [c_1; \dots; c_k]$, interleaved with a measure of the number j of head 0's in $v[s] = 0^j 1 \dots$ to branch on the appropriate tile_{c_j} . This decoders also detects and rejects a value $j \geq |R|$ or a stack $v[s]$ of the form 0^* not containing any 1. This card index decoder is implemented by structural induction on R . We refer to the Coq term `full_decoder` in file `bsm_utils.v` for more detailed explanations. This solves part (2) of pcp_bsm_R .

To solve part (3), we implement stack comparison.

Lemma 5.4. *For all $n : \mathbb{N}, x \neq y : \mathbb{F}_n$ and $i, p, q : \mathbb{N}$, there is a program $\text{compare\_stacks} : \mathbb{L}\mathbb{I}_n$ (of length 10) which satisfies the following specification:*

$$\begin{aligned} \forall v \exists j w, (i, \text{compare\_stacks}) \parallel (i, v) >^* (j, w) \\ \wedge \forall z, z \neq x \rightarrow z \neq y \rightarrow v[z] = w[z] \\ \wedge (v[x] = v[y] \wedge j = p \vee v[x] \neq v[y] \wedge j = q). \end{aligned}$$

Hence, compare_stacks can distinguish between $v[x] = v[y]$ and $v[x] \neq v[y]$ by jumping to either p or (another) PC value q . Notice that the resulting state w is not affected too much by the computation since it preserves any stack register except x and y . Since the values of $w[x]$ and $w[y]$ are unspecified, we might have to erase them afterwards:

Lemma 5.5. *The program $(i, \text{empty\_stack})$ with*

$$\text{empty\_stack} := [\text{POP } x \ i \ (3 + i); \text{PUSH } x \ 0; \text{POP } x \ i \ i]$$

composed of 3 instructions satisfies the specification

$$\forall v, (i, \text{empty\_stack}) \parallel (i, v) >^* (3 + i, v[x := []]).$$

This solves part (3) of pcp_bsm_R . To complete the reduction of an iBPCP instance to a BSM instance we have to prove soundness and completeness of the pcp_bsm_R program. We do this by combining the previous results using the tools of compositional reasoning described in Section 4.3.

Theorem 5.6 (iBPCP simulator). *For any configuration v and any state s , the two following properties hold:*

$$\begin{aligned} \text{iBPCP } R \rightarrow (1, \text{pcp\_bsm}_R) \parallel (1, v) >^* s_0 & \quad (\text{sound}) \\ (1, \text{pcp\_bsm}_R) \parallel (1, v) \rightsquigarrow s \rightarrow \text{iBPCP } R \wedge s = s_0 & \quad (\text{complete}) \end{aligned}$$

where $s_0 = (0, v_0)$ and $v_0 = [s, u, l, a := []]$.

As a side remark, we show that the length of pcp_bsm_R is $86 + 3 \cdot |R| + \text{size } R$ where $\text{size } R$ is the length of all concatenated upper and lower traces in R . This allows us to conclude:

Theorem 5.7. $\text{iBPCP} \leq \text{BSM}$.

6 A Certified Low-Level Compiler

We present the implementation of a certified compiler for low-level languages as defined in Section 4, to be used in later Section 7.3 to reduce BSM-termination to MM-termination.

Let us start with a description of the problem. We are given two models of computation X and Y as presented in Section 4, each with their own set of instructions and their own semantics described by step relations. We adapt the notation of Section 4 to discuss about those two different models at the same time by decorating the notations with either X or Y , hence $\mathcal{P} \parallel_X s_1 \succ^3 s_2$ means program $\mathcal{P}$ of model X can compute state s_2 in 3 steps when starting from state s_1 , while $Q \parallel_Y s \downarrow$ means that the computation of program Q of model Y terminates when starting from state s . The model X will be the source while Y will be the target model of the compiler.

The compiler consists of transforming any program $\mathcal{P}$ in model X into a program Q in model Y such that Q simulates the behaviour of $\mathcal{P}$ faithfully. We model the correspondence between the configurations of machines in model X and in model Y by a *simulation* binary relation $\bowtie : \mathbb{C}_X \rightarrow \mathbb{C}_Y \rightarrow \mathbb{P}$.

A typical soundness result would be

$$\begin{aligned} v_1 \bowtie w_1 \wedge \mathcal{P} \parallel_X (i_1, v_1) >^* (i_2, v_2) \\ \rightarrow \exists (j_2, w_2), v_2 \bowtie w_2 \wedge \mathcal{Q} \parallel_Y (j_1, w_1) >^* (j_2, w_2) \end{aligned}$$

meaning that this particular computation of $\mathcal{P}$ from (i_1, v_1) is simulated by some computation of $\mathcal{Q}$ from (j_1, w_1) preserving the simulation relation between the configurations in X and Y . However, such a result only expresses a relation between configurations, not between PC values. The relation between PC values is going to be made more precise using a linker.

As it is generic, our compiler cannot cope with the particular behaviour of individual instructions because that really depends on the choice of X and Y : for instance, there cannot exist a correct compiler if model X is more expressive than model Y . As a result, the generic compiler assumes as input a *one-to-many instruction compiler* that transforms a single instruction in $\mathbb{I}_X$ into a list of instructions in $\mathbb{L}\mathbb{I}_Y$. The problem solved by our generic compiler is transforming this individual instruction compiler into a program compiler while preserving soundness.

Because the replacement of X instructions by Y instructions is one-to-many, absolute jumps must be relocated at addresses that depend on the whole program, not only on the instruction that is being compiled. For instance, consider the absolute jumps p and q embedded in the BSM-instruction $\text{POP } \alpha \ p \ q$ which is compiled in 16 MM-instructions in Section 7.3. These PC addresses p/q must be recomputed before compiling the $\text{POP } \alpha \ p \ q$ instruction. But their new values depend on the content of the compiled program *ahead of the instructions they point at*. If they point forward, a chicken-and-egg problem appears. To solve it, traditional compilers replace p/q with symbolic addresses and the whole program is compiled in a symbolic address space of graph-like structure. Symbols are resolved by a *linker* in a later phase that maps the graph of instructions into a sequence again.

Here, we can avoid this quite complex machinery by assuming a specific feature of our instruction compiler: even though the list $ll_\rho : \mathbb{L}\mathbb{I}_Y$ of Y -instructions corresponding to some X -instruction $\rho : \mathbb{I}_X$ might depend on where ρ is located in the compiled code (and also on where the absolute and relative² jumps which ρ contains are remapped), the *length* $|ll_\rho|$ *does not depend on those actual positions*. So there is a first abstract compilation phase that only computes the length of codes, from which we can deduce a linker that is going to be compatible with a later concrete compilation phase that computes actual instructions at their definitive location as given by the linker.

²Although instructions might not refer to specific relative addresses explicitly, the default incrementation of the PC by one makes +1 an almost ubiquitous relative address.

Let us now describe the compiler more formally. We assume an instruction compiler and a code length function:

$$\begin{aligned} \text{icomp} &: (\mathbb{N} \rightarrow \mathbb{N}) \rightarrow \mathbb{N} \rightarrow \mathbb{I}_X \rightarrow \mathbb{L}\mathbb{I}_Y \\ \text{ilen} &: \mathbb{I}_X \rightarrow \mathbb{N} \end{aligned}$$

where $\text{icomp } \text{lnk } i \ \rho$ compiles instruction ρ into $ll_\rho : \mathbb{L}\mathbb{I}_Y$ at source position i using linker lnk to remap absolute and relative jumps, and the length of ll_ρ is given by $\text{ilen } \rho$. Hence, we assume the hypothesis

$$H_{\text{ilen}} : \forall \text{lnk } i \ \rho, |\text{icomp } \text{lnk } i \ \rho| = \text{ilen } \rho$$

so that the length of compiled instructions only depends on instruction ρ , not where it is positioned (i.e. i) or where jumps are remapped (i.e. lnk). In contrast, the actual code $\text{icomp } \text{lnk } i \ \rho$ almost always depends on lnk and i .

Let us now come to the semantic soundness hypothesis on icomp . For this, a step relation for X and a step relation for Y must be given. For simplicity of arguments, we assume that the step relation for X is total and that the step relation for Y is deterministic (i.e. functional).

H_X : the step relation $>_X$ is total

H_Y : the step relation $>_Y$ is deterministic

We do not view totality as a strong requirement. It could be noted however that determinism would not be satisfied in case the target code involves some form of parallelism.

Given a simulation relation $\bowtie : \mathbb{C}_X \rightarrow \mathbb{C}_Y \rightarrow \mathbb{P}$ between the configurations of X and Y , the critical soundness assumption we make on icomp is the following:

$$\begin{aligned} H_{\text{icomp}} &: \forall \text{lnk } \rho (i_1, v_1) (i_2, v_2) \ w_1, \\ &\quad \rho \parallel_X (i_1, v_1) > (i_2, v_2) \\ &\quad \wedge v_1 \bowtie w_1 \wedge \text{lnk}(1 + i_1) = \text{ilen } \rho + \text{lnk}(i_1) \\ &\rightarrow \exists w_2, (\text{lnk } i_1, ll_\rho) \parallel_Y (\text{lnk } i_1, w_1) >^+ (\text{lnk } i_2, w_2) \\ &\quad \wedge v_2 \bowtie w_2 \end{aligned}$$

where ll_ρ is shorthand for $ll_\rho := \text{icomp } \text{lnk } i_1 \ \rho$. It requires some explanations: H_{icomp} states that if ll_ρ is the result of the compilation of instruction ρ at source position i_1 using linker lnk , if the X -step relation for ρ transforms state (i_1, v_1) into state (i_2, v_2) and if w_1 simulates v_1 , then the compiled code ll_ρ at its target address as given by lnk computes some state $(\text{lnk } i_2, w_2)$ from state $(\text{lnk } i_1, w_1)$ in the Y semantics and w_2 simulates v_2 . The simulated computation must make progress, hence ll_ρ can never be empty. Additionally, the pre-condition $\text{lnk}(1 + i_1) = \text{ilen } \rho + \text{lnk}(i_1)$ assumes that the linker lnk preserves the +1 (implicit) relative address.

These assumptions, in particular H_{icomp} which is the most important, are sufficient to build our program compiler and show its correctness, of which the proof is decomposed in two phases: a semantic one and a syntactic one.

6.1 Semantic Correctness of Compiled Code

Under the previous assumptions, suppose that we are given a source X -program $\mathcal{P}$, a linker lnk and a compiled Y -program

Q . We give a syntactic correctness criterion which is sufficient to establish that Q correctly simulates $\mathcal{P}$. The criterion

$$H_{\mathcal{P}Q} : \forall i \rho, (i, [\rho]) <_{sc} \mathcal{P} \rightarrow \wedge \left\{ \begin{array}{l} (\text{lnk } i, \text{icomp } \text{lnk } i \rho) <_{sc} Q \\ \text{lnk}(1 + i) = \text{ilen } \rho + \text{lnk}(i) \end{array} \right.$$

means that whenever instruction ρ occurs in $\mathcal{P}$ at position i then its corresponding compiled code $\text{icomp } \text{lnk } i \rho$ occurs at position $\text{lnk } i$ in Q and lnk respects the $1 + i$ relative address. This condition does not involve the semantics chosen for X and Y at all, it is a purely syntactic assumption which, in combination with H_{icomp} , is sufficient to ensure that Q correctly simulates $\mathcal{P}$.

Theorem 6.1. *Under assumptions $H_{\text{icomp}}, H_{\text{ilen}}, H_X, H_Y$ and $H_{\mathcal{P}Q}$, the program Q soundly simulates $\mathcal{P}$, that is for any $(i_1, v_1), (i_2, v_2)$ and w_1 , the following property holds:*

$$\begin{aligned} v_1 \bowtie w_1 \wedge \mathcal{P} \parallel_X (i_1, v_1) >^* (i_2, v_2) \\ \rightarrow \exists w_2, v_2 \bowtie w_2 \wedge Q \parallel_Y (\text{lnk } i_1, w_1) >^* (\text{lnk } i_2, w_2) \end{aligned}$$

Proof. By induction on the computation of $\mathcal{P}$. $\square$

The completeness of Q is more complicated to express and prove because *not every computation of Q corresponds to a computation of $\mathcal{P}$* . To ensure that property, we must restrict the computations of Q to those which start at some $\text{lnk } i$ (obviously) but we must also assume termination using big-step semantics.

Theorem 6.2. *Under assumptions $H_{\text{icomp}}, H_{\text{ilen}}, H_X, H_Y$ and $H_{\mathcal{P}Q}$, the program Q completely simulates $\mathcal{P}$, that is for any (i_1, v_1) and w_1 , the following property holds:*

$$v_1 \bowtie w_1 \wedge Q \parallel_Y (\text{lnk } i_1, w_1) \downarrow \rightarrow \mathcal{P} \parallel_X (i_1, v_1) \downarrow$$

Proof. By complete induction on the length q of the terminating computation $Q \parallel_Y (\text{lnk } i_1, w_1) \rightsquigarrow^q s$. $\square$

Hence, to get a correct (sound and complete) compiler from a sound instruction compiler, it is enough to produce code that satisfies $H_{\mathcal{P}Q}$ which is a purely syntactic criterion.

6.2 A Syntactically Correct Compiler

Obtaining lnk and Q satisfying condition $H_{\mathcal{P}Q}$ from $\mathcal{P}$ is not completely trivial because as said earlier, the linker lnk which represents the remapping of the address space of the source program $\mathcal{P}$, depends on the length of compiled code (hence Q) and $\mathcal{P}$ needs lnk to be compiled correctly. The circularity is broken under the H_{ilen} assumption that isolates code length from target address space considerations.

More formally, given source code $\mathcal{P} = (i, [\rho_0; \dots; \rho_{n-1}])$ and a destination address j where the compiled code is supposed to start, we first build lnk such that

$$\text{lnk}(k + i) = j + \text{ilen } \rho_0 + \dots + \text{ilen } \rho_{k-1} \quad \text{for } k \in [0, n]$$

and then the compiled code Q as

$$Q := (j, \text{icomp } \text{lnk } i \rho_0 \# \dots \# \text{icomp } \text{lnk } (n - 1 + i) \rho_{n-1}).$$

There are some subtleties of secondary importance related to where source PC values outside of $\mathcal{P}$ should be remapped.³ For our purpose, we simply map them all to $\text{lnk}(n + i)$. This simple linker/compiler satisfies the syntactic correctness criterion $H_{\mathcal{P}Q}$ and we deduce the following correction theorem.

Theorem 6.3 (Compiler). *Assuming $H_{\text{icomp}}, H_{\text{ilen}}, H_X, H_Y$, for any X -program $\mathcal{P}$ and any destination address j , one can compute a linker $\text{lnk} : \mathbb{N} \rightarrow \mathbb{N}$ and a Y -program Q with the following properties:*

1. $j = \text{start } Q = \text{lnk}(\text{start } \mathcal{P})$;
2. *for any i , out $i \mathcal{P} \rightarrow \text{lnk } i = \text{end } Q$;*
3. *the following holds for any i_1, v_1, i_2, v_2, w_1 :*

$$v_1 \bowtie w_1 \wedge \mathcal{P} \parallel_X (i_1, v_1) \rightsquigarrow (i_2, v_2)$$

$$\rightarrow \exists w_2, v_2 \bowtie w_2 \wedge Q \parallel_Y (\text{lnk } i_1, w_1) \rightsquigarrow (\text{lnk } i_2, w_2)$$

4. *the following holds for any i_1, v_1, w_1, j_2, w_2 :*

$$v_1 \bowtie w_1 \wedge Q \parallel_Y (\text{lnk } i_1, w_1) \rightsquigarrow (j_2, w_2)$$

$$\rightarrow \exists i_2 v_2, v_2 \bowtie w_2 \wedge \mathcal{P} \parallel_X (i_1, v_1) \rightsquigarrow (i_2, v_2) \wedge j_2 = \text{lnk } i_2.$$

This concludes our implementation of a generic and certified compiler, as a functor that maps a sound instruction compiler into a correct program compiler.

6.3 Discussion

Our generic compiler is certainly not the only certified compiler around, CakeML [23] (ML compiler to assembly, certified in HOL) and CompCert [28] (C compiler to assembly, certified in Coq) being among the most famous ones. We make some comparisons with two Coq-certified compilers [6, 28]. Unfortunately for our purpose, these two compilers both involve specific source and target languages which are of little use for the BSM-MM reduction we need. These compilers are more ambitious than our intended BSM-MM compiler: source and target languages are more complex both syntactically and semantically. In the case of CompCert [28], there is also an aim at certifying most of the compilation tool-chain, up to near state-of-the-art optimisations.

Hence, we do not claim to have solved a very difficult open problem with our generic compiler. Nonetheless, we want to point out two of its main specificities. First, we have identified a small set of assumptions that suffice at proving correctness of the produced compiled code. We hope that these could serve as guidelines for other developments to come. More critically in our setting which consists in building reductions between undecidable problems, our compiler comes not only with a soundness theorem 6.1 comparable to those of [6, 28], but it also comes with a proof of *completeness* (Theorem 6.2) stating that termination of compiled code implies termination of source code.

In the cases of [6] or CompCert [28], this might be a non-issue since the intended use is only for terminating programs.

³Depending of the intended application, it could be useful to distinguish the remappings of source PC addresses outside of $\mathcal{P}$.

To the best of our understanding, this question of completeness does not seem to have been considered by any of these projects. However, to establish undecidability results — all the more in a constructive setting, — it is not possible to assume termination and thus we really need a completeness result to show the correctness of compiler-based reductions.

7 Minsky Machines

Minsky Machines (MM) [32] have n registers with natural number values. Similar to BSM configurations, we define MM configurations as vectors of registers values, i.e. $\mathbb{C} := \mathbb{N}^n$. Instructions are defined by

$$\mathbb{I}_n ::= \text{INC } (\alpha : \mathbb{F}_n) \mid \text{DEC } (\alpha : \mathbb{F}_n) (p : \mathbb{N}).$$

7.1 Semantics for MM

The rules defining the step relation for MM instructions are:

$$\begin{aligned} \text{INC } \alpha &\parallel (i, v) > (1 + i, v[\alpha := 1 + n]) && \text{if } v[\alpha] = n \\ \text{DEC } \alpha p &\parallel (i, v) > (p, v) && \text{if } v[\alpha] = 0 \\ \text{DEC } \alpha p &\parallel (i, v) > (1 + i, v[\alpha := n]) && \text{if } v[\alpha] = 1 + n \end{aligned}$$

That means, $\text{INC } \alpha$ increases the value of register α by one. $\text{DEC } \alpha p$ decreases the value of register α by one if that is possible and increases the PC, or, if the register is already 0, jumps to PC value p . As is the case for BSM, the step relation for MM is both deterministic and total.

Finally, we use a more restricted halting problem for MM than for BSM, where we start execution with PC at 1 and expect the machine to halt with PC at 0 and on a null register configuration:

$$\text{MM}(n : \mathbb{N}, P : \mathbb{L} \mathbb{I}_n, v : \mathbb{N}^n) := (1, P) \parallel (1, v) \rightsquigarrow (0, \vec{0}).$$

This restricted halting problem is well-suited for a reduction to linear logic. We could use a more general halting problem and then reduce this more general problem to the special problem. However, this would require a dedicated transformation on MM programs, which we circumvent by reducing BSM to the special version directly.

7.2 Simulating Binary Stacks with Numbers

The reduction from BSM to MM is derived as a direct consequence of our implementation of a generic certified low-level compiler. Basically to get a correct compiler from BSM to MM, the only remaining tasks are: (a) simulate binary stacks with natural numbers, and (b) simulate individual BSM-instructions with lists of MM-instructions. We give a high-level overview of the principal ideas and the main soundness and completeness results.

BSM programs manipulate binary stacks in $\mathbb{L}\mathbb{B}$ while MM programs manipulate natural numbers in $\mathbb{N}$. Hence to represent $\mathbb{L}\mathbb{B}$ with $\mathbb{N}$, we implement multiplication and division by 2. For instance, we show the following fact.

Fact 7.1. *For any $n : \mathbb{N}$, $x \neq q \neq r : \mathbb{F}_n$ and $i : \mathbb{N}$, the list mm_div2 of 6 MM-instructions*

[DEC $x (6 + i)$; INC r ; DEC $x (6 + i)$; DEC $r (4 + i)$; INC q ; DEC $r i$] implements division by 2 because it satisfies the following specification for all k, b, v :

$$\begin{aligned} &v[q] = 0 \wedge v[r] = 0 \wedge v[x] = b + 2 \cdot k \\ &\rightarrow (i, \text{mm\_div2}) \parallel (i, v) >^+ (6 + i, v[x := 0, q := k, r := b]) \end{aligned}$$

Note that $b : \mathbb{B}$ is cast into $0/1 : \mathbb{N}$ to simplify the above statement. We also need transfer (resp. nullify) programs which adds the value of register x to y , nullifying x at the same time (resp. nullifies one or several registers). These programs are (somewhat) obvious to write and prove correct.

We can then encode a binary stack $l : \mathbb{L}\mathbb{B}$ as a natural number $\text{s2n } l : \mathbb{N}$ using:

$$\text{s2n } [] := 1 \quad \text{s2n } (0 :: l) := 2 \cdot \text{s2n } l \quad \text{s2n } (1 :: l) := 1 + 2 \cdot \text{s2n } l$$

and simulate $\text{PUSH } x 0$, $\text{PUSH } x 1$ and $\text{POP } x j e$ instructions.

Lemma 7.2. *For any $n : \mathbb{N}$, $x \neq z \neq t : \mathbb{F}_n$ and $i : \mathbb{N}$, there is a list mm_push_0 of 7 instructions which satisfies the following specification for all l, v :*

$$\begin{aligned} &v[t] = 0 \wedge v[z] = 0 \wedge v[x] = \text{s2n } l \\ &\rightarrow (i, \text{mm\_push\_0}) \parallel (i, v) >^+ (7 + i, v[x := \text{s2n}(0 :: l)]) \end{aligned}$$

Lemma 7.3. *For any $n : \mathbb{N}$, $x \neq z \neq t : \mathbb{F}_n$ and $i : \mathbb{N}$, there is a list mm_push_1 of 8 instructions which satisfies the following specification for all l, v :*

$$\begin{aligned} &v[t] = 0 \wedge v[z] = 0 \wedge v[x] = \text{s2n } l \\ &\rightarrow (i, \text{mm\_push\_1}) \parallel (i, v) >^+ (8 + i, v[x := \text{s2n}(1 :: l)]) \end{aligned}$$

Lemma 7.4. *For any $n : \mathbb{N}$, $x \neq z \neq t : \mathbb{F}_n$ and $i, j, k, e : \mathbb{N}$, there is a list mm_pop of 16 instructions which satisfies the following specifications for any l and v s.t. $v[t] = v[z] = 0$:*

$$\begin{aligned} &v[x] = \text{s2n } [] \rightarrow (i, \text{mm\_pop}) \parallel (i, v) >^+ (e, v) \\ &v[x] = \text{s2n}(0 :: l) \rightarrow (i, \text{mm\_pop}) \parallel (i, v) >^+ (j, v[x := \text{s2n } l]) \\ &v[x] = \text{s2n}(1 :: l) \rightarrow (i, \text{mm\_pop}) \parallel (i, v) >^+ (k, v[x := \text{s2n } l]) \end{aligned}$$

Hence, we simulate any individual BSM instruction by a (small) list of MM instructions that operate on stacks encoded using s2n . Using Theorem 6.3 of Section 6, we obtain a correct compiler for whole BSM programs to MM programs.

7.3 BSM reduces to MM

We describe the correctness and completeness results we obtain from the generic compiler of Theorem 6.3. Assume that we are given a BSM program (i, P) with n binary stack registers, we compile (i, P) at address 1 into a linker $\text{lnk} : \mathbb{N} \rightarrow \mathbb{N}$ and a MM program Q starting at address 1 with $2 + n$ registers. Each stack register $p : \mathbb{F}_n$ is s2n encoded into the register $2 + p : \mathbb{F}_{2+n}$ of Q . The registers $0, 1 : \mathbb{F}_{2+n}$ of Q serve as spare registers for internal computations. The simulation invariant

$$v \bowtie w := w[0] = 0 \wedge w[1] = 0 \wedge \forall p, w[2 + p] = \text{s2n}(v[p])$$

ensures the preservation of the encoding during execution. Ink satisfies $\text{Ink } i = 1$ and for any j for which $\text{out } j (i, P)$ holds, we have $\text{Ink } j = \text{end}(1, Q)$. Moreover, the code $(1, Q)$ satisfies the following soundness and completeness results:

Fact 7.5. *For any i_1, i_2, v_1, v_2, w_1 , the following holds:*

$$\begin{aligned} v_1 \bowtie w_1 \wedge (i, P) \parallel (i_1, v_1) \rightsquigarrow (i_2, v_2) \\ \rightarrow \exists w_2, v_2 \bowtie w_2 \wedge (1, Q) \parallel (\text{Ink } i_1, w_1) \rightsquigarrow (\text{Ink } i_2, w_2). \end{aligned}$$

Fact 7.6. *For any i_1, v_1, w_1, j_2, w_2 , the following holds:*

$$\begin{aligned} v_1 \bowtie w_1 \wedge (1, Q) \parallel_Y (\text{Ink } i_1, w_1) \rightsquigarrow (j_2, w_2) \\ \rightarrow \exists i_2, v_2, v_2 \bowtie w_2 \wedge (i, P) \parallel_X (i_1, v_1) \rightsquigarrow (i_2, v_2). \end{aligned}$$

From those correctness results, it is easy to derive a simulator for BSM termination as MM termination.

Theorem 7.7 (BSM simulator). *Let (i, P) be a BSM-program with n binary stacks. One can compute a $2 + n$ registers MM-program bsm_mm with the following specification. For any v , with $w := [s2n \mid l \in v]$, we have:*

$$(i, P) \parallel (i, v) \downarrow \leftrightarrow (1, \text{bsm\_mm}) \parallel (1, 0 :: 0 :: w) \rightsquigarrow (0, \vec{0}).$$

Proof. The code of bsm_mm consists of the compiled code $(1, Q)$ of (i, P) followed by some obvious code nullifying all the variables after the execution of $(1, Q)$ is terminated and then jumping to PC value 0. The nullifying code uses the fact that any output state (j_2, w_2) of $(1, Q)$ satisfies $j_2 = \text{end}(1, Q)$ and $w_2[0] = 0$ to nullify the other registers using register 0 as empty register. We critically need both the soundness Fact 7.5 and completeness Fact 7.6 to get the above equivalence. $\square$

This allows us to conclude:

Theorem 7.8. $\text{BSM} \leq \text{MM}$.

8 Intuitionistic Linear Logic

$$\begin{array}{c} \frac{}{A \vdash A} \quad \frac{\Gamma \vdash A \quad \Delta, B \vdash C}{\Gamma, \Delta, A \multimap B \vdash C} \quad \frac{\Gamma, A \vdash B}{\Gamma \vdash A \multimap B} \\ \frac{\Gamma, A \vdash C}{\Gamma, A \& B \vdash C} \quad \frac{\Gamma, B \vdash C}{\Gamma, A \& B \vdash C} \quad \frac{\Gamma \vdash A \quad \Gamma \vdash B}{\Gamma \vdash A \& B} \\ \frac{\Gamma, A \vdash B}{\Gamma, !A \vdash B} \quad \frac{\Gamma \vdash B}{\Gamma, !A \vdash B} \quad \frac{\Gamma, !A, !A \vdash B}{\Gamma, !A \vdash B} \quad \frac{! \Gamma \vdash B}{! \Gamma \vdash !B} \end{array}$$

Figure 1. The S_{ILL} sequent calculus.

The reduction of MM to entailment in *Intuitionistic Linear Logic (ILL)* [16] implements the method presented in [26] and fully described in [27]. The undecidability of (I)LL was first established in [30] and a rich literature followed [20, 21, 29], but [26] singled out the elementary fragment of ILL (eILL) together with a short proof of undecidability by reducing the termination of two counters Minsky machines.

Here we present the reduction of termination of arbitrary MM (with n counters) and we verify it in Coq. The proof consists first of the reduction of MM to entailment in eILL, and then the proof of the (faithful) embedding of eILL into ILL. The latter proof actually takes place in the $\{\&, \multimap, !\}$ cut-free fragment of ILL, i.e. we do not need the $\{\otimes, \oplus, \perp, \top, 1\}$ logical connectives nor the cut rule to build the embedding.

As a way to limit the technical material we need to present, we only introduce the $\{\&, \multimap, !\}$ fragment in this paper. However, the Coq code contains the proof of the embedding of eILL into the whole ILL.⁴ Hence from now on, we will abusively denote ILL for its restriction to the $\{\&, \multimap, !\}$ fragment.

The syntax of (the $\{\&, \multimap, !\}$ fragment of) ILL is described by the BNF-grammar $A ::= v \mid A \& A \mid A \multimap A \mid !A$ where v belongs to a *countably infinite* set of logical variables. A *sequent* is a pair denoted $\Gamma \vdash A$ composed of a multiset Γ of formulae and a single formula A on the right of the $\vdash$ sign. The *sequent calculus* S_{ILL} is described in Fig. 1. In the last rule, $! \Gamma$ denotes $! \Gamma := !A_1, \dots, !A_n$ when $\Gamma = A_1, \dots, A_n$. Using S_{ILL} , we define the entailment relation: the hypotheses in Γ entail A when $\Gamma \vdash A$ has a derivation in S_{ILL} .

8.1 eILL reduces to ILL

ILL contains a sub-fragment called *elementary Intuitionistic Linear Logic (eILL)* first spotted in [26]. It is not a syntactic fragment at the level of formulae, but it is a syntactic fragment at the level of sequents. A sequent in eILL is a sequent of ILL of the form $! \Sigma, \Gamma \vdash u$, where Γ is a multiset composed of logical variables only and u is a unique logical variable. On the contrary, Σ is composed of formulae called *command formulae* [27] of the form $(u \multimap v) \multimap w$, $u \multimap (v \multimap w)$ and $(u \& v) \multimap w$ where u, v and w are logical variables. So eILL sequents have a very simple/flat syntactic structure.

$$\begin{array}{c} \frac{}{! \Sigma, u \vdash u} \text{Ax} \quad \frac{! \Sigma, \Gamma, u \vdash v}{! \Sigma, \Gamma \vdash w} \quad (u \multimap v) \multimap w \in \Sigma \\ \frac{! \Sigma, \Gamma \vdash u \quad ! \Sigma, \Delta \vdash v}{! \Sigma, \Gamma, \Delta \vdash w} \quad u \multimap (v \multimap w) \in \Sigma \\ \frac{! \Sigma, \Gamma \vdash u \quad ! \Sigma, \Gamma \vdash v}{! \Sigma, \Gamma \vdash w} \quad (u \& v) \multimap w \in \Sigma \end{array}$$

Figure 2. The G_{eILL} goal-directed sequent calculus.

Although entailment in eILL could be described by S_{ILL} , for the purpose of reducing MM, it is more conveniently described by the goal directed rules of G_{eILL} in Fig. 2. The reason is that they directly encode MM small-step semantics as shall be explained in Section 8.2. Notice however that

⁴This means that our method proves the embedding of eILL into both the $\{\&, \multimap, !\}$ fragment and the whole ILL. However, we do not show or formalise the embedding of the fragment into the whole: it involves a proof of either cut-elimination, or a larger phase semantics development.

while sequents of eILL are stable under (backwards application of) the rules of G_{eILL} , they are not stable under the rules of S_{ILL} , which means that proof-search from an eILL sequent in S_{ILL} might produce sequents which do not belong to eILL.

Given a commutative monoid $(M, \bullet, \epsilon)$, for any $X, Y \subseteq M$, we define the *point-wise extension* by $X \bullet Y := \{x \bullet y \mid x \in X \wedge y \in Y\}$ and its *linear adjunct* as $X \multimap Y := \{k \in M \mid \{k\} \bullet X \subseteq Y\}$ providing a residuated monoidal structure on the powerset $M \rightarrow \mathbb{P}$. Given any semantic interpretation $[[\cdot]] \subseteq M$ of logical variables, we extend it inductively to ILL via *trivial phase semantics*:

$$\begin{aligned} [[A \& B]] &:= [[A]] \cap [[B]] & [[!A]] &:= \{\epsilon\} \cap [[A]] \\ [[A \multimap B]] &:= [[A]] \multimap [[B]] \end{aligned}$$

and to multisets by $[[A_1, \dots, A_n]] := [[A_1]] \bullet \dots \bullet [[A_n]]$. An ILL sequent $\Gamma \vdash A$ is *valid in that interpretation* if $[[\Gamma]] \subseteq [[A]]$. A sequent is *valid in a commutative monoid* $(M, \bullet, \epsilon)$ if it is valid for any trivial interpretation $[[\cdot]] \subseteq M$ in that monoid.

Theorem 8.1. *The three following results hold:*

1. *If derivable in G_{eILL} , $! \Sigma, \Gamma \vdash u$ is also derivable in S_{ILL} ;*
2. *An ILL sequent derivable in S_{ILL} is valid in every commutative monoid;*
3. *Let n be greater than the number of logical variables of Σ, Γ . If $! \Sigma, \Gamma \vdash w$ is valid in the free commutative monoid $(\mathbb{N}^n, +, \vec{0})$ then $! \Sigma, \Gamma \vdash w$ is provable in G_{eILL} .*

Notice that trivial phase semantics is also sound for the whole ILL, not just the $(\&, \multimap, !)$ -fragment ILL or eILL. But it is only complete for eILL, not for the $(\&, \multimap, !)$ -fragment nor for the whole ILL because the interpretation of the bang $!$ is too restricted compared to general phase semantics [27].

Theorem 8.2. $eILL \leq ILL$.

Proof. We simply need to show that a sequent of eILL has a proof in G_{eILL} if and only if it has a proof in S_{ILL} and this is a direct consequence of Theorem 8.1. $\square$

8.2 MM reduces to eILL

We show how to faithfully encode MM-termination by entailment in G_{eILL} . For this, we describe how to transform a n registers MM-program P and a vector $v : \mathbb{N}^n$ into an eILL-sequent, s.t. the G_{eILL} -entailment of the sequent simulates the predicate $(1, P) \parallel (1, v) \rightsquigarrow (0, \vec{0})$.

We fix n and $P = [\rho_1; \dots; \rho_l]$. Let us partition the countable infinite set of logical variables into disjoint sets

$$\{x_0, \dots, x_{n-1}\} \uplus \{\bar{x}_0, \dots, \bar{x}_{n-1}\} \uplus \{q_0, q_1, \dots\}.$$

Let us describe the eILL-sequent $! \Sigma, \Gamma_v \vdash q_1$ simulating $(1, P) \parallel (1, v) \rightsquigarrow (0, \vec{0})$. We start with the commands in Σ . For every labelled instruction $i : \rho_i$ of $(1, P)$, we add one or two commands in Σ according to:

$$\begin{aligned} i : \text{INC } p &\rightsquigarrow (x_p \multimap q_{i+1}) \multimap q_i \\ i : \text{DEC } p j &\rightsquigarrow x_p \multimap (q_{i+1} \multimap q_i) \text{ and } (\bar{x}_p \& q_j) \multimap q_i. \end{aligned}$$

In addition, we add the following $1 + n^2$ commands to Σ independently of P . These just depend on n :

$$\begin{aligned} \{(q_0 \multimap q_0) \multimap q_0\} \uplus \{(\bar{x}_i \multimap \bar{x}_i) \multimap \bar{x}_i \mid i < n\} \\ \uplus \{x_j \multimap (\bar{x}_i \multimap \bar{x}_i) \mid i \neq j < n\}. \end{aligned}$$

We define the multiset $\Gamma_v := \{v[0].x_0, \dots, v[n-1].x_{n-1}\}$ where $p.x_i$ means repeating p times x_i in the multiset. We show the following reduction.

Lemma 8.3. *We have $(1, P) \parallel (i, v) \rightsquigarrow (0, \vec{0})$ if and only if the sequent $! \Sigma, \Gamma_v \vdash q_i$ is provable in G_{eILL} .*

Proof. We show the *only if* part by structural induction on the predicate $(1, P) \parallel (i, v) \succ^* (0, \vec{0})$. In the base case we have $i = 0$ and $v = \vec{0}$ and we check $! \Sigma, \emptyset \vdash q_0$:

$$\frac{\frac{\text{Ax}}{! \Sigma, q_0 \vdash q_0}}{! \Sigma, \emptyset \vdash q_0} (q_0 \multimap q_0) \multimap q_0 \in \Sigma$$

If the last rule is

$$\frac{\text{INC } p \parallel (i, v) \succ s \quad (1, P) \parallel s \succ^* (0, \vec{0})}{(i, P) \parallel (i, v) \succ^* (0, \vec{0})}$$

then $s = (i + 1, v + 1.x_p)$ and we build the following proof

$$\frac{\dots \quad \frac{! \Sigma, \Gamma_v, x_p \vdash q_{i+1}}{! \Sigma, \Gamma_v \vdash q_i} (x_p \multimap q_{i+1}) \multimap q_i \in \Sigma}{! \Sigma, \Gamma_v \vdash q_i}$$

where the top dots symbolize a sub-proof corresponding to the induction hypothesis. Notice that $(x_p \multimap q_{i+1}) \multimap q_i$ was explicitly added to Σ to encode instruction $i : \text{INC } p$.

We do the same for $i : \text{DEC } p j$. If the last rule is

$$\frac{\text{DEC } p j \parallel (i, v) \succ s \quad (1, P) \parallel s \succ^* (0, \vec{0})}{(i, P) \parallel (i, v) \succ^* (0, \vec{0})}$$

then there are two cases depending on whether the value of v at coordinate p is 0 or not. If $v[p] \neq 0$ then $v = w + 1.x_p$, $\Gamma_v = \Gamma_w, x_p$ and $s = (i + 1, w)$ and we build the following proof using the induction hypothesis for the top-right dots:

$$\frac{\frac{\text{Ax}}{! \Sigma, x_p \vdash x_p} \quad \dots \quad \frac{! \Sigma, \Gamma_w \vdash q_{i+1}}{! \Sigma, \Gamma_w, x_p \vdash q_i} x_p \multimap (q_{i+1} \multimap q_i) \in \Sigma}{! \Sigma, \Gamma_w, x_p \vdash q_i}$$

The case where $v[p] = 0$ is a bit more complicated. In that case, $x_p \notin \Gamma_v$ and we have $s = (j, v)$. We build the proof as:

$$\frac{\dots \quad \frac{! \Sigma, \Gamma_v \vdash \bar{x}_p}{! \Sigma, \Gamma_v \vdash q_i} x_p \notin \Gamma_v \quad \dots \quad \frac{! \Sigma, \Gamma_v \vdash q_j}{! \Sigma, \Gamma_v \vdash q_i} (\bar{x}_p \& q_j) \multimap q_i \in \Sigma}{! \Sigma, \Gamma_v \vdash q_i}$$

where the top-right dots represent the induction hypothesis. For the top-left dots, we build a proof that can only succeed

when Γ_v does not contains any occurrence of x_p , implementing a zero-test by proof-search. For $q \neq p < n$, we can remove the occurrence of x_q in Δ, x_q using $\bar{x}_p$ as a goal:

$$\frac{\frac{\text{Ax}}{! \Sigma, x_q \vdash x_q} \quad \frac{\dots}{! \Sigma, \Delta \vdash \bar{x}_p}}{! \Sigma, \Delta, x_q \vdash \bar{x}_p} x_q \multimap (\bar{x}_p \multimap \bar{x}_p) \in \Sigma$$

As Γ_v is only composed of variables x_q with $q \neq p$, we repeat the previous step until Γ_v becomes empty and we end up with $! \Sigma, \emptyset \vdash \bar{x}_p$ which is then proved by

$$\frac{\frac{\text{Ax}}{! \Sigma, \bar{x}_p \vdash \bar{x}_p}}{! \Sigma, \emptyset \vdash \bar{x}_p} (\bar{x}_p \multimap \bar{x}_p) \multimap \bar{x}_p \in \Sigma$$

This conclude the *only if part* of the argument.

The *if part* is obtained using soundness of trivial phase semantics (see Theorem 8.1). To simplify notations, we denote by $(\vec{x}_0, \dots, \vec{x}_{n-1})$ the canonical basis of $(\mathbb{N}^n, +, \vec{0})$, i.e. we write $\mathbb{N}^n = \mathbb{N}.\vec{x}_0 \oplus \dots \oplus \mathbb{N}.\vec{x}_{n-1}$ and the correspondence $x_i \rightleftharpoons \vec{x}_i$ induces a monoid isomorphism between the multisets generated by $\{x_0, \dots, x_{n-1}\}$ and $(\mathbb{N}^n, +, \vec{0})$. We define the following interpretation of logical variables in $(\mathbb{N}^n, +, \vec{0})$:

$$\begin{aligned} \llbracket x_i \rrbracket &:= \{\vec{x}_i\} & \llbracket \bar{x}_i \rrbracket &:= \{\vec{x}_i\}^\perp = \{w \mid w[i] = 0\} \\ \llbracket q_i \rrbracket &:= \{w \in \mathbb{N}^n \mid (1, P) \parallel (i, w) \rightsquigarrow (0, \vec{0})\} \end{aligned}$$

and we show that $\vec{0} \in \llbracket \sigma \rrbracket$ for any $\sigma \in \Sigma$. Hence the identity $\llbracket ! \Sigma \rrbracket = \{\vec{0}\}$ holds. Writing $v = \alpha_1.\vec{x}_1 + \dots + \alpha_n.\vec{x}_n$ we deduce $\llbracket \Gamma_v \rrbracket = \llbracket \alpha_1.x_1, \dots, \alpha_n.x_n \rrbracket = \alpha_1.\{\vec{x}_1\} + \dots + \alpha_n.\{\vec{x}_n\} = \{v\}$. As a consequence, we get $\llbracket ! \Sigma, \Gamma_v \rrbracket = \llbracket ! \Sigma \rrbracket + \llbracket \Gamma_v \rrbracket = \{v\}$. Thus, if $! \Sigma, \Gamma_v \vdash q_i$ has a proof in $\mathcal{G}_{\text{eILL}}$, by soundness of trivial phase semantics we deduce $v \in \llbracket ! \Sigma, \Gamma_v \rrbracket \subseteq \llbracket q_i \rrbracket$, i.e. the computation $(1, P) \parallel (i, v) \rightsquigarrow (0, \vec{0})$ terminates. $\square$

This allows us to conclude:

Theorem 8.4. $\text{MM} \leq \text{eILL}$.

9 Conclusion

Our formalisation consists of about 5200 lines of code. Around 200 loc are needed for the definition and reductions concerning PCP. The generic compiler needs 1500 loc. We have 1600 loc for the BSM section, 900 loc for the MM section, and 1000 loc for the ILL section.

Related work. There is not much related work concerning formalised undecidability proofs. Forster, Heiter, and Smolka [11] reduce the halting problem for Turing machines via string rewriting systems to PCP and further on to problems concerning context-free grammars.

Forster, Kirst and Smolka [12] reduce PCP to validity, satisfiability and provability in both classical and intuitionistic first-order logic (FOL). They have undecidability proofs not

requiring intermediate steps, in striking contrast to the involved formalisations needed to complete our reduction.⁵ They also include the notion of enumerability, connect (non)-enumerability with (un)decidability and give enumerability proofs for PCP and provable formulas.

With regard to the formalisation of models of computations in constructive type theory, there is slightly more related work, starting with a proof of the s_{mn} theorem by Zammit [38]. Forster and Smolka [15] formalise the call-by-value λ -calculus as model of computation in Coq. They constructively develop the standard theory of computation, but only consider undecidability problems related to the λ -calculus.

Larchey-Wendling [24] formalises another standard model of computation, that of μ -recursive functions, which are shown to be representable by guarded Coq-terms. In particular, total μ -recursive predicates $\mathbb{N}^k \rightarrow \mathbb{N} \rightarrow \mathbb{P}$ can be constructively reified into functions of type $\mathbb{N}^k \rightarrow \mathbb{N}$.

Asperti and Ricciotti [1] formalise multi-tape Turing machines in the proof assistant Matita and construct a universal Turing machine. Ciaffaglione [7] formalises Turing machines in Coq using coinductive definitions and proves the halting problem to be undecidable. Vinogradova [37] gives an abstract discussion of undecidability using Turing categories.

Forster and Kunze [13] present a framework to extract Coq functions to the call-by-value λ -calculus from [15] in a verified way. We could use this framework to prove the computability of all reductions explicitly, allowing for the assumption of non-constructive axioms, but thereby committing to a fixed model of computation.

Future work. We envision a library of undecidable problems in Coq. The library should contain basic problems that can be used for further reductions, like PCP, BSM, MM, and Halt as well as advanced problems with non-trivial proofs like ILL. We want to add the μ -recursive functions from [24], the halting problem for the call-by-value λ -calculus from [15], and more computational models, like for instance the WHILE language used by Jones [19]. Other basic problems that can serve as stepping stones for other reductions include tiling problems from e.g. Berger [3].

Formalising Hilbert's 10th problem — the undecidability of diophantine equations, — could be a challenge, e.g. following [18]. Finally, concerning linear logic, there are still fragments where the decidability status is unclear or debated: the entailment of multiplicative and exponential linear logic (MELL) was claimed to be decidable in [4] but this proof and hence the result is highly contested, see e.g. [10, footnote 1] and [35, footnote 4] and [25]. We hope that using the library, future research in this direction can be formalised.

⁵As it is highly expressive, FOL can directly express a characteristic formula for PCP. On the opposite side, directly giving a characterising ILL formula is quite infeasible, making intermediate steps crucial.

References

- [1] Andrea Asperti and Wilmer Ricciotti. 2015. A formalization of multi-tape Turing machines. *Theoret. Comput. Sci.* 603 (2015), 23–42.
- [2] Mauricio Ayala-Rincón and César A. Muñoz (Eds.). 2017. *Interactive Theorem Proving - 8th International Conference, ITP 2017, Proceedings*. Lecture Notes in Computer Science, Vol. 10499. Springer.
- [3] Robert Berger. 1966. *The undecidability of the domino problem*. Number 66 in Memoirs of the American Mathematical Society. American Mathematical Soc.
- [4] Katalin Bimbó. 2015. The decidability of the intensional fragment of classical linear logic. *Theoret. Comput. Sci.* 597 (2015), 1–17.
- [5] Thomas Braibant and Damien Pous. 2010. An efficient Coq tactic for deciding Kleene algebras. In *International Conference on Interactive Theorem Proving*. Springer, 163–178.
- [6] Adam Chlipala. 2010. A verified compiler for an impure functional language. In *Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, 93–106.
- [7] Alberto Ciaffaglione. 2016. Towards Turing computability via coinduction. *Science of Computer Programming* 126 (2016), 31–51.
- [8] Christian Doczkal and Joachim Bard. 2018. Completeness and decidability of converse PDL in the constructive type theory of Coq. In *Proceedings of the 7th ACM SIGPLAN International Conference on Certified Programs and Proofs*. ACM, 42–52.
- [9] Christian Doczkal and Gert Smolka. 2016. Completeness and decidability results for CTL in constructive type theory. *J. Automat. Reason.* 56, 3 (2016), 343–365.
- [10] Diego Figueira, Ranko Lazic, Jérôme Leroux, Filip Mazowiecki, and Grégoire Sutre. 2017. Polynomial-Space Completeness of Reachability for Succinct Branching VASS in Dimension One. In *ICALP 2017 (LIPIcs)*, Vol. 80. Schloss Dagstuhl, 119:1–14.
- [11] Yannick Forster, Edith Heiter, and Gert Smolka. 2018. Verification of PCP-Related Computational Reductions in Coq. In *International Conference on Interactive Theorem Proving*. Springer, 253–269.
- [12] Yannick Forster, Dominik Kirst, and Gert Smolka. 2018. On Synthetic Undecidability in Coq, with an Application to the Entscheidungsproblem. (2018). https://www.ps.uni-saarland.de/Publications/documents/ForsterEtAl_2018_On-Synthetic-Undecidability.pdf Submitted to CPP'19.
- [13] Yannick Forster and Fabian Kunze. 2016. Verified Extraction from Coq to a Lambda-Calculus. In *Coq Workshop 2016*.
- [14] Yannick Forster and Dominique Larchey-Wendling. 2018. Towards a library of formalised undecidable problems in Coq: The undecidability of intuitionistic linear logic. *Workshop on Syntax and Semantics of Low-level Languages, Oxford* (2018).
- [15] Yannick Forster and Gert Smolka. 2017. Weak call-by-value lambda calculus as a model of computation in Coq. See [2], 189–206.
- [16] Jean-Yves Girard. 1987. Linear logic. *Theoret. Comput. Sci.* 50, 1 (1987), 1–102.
- [17] Jean-Yves Girard, Yves Lafont, and Laurent Regnier (Eds.). 1995. *Advances in Linear Logic*. London Mathematical Society Lecture Note Series, Vol. 222. Cambridge University Press.
- [18] J. P. Jones and Y. V. Matijasevič. 1984. Register Machine Proof of the Theorem on Exponential Diophantine Representation of Enumerable Sets. *J. Symb. Log.* 49, 3 (1984), 818–829.
- [19] Neil D Jones. 1997. *Computability and complexity: from a programming perspective*. Vol. 21. MIT press.
- [20] Max Kanovich. 1994. Linear Logic as a Logic of Computations. *Ann. Pure Appl. Logic* 67, 1–3 (1994), 183–212.
- [21] Max Kanovich. 1995. The direct simulation of Minsky machines in linear logic. See [17], Chapter 2, 123–145.
- [22] Chantal Keller and Thorsten Altenkirch. 2010. Hereditary substitutions for simple types, formalized. In *Proceedings of the third ACM SIGPLAN workshop on Mathematically structured functional programming*. ACM, 3–10.
- [23] Ramana Kumar, Magnus O. Myreen, Michael Norrish, and Scott Owens. 2014. CakeML: a verified implementation of ML. In *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, 179–192.
- [24] Dominique Larchey-Wendling. 2017. Typing Total Recursive Functions in Coq. See [2], 371–388.
- [25] Dominique Larchey-Wendling. 2018. Constructive Decision via Redundancy-Free Proof-Search. In *Automated Reasoning - 9th International Joint Conference, IJCAR, Proceedings (Lecture Notes in Computer Science)*, Vol. 10900. Springer, 422–438.
- [26] Dominique Larchey-Wendling and Didier Galmiche. 2010. The undecidability of boolean BI through phase semantics. In *LICS 2010*. IEEE, 140–149.
- [27] Dominique Larchey-Wendling and Didier Galmiche. 2013. Non-deterministic Phase Semantics and the Undecidability of Boolean BI. *ACM Trans. Comput. Log.* 14, 1 (2013), 6:1–6:41.
- [28] Xavier Leroy. 2006. Formal certification of a compiler back-end or: programming a compiler with a proof assistant. In *Proceedings of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, 42–54.
- [29] Patrick Lincoln. 1995. Deciding provability of linear logic formulas. See [17], Chapter 2, 109–122.
- [30] Patrick Lincoln, John C. Mitchell, Andre Scedrov, and Natarajan Shankar. 1990. Decision Problems for Propositional Linear Logic. In *31st Annual Symposium on Foundations of Computer Science*, Vol. 2. IEEE Computer Society, 662–671.
- [31] Petar Maksimović and Alan Schmitt. 2015. HOCORE in coq. In *International Conference on Interactive Theorem Proving*. Springer, 278–293.
- [32] Marvin L Minsky. 1967. *Computation: finite and infinite machines*. Prentice-Hall, Inc.
- [33] Hartley Rogers. 1967. *Theory of recursive functions and effective computability*. Vol. 5. McGraw-Hill New York.
- [34] Steven Schäfer, Gert Smolka, and Tobias Tebbi. 2015. Completeness and decidability of de Bruijn substitution algebra in Coq. In *Proceedings of the 2015 Conference on Certified Programs and Proofs*. ACM, 67–73.
- [35] Sylvain Schmitz. 2016. The Complexity of Reachability in Vector Addition Systems. *ACM SIGLOG News* 3, 1 (2016), 4–21.
- [36] The Coq proof assistant. 2018. <http://coq.inria.fr>. (2018).
- [37] Polina Vinogradova. 2017. *Formalizing Abstract Computability: Turing Categories in Coq*. Ph.D. Dissertation. University of Ottawa.
- [38] Vincent Zammitt. 1997. *A Proof of the S-m-n theorem in Coq*. Technical report. The Computing Laboratory, The University of Kent, Canterbury, Kent, UK. <http://kar.kent.ac.uk/21524/>