

Mechanising Turing Machines and Computability Theory in Isabelle/HOL

Jian Xu¹, Xingyuan Zhang¹, and Christian Urban²

¹ PLA University of Science and Technology, China

² King's College London, UK

Abstract. We present a formalised theory of computability in the theorem prover Isabelle/HOL. This theorem prover is based on classical logic which precludes *direct* reasoning about computability: every boolean predicate is either true or false because of the law of excluded middle. The only way to reason about computability in a classical theorem prover is to formalise a concrete model for computation. We formalise Turing machines and relate them to abacus machines and recursive functions. We also formalise a universal Turing machine and Hoare-style reasoning techniques that allow us to reason about Turing machine programs. Our theory can be used to formalise other computability results. We give one example about the computational equivalence of single-sided Turing machines.

1 Introduction

Suppose you want to mechanise a proof for whether a predicate P , say, is decidable or not. Decidability of P usually amounts to showing whether $P \vee \neg P$ holds. But this does *not* work in Isabelle/HOL and other HOL theorem provers, since they are based on classical logic where the law of excluded middle ensures that $P \vee \neg P$ is always provable no matter whether P is constructed by computable means. We hit on this limitation previously when we mechanised the correctness proofs of two algorithms [9,10], but were unable to formalise arguments about decidability or undecidability.

The only satisfying way out of this problem in a theorem prover based on classical logic is to formalise a theory of computability. Norrish provided such a formalisation for HOL4. He chose the λ -calculus as the starting point for his formalisation because of its “simplicity” [6, Page 297]. Part of his formalisation is a clever infrastructure for reducing λ -terms. He also established the computational equivalence between the λ -calculus and recursive functions. Nevertheless he concluded that it would be appealing to have formalisations for more operational models of computations, such as Turing machines or register machines. One reason is that many proofs in the literature use them. He noted however that [6, Page 310]:

“If register machines are unappealing because of their general fiddliness, Turing machines are an even more daunting prospect.”

In this paper we take on this daunting prospect and provide a formalisation of Turing machines, as well as abacus machines (a kind of register machines) and recursive functions. To see the difficulties involved with this work, one has to understand that Turing

machine programs can be completely *unstructured*, behaving similar to Basic programs containing the infamous goto-statements [3]. This precludes in the general case a compositional Hoare-style reasoning about Turing programs. We provide such Hoare-rules for when it *is* possible to reason in a compositional manner (which is fortunately quite often), but also tackle the more complicated case when we translate abacus programs into Turing programs. This reasoning about concrete Turing machine programs is usually left out in the informal literature, e.g. [2].

We are not the first who formalised Turing machines: we are aware of the work by Asperti and Ricciotti [1]. They describe a complete formalisation of Turing machines in the Matita theorem prover, including a universal Turing machine. However, they do *not* formalise the undecidability of the halting problem since their main focus is complexity, rather than computability theory. They also report that the informal proofs from which they started are not “sufficiently accurate to be directly usable as a guideline for formalization” [1, Page 2]. For our formalisation we follow mainly the proofs from the textbook by Boolos et al [2] and found that the description there is quite detailed. Some details are left out however: for example, constructing the *copy Turing machine* is left as an exercise to the reader—a correctness proof is not mentioned at all; also [2] only shows how the universal Turing machine is constructed for Turing machines computing unary functions. We had to figure out a way to generalise this result to n -ary functions. Similarly, when compiling recursive functions to abacus machines, the textbook again only shows how it can be done for 2- and 3-ary functions, but in the formalisation we need arbitrary functions. But the general ideas for how to do this are clear enough in [2].

The main difference between our formalisation and the one by Asperti and Ricciotti is that their universal Turing machine uses a different alphabet than the machines it simulates. They write [1, Page 23]:

“In particular, the fact that the universal machine operates with a different alphabet with respect to the machines it simulates is annoying.”

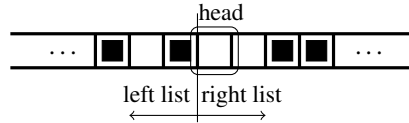
In this paper we follow the approach by Boolos et al [2], which goes back to Post [7], where all Turing machines operate on tapes that contain only *blank* or *occupied* cells. Traditionally the content of a cell can be any character from a finite alphabet. Although computationally equivalent, the more restrictive notion of Turing machines in [2] makes the reasoning more uniform. In addition some proofs *about* Turing machines are simpler. The reason is that one often needs to encode Turing machines—consequently if the Turing machines are simpler, then the coding functions are simpler too. Unfortunately, the restrictiveness also makes it harder to design programs for these Turing machines. In order to construct a universal Turing machine we therefore do not follow [1], instead follow the proof in [2] by translating abacus machines to Turing machines and in turn recursive functions to abacus machines. The universal Turing machine can then be constructed as a recursive function.

Contributions: We formalised in Isabelle/HOL Turing machines following the description of Boolos et al [2] where tapes only have blank or occupied cells. We mechanise the undecidability of the halting problem and prove the correctness of concrete Turing machines that are needed in this proof; such correctness proofs are left out in the informal literature. For reasoning about Turing machine programs we derive Hoare-rules.

We also construct the universal Turing machine from [2] by translating recursive functions to abacus machines and abacus machines to Turing machines. Since we have set up in Isabelle/HOL a very general computability model and undecidability result, we are able to formalise other results: we describe a proof of the computational equivalence of single-sided Turing machines, which is not given in [2], but needed for example for formalising the undecidability proof of Wang’s tiling problem [8].

2 Turing Machines

Turing machines can be thought of as having a *head*, “gliding” over a potentially infinite tape. Boolos et al [2] only consider tapes with cells being either blank or occupied, which we represent by a datatype having two constructors, namely *Bk* and *Oc*. One way to represent such tapes is to use a pair of lists, written (l, r) , where l stands for the tape on the left-hand side of the head and r for the tape on the right-hand side. We use the notation Bk^n (similarly Oc^n) for lists composed of n elements of *Bks*. We also have the convention that the head, abbreviated *hd*, of the right list is the cell on which the head of the Turing machine currently scans. This can be pictured as follows:



Note that by using lists each side of the tape is only finite. The potential infinity is achieved by adding an appropriate blank or occupied cell whenever the head goes over the “edge” of the tape. To make this formal we define five possible *actions* the Turing machine can perform:

$$\begin{aligned}
 a ::= & W_{Bk} \text{ (write blank, } Bk) & | & L \text{ (move left)} & | & Nop \text{ (do-nothing operation)} \\
 & W_{Oc} \text{ (write occupied, } Oc) & | & R \text{ (move right)}
 \end{aligned}$$

We slightly deviate from the presentation in [2] (and also [1]) by using the *Nop* operation; however its use will become important when we formalise halting computations and also universal Turing machines. Given a tape and an action, we can define the following tape updating function:

$$\begin{aligned}
 \text{update } (l, r) W_{Bk} & \stackrel{\text{def}}{=} (l, Bk::tl \ r) \\
 \text{update } (l, r) W_{Oc} & \stackrel{\text{def}}{=} (l, Oc::tl \ r) \\
 \text{update } (l, r) L & \stackrel{\text{def}}{=} \text{if } l = [] \text{ then } ([], Bk::r) \text{ else } (tl \ l, hd \ l::r) \\
 \text{update } (l, r) R & \stackrel{\text{def}}{=} \text{if } r = [] \text{ then } (Bk::l, []) \text{ else } (hd \ r::l, tl \ r) \\
 \text{update } (l, r) Nop & \stackrel{\text{def}}{=} (l, r)
 \end{aligned}$$

The first two clauses replace the head of the right list with a new *Bk* or *Oc*, respectively. To see that these two clauses make sense in case where r is the empty list, one has to

know that the tail function, tl , is defined such that $tl [] \stackrel{def}{=} []$ holds. The third clause implements the move of the head one step to the left: we need to test if the left-list l is empty; if yes, then we just prepend a blank cell to the right list; otherwise we have to remove the head from the left-list and prepend it to the right list. Similarly in the fourth clause for a right move action. The *Nop* operation leaves the the tape unchanged.

Next we need to define the *states* of a Turing machine. We follow the choice made in [1] by representing a state with a natural number and the states in a Turing machine program by the initial segment of natural numbers starting from 0. In doing so we can compose two Turing machine programs by shifting the states of one by an appropriate amount to a higher segment and adjusting some “next states” in the other.

An *instruction* of a Turing machine is a pair consisting of an action and a natural number (the next state). A *program* p of a Turing machine is then a list of such pairs. Using as an example the following Turing machine program, which consists of four instructions

$$dither \stackrel{def}{=} \underbrace{[(W_{Bk}, 1), (R, 2)]}_{\substack{\text{1st state} \\ = \text{starting state}}} \underbrace{[(L, 1), (L, 0)]}_{\text{2nd state}} \quad (1)$$

the reader can see we have organised our Turing machine programs so that segments of two belong to a state. The first component of such a segment determines what action should be taken and which next state should be transitioned to in case the head reads a *Bk*; similarly the second component determines what should be done in case of reading *Oc*. We have the convention that the first state is always the *starting state* of the Turing machine. The 0-state is special in that it will be used as the “halting state”. There are no instructions for the 0-state, but it will always perform a *Nop*-operation and remain in the 0-state. Unlike Asperti and Ricciotti [1], we have chosen a very concrete representation for programs, because when constructing a universal Turing machine, we need to define a coding function for programs. This can be directly done for our programs-as-lists, but is slightly more difficult for the functions used by Asperti and Ricciotti.

Given a program p , a state and the cell being read by the head, we need to fetch the corresponding instruction from the program. For this we define the function *fetch*

$$\begin{aligned} fetch\ p\ 0\ _ &= (Nop, 0) \\ fetch\ p\ (Suc\ s)\ Bk &\stackrel{def}{=} \text{case } nth_of\ p\ (2 * s)\ \text{of} \\ &\quad None \Rightarrow (Nop, 0) \mid \text{Some } i \Rightarrow i \\ fetch\ p\ (Suc\ s)\ Oc &\stackrel{def}{=} \text{case } nth_of\ p\ (2 * s + 1)\ \text{of} \\ &\quad None \Rightarrow (Nop, 0) \mid \text{Some } i \Rightarrow i \end{aligned} \quad (2)$$

In this definition the function *nth_of* returns the n th element from a list, provided it exists (*Some*-case), or if it does not, it returns the default action *Nop* and the default state 0 (*None*-case). We often have to restrict Turing machine programs to be well-formed: a program p is *well-formed* if it satisfies the following three properties:

$$wf\ p \stackrel{def}{=} 2 \leq length\ p \wedge is_even\ (length\ p) \wedge (\forall (a, s) \in p. s \leq length\ p\ div\ 2)$$

The first states that p must have at least an instruction for the starting state; the second that p has a Bk and Oc instruction for every state, and the third that every next-state is one of the states mentioned in the program or being the 0 -state.

We need to be able to sequentially compose Turing machine programs. Given our concrete representation, this is relatively straightforward, if slightly fiddly. We use the following two auxiliary functions:

$$\begin{aligned} \text{shift } p \ n &\stackrel{\text{def}}{=} \text{map } (\lambda(a, s). (a, \text{if } s = 0 \text{ then } 0 \text{ else } s + n)) \ p \\ \text{adjust } p &\stackrel{\text{def}}{=} \text{map } (\lambda(a, s). (a, \text{if } s = 0 \text{ then } \text{Suc } (\text{length } p \ \text{div } 2) \text{ else } s)) \ p \end{aligned}$$

The first adds n to all states, except the 0 -state, thus moving all “regular” states to the segment starting at n ; the second adds $\text{Suc } (\text{length } p \ \text{div } 2)$ to the 0 -state, thus redirecting all references to the “halting state” to the first state after the program p . With these two functions in place, we can define the *sequential composition* of two Turing machine programs p_1 and p_2 as

$$p_1 \oplus p_2 \stackrel{\text{def}}{=} \text{adjust } p_1 \ @ \ \text{shift } p_2 \ (\text{length } p_1 \ \text{div } 2)$$

A *configuration* c of a Turing machine is a state together with a tape. This is written as $(s, (l, r))$. We say a configuration is *final* if $s = 0$ and we say a predicate P holds for a configuration if P holds for the tape (l, r) . If we have a configuration and a program, we can calculate what the next configuration is by fetching the appropriate action and next state from the program, and by updating the state and tape accordingly. This single step of execution is defined as the function *step*

$$\text{step } (s, (l, r)) \ p \stackrel{\text{def}}{=} \text{let } (a, s') = \text{fetch } p \ s \ (\text{read } r) \\ \text{in } (s', \text{update } (l, r) \ a)$$

where *read* r returns the head of the list r , or if r is empty it returns Bk . It is impossible in Isabelle/HOL to lift the *step*-function in order to realise a general evaluation function for Turing machines programs. The reason is that functions in HOL-based provers need to be terminating, and clearly there are programs that are not. We can however define a recursive evaluation function that performs exactly n steps:

$$\begin{aligned} \text{steps } c \ p \ 0 &\stackrel{\text{def}}{=} c \\ \text{steps } c \ p \ (\text{Suc } n) &\stackrel{\text{def}}{=} \text{steps } (\text{step } c \ p) \ p \ n \end{aligned}$$

Recall our definition of *fetch* (shown in (2)) with the default value for the 0 -state. In case a Turing program takes according to the usual textbook definition, say [2], less than n steps before it halts, then in our setting the *steps*-evaluation does not actually halt, but rather transitions to the 0 -state (the final state) and remains there performing *Nop*-actions until n is reached.

We often need to restrict tapes to be in standard form, which means the left list of the tape is either empty or only contains *Bks*, and the right list contains some “clusters” of *Ocs* separated by single blanks. To make this formal we define the following overloaded function encoding natural numbers into lists of *Ocs* and *Bks*.

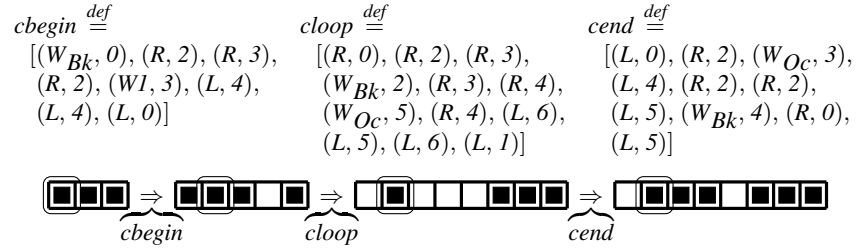


Fig. 1. The three components of the *copy Turing machine* (above). If started (below) with the tape $([], \langle 2 \rangle)$ the first machine appends $[Bk, Oc]$ at the end of the right tape; the second then “moves” all Ocs except the first from the beginning of the tape to the end; the third “refills” the original block of Ocs . The resulting tape is $([Bk], \langle (2, 2) \rangle)$.

$$\begin{aligned}
\langle n \rangle &\stackrel{def}{=} Oc^n + 1 & \langle [] \rangle &\stackrel{def}{=} [] \\
\langle (n, m) \rangle &\stackrel{def}{=} \langle n \rangle @ [Bk] @ \langle m \rangle & \langle [n] \rangle &\stackrel{def}{=} \langle n \rangle \\
&& \langle n::ns \rangle &\stackrel{def}{=} \langle (n, ns) \rangle
\end{aligned}$$

A *standard tape* is then of the form $(Bk^l, \langle [n_1, \dots, n_m] \rangle)$ for some l and n_i . Note that the head in a standard tape “points” to the leftmost Oc on the tape. Note also that the natural number 0 is represented by a single filled cell on a standard tape, 1 by two filled cells and so on.

Before we can prove the undecidability of the halting problem for our Turing machines working on standard tapes, we have to analyse two concrete Turing machine programs and establish that they are correct—that means they are “doing what they are supposed to be doing”. Such correctness proofs are usually left out in the informal literature, for example [2]. The first program we need to prove correct is the *dither* program shown in (1) and the second program is *copy* defined as

$$copy \stackrel{def}{=} cbegin \oplus cloop \oplus cend \quad (3)$$

whose three components are given in Figure 1. For our correctness proofs, we introduce the notion of total correctness defined in terms of *Hoare-triples*, written $\{P\} p \{Q\}$. They implement the idea that a program p started in state l with a tape satisfying P will after some n steps halt (have transitioned into the halting state) with a tape satisfying Q . This idea is very similar to the notion of *realisability* in [1]. We also have *Hoare-pairs* of the form $\{P\} p \uparrow$ implementing the case that a program p started with a tape satisfying P will loop (never transition into the halting state). Both notion are formally defined as

$ \begin{aligned} \{P\} p \{Q\} &\stackrel{def}{=} \\ &\forall (l, r). \\ &\text{if } P(l, r) \text{ holds then} \\ &\exists n. \text{ such that} \\ &is_final(steps(I, (l, r)) p n) \wedge \\ &Q \text{ holds for } (steps(I, (l, r)) p n) \end{aligned} $	$ \begin{aligned} \{P\} p \uparrow &\stackrel{def}{=} \\ &\forall (l, r). \\ &\text{if } P(l, r) \text{ holds then} \\ &\forall n. \neg is_final(steps(I, (l, r)) p n) \end{aligned} $
--	---

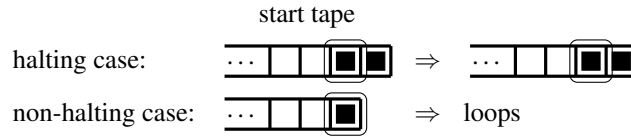
For our Hoare-triples we can easily prove the following Hoare-consequence rule

$$\frac{P' \mapsto P \quad \{P\} p \{Q\} \quad Q \mapsto Q'}{\{P'\} p \{Q'\}} \quad (4)$$

where $P' \mapsto P$ stands for the fact that for all tapes tp , $P' tp$ implies $P tp$ (similarly for Q and Q').

Like Asperti and Ricciotti with their notion of realisability, we have set up our Hoare-rules so that we can deal explicitly with total correctness and non-termination, rather than have notions for partial correctness and termination. Although the latter would allow us to reason more uniformly (only using Hoare-triples), we prefer our definitions because we can derive below some simple Hoare-rules for sequentially composed Turing programs. In this way we can reason about the correctness of *cbegin*, for example, completely separately from *cloop* and *ccend*.

It is relatively straightforward to prove that the Turing program *dither* shown in (1) is correct. This program should be the “identity” when started with a standard tape representing I but loops when started with 0 instead, as pictured below.



We can prove the following Hoare-statements:

$$\begin{aligned} & \{\lambda tp. \exists k. tp = (Bk^k, \langle I \rangle)\} \textit{dither} \{\lambda tp. \exists k. tp = (Bk^k, \langle I \rangle)\} \\ & \{\lambda tp. \exists k. tp = (Bk^k, \langle 0 \rangle)\} \textit{dither} \uparrow \end{aligned}$$

The first is by a simple calculation. The second is by an induction on the number of steps we can perform starting from the input tape.

The program *copy* defined in (3) has 15 states; its purpose is to produce the standard tape $(Bks, \langle (n, n) \rangle)$ when started with $(Bks, \langle n \rangle)$, that is making a copy of a value on the tape. Reasoning about this program is substantially harder than about *dither*. To ease the burden, we derive the following two Hoare-rules for sequentially composed programs.

$$\frac{\{P\} p_1 \{Q\} \quad \{Q\} p_2 \{R\}}{\{P\} p_1 \oplus p_2 \{R\}} \textit{wfp}_1 \quad \frac{\{P\} p_1 \{Q\} \quad \{Q\} p_2 \uparrow}{\{P\} p_1 \oplus p_2 \uparrow} \textit{wfp}_1$$

The first corresponds to the usual Hoare-rule for composition of two terminating programs. The second rule gives the conditions for when the first program terminates generating a tape for which the second program loops. The side-conditions about \textit{wfp}_1 are needed in order to ensure that the redirection of the halting and initial state in p_1 and p_2 , respectively, match up correctly. These Hoare-rules allow us to prove the correctness of *copy* by considering the correctness of the components *cbegin*, *cloop* and *ccend* in isolation. This simplifies the reasoning considerably, for example when designing decreasing measures for proving the termination of the programs. We will show the

$I_1 n(l, r)$	$\stackrel{def}{=} (l, r) = (\[], Oc^n)$	(starting state)
$I_2 n(l, r)$	$\stackrel{def}{=} \exists i j. 0 < i \wedge i + j = n \wedge (l, r) = (Oc^i, Oc^j)$	
$I_3 n(l, r)$	$\stackrel{def}{=} 0 < n \wedge (l, tl\ r) = (Bk::Oc^n, \[])$	
$I_4 n(l, r)$	$\stackrel{def}{=} 0 < n \wedge (l, r) = (Oc^n, [Bk, Oc]) \vee (l, r) = (Oc^{n-1}, [Oc, Bk, Oc])$	
$I_0 n(l, r)$	$\stackrel{def}{=} 1 < n \wedge (l, r) = (Oc^{n-2}, [Oc, Oc, Bk, Oc]) \vee$ $n = 1 \wedge (l, r) = (\[], [Bk, Oc, Bk, Oc])$	(halting state)
$J_1 n(l, r)$	$\stackrel{def}{=} \exists i j. i + j + 1 = n \wedge (l, r) = (Oc^i, Oc::Oc::Bk^j @ Oc^j) \wedge 0 < j \vee$ $0 < n \wedge (l, r) = (\[], Bk::Oc::Bk^n @ Oc^n)$	(starting state)
$J_0 n(l, r)$	$\stackrel{def}{=} 0 < n \wedge (l, r) = ([Bk], Oc::Bk^n @ Oc^n)$	(halting state)
$K_1 n(l, r)$	$\stackrel{def}{=} 0 < n \wedge (l, r) = ([Bk], Oc::Bk^n @ Oc^n)$	(starting state)
$K_0 n(l, r)$	$\stackrel{def}{=} 0 < n \wedge (l, r) = ([Bk], Oc^n @ (Bk::Oc^n))$	(halting state)

Fig. 2. The invariants I_0, \dots, I_4 are for the states of *cbegin*. Below, the invariants only for the starting and halting states of *cloop* and *cend* are shown. In each invariant the parameter n stands for the number of *Ocs* with which the Turing machine is started.

details for the program *cbegin*. For the two other programs we refer the reader to our formalisation.

Given the invariants I_0, \dots, I_4 shown in Figure 2, which correspond to each state of *cbegin*, we define the following invariant for the whole *cbegin* program:

$$\begin{aligned}
 I_{cbegin} n(s, tp) &\stackrel{def}{=} \text{if } s = 0 \text{ then } I_0 n \text{ tp} \\
 &\quad \text{else if } s = 1 \text{ then } I_1 n \text{ tp} \\
 &\quad \text{else if } s = 2 \text{ then } I_2 n \text{ tp} \\
 &\quad \text{else if } s = 3 \text{ then } I_3 n \text{ tp} \\
 &\quad \text{else if } s = 4 \text{ then } I_4 n \text{ tp} \\
 &\quad \text{else False}
 \end{aligned}$$

This invariant depends on n representing the number of *Ocs* (or encoded number) on the tape. It is not hard (26 lines of automated proof script) to show that for $0 < n$ this invariant is preserved under the computation rules *step* and *steps*. This gives us partial correctness for *cbegin*.

We next need to show that *cbegin* terminates. For this we introduce lexicographically ordered pairs (n, m) derived from configurations $(s, (l, r))$: n is the state s , but ordered according to how *cbegin* executes them, that is $1 > 2 > 3 > 4 > 0$; in order to have a strictly decreasing measure, m takes the data on the tape into account and is calculated according to the following measure function:

$$\begin{aligned}
 M_{cbegin}(s, (l, r)) &\stackrel{def}{=} \text{if } s = 2 \text{ then length } r \\
 &\quad \text{else if } s = 3 \text{ then (if } r = \[] \vee r = [Bk] \text{ then } 1 \text{ else } 0) \\
 &\quad \text{else if } s = 4 \text{ then length } l \\
 &\quad \text{else } 0
 \end{aligned}$$

With this in place, we can show that for every starting tape of the form $([], Oc^n)$ with $0 < n$, the Turing machine *cbegin* will eventually halt (the measure decreases in each step). Taking this and the partial correctness proof together, we obtain the left-most Hoare-triple for *cbegin*:

$$\{I_1 n\} cbegin \{I_0 n\} \quad \{J_1 n\} cloop \{J_0 n\} \quad \{K_1 n\} cend \{K_0 n\}$$

where we assume $0 < n$ (similar reasoning is needed for the Hoare-triples for *clloop* and *cend*). Since the invariant of the halting state of *cbegin* implies the invariant of the starting state of *clloop*, that is $I_0 n \mapsto J_1 n$ holds, and also $J_0 n = K_1 n$, we can derive the following Hoare-triple for the correctness of *copy*:

$$\{\lambda tp. tp = ([], \langle n \rangle)\} copy \{\lambda tp. tp = ([Bk], \langle (n, n) \rangle)\}$$

That means if we start with a tape of the form $([], \langle n \rangle)$ then *copy* will halt with the tape $([Bk], \langle (n, n) \rangle)$, as desired.

Finally, we are in the position to prove the undecidability of the halting problem. A program p started with a standard tape containing the (encoded) numbers ns will *halt* with a standard tape containing a single (encoded) number is defined as

$$halts p ns \stackrel{def}{=} \{\lambda tp. tp = ([], \langle ns \rangle)\} p \{\lambda tp. \exists k n. tp = (Bk^k, \langle n \rangle)\}$$

This roughly means we considering only Turing machine programs representing functions that take some numbers as input and produce a single number as output. For undecidability, the property we are proving is that there is no Turing machine that can decide in general whether a Turing machine program halts (answer either 0 for halting and 1 for looping). Given our correctness proofs for *dither* and *copy* shown above, this non-existence is now relatively straightforward to establish. We first assume there is a coding function, written *code M*, which represents a Turing machine M as a natural number. No further assumptions are made about this coding function. Suppose a Turing machine H exists such that if started with the standard tape $([Bk], \langle (code M, ns) \rangle)$ returns 0, respectively 1, depending on whether M halts when started with the input tape containing ns . This assumption is formalised as follows—for all M and all lists of natural numbers ns :

$$\begin{aligned} &halts M ns \text{ implies } \{\lambda tp. tp = ([Bk], \langle (code M, ns) \rangle)\} H \{\lambda tp. \exists k. tp = (Bk^k, \langle 0 \rangle)\} \\ &\neg halts M ns \text{ implies } \{\lambda tp. tp = ([Bk], \langle (code M, ns) \rangle)\} H \{\lambda tp. \exists k. tp = (Bk^k, \langle 1 \rangle)\} \end{aligned}$$

The contradiction can be derived using the following Turing machine

$$contra \stackrel{def}{=} copy \oplus H \oplus dither$$

Suppose *halts contra* [*code contra*] holds. Given the invariants on the left, we can derive the following Hoare-pair for *contra* on the right.

$$\begin{array}{l} P_1 \stackrel{def}{=} \lambda tp. tp = ([], \langle code contra \rangle) \\ P_2 \stackrel{def}{=} \lambda tp. tp = ([Bk], \langle (code contra, code contra) \rangle) \\ P_3 \stackrel{def}{=} \lambda tp. \exists k. tp = (Bk^k, \langle 0 \rangle) \end{array} \quad \frac{\frac{\{P_1\} copy \{P_2\} \quad \{P_2\} H \{P_3\}}{\{P_1\} copy \oplus H \{P_3\}} \quad \{P_3\} dither \uparrow}{\{P_1\} contra \uparrow}$$

This Hoare-pair contradicts our assumption that *contra* started with $\langle \text{code contra} \rangle$ halts.

Suppose $\neg \text{halts contra } [\text{code contra}]$ holds. Again given the invariants on the left, we can derive the Hoare-triple for *contra* on the right.

$$\begin{array}{l}
Q_1 \stackrel{\text{def}}{=} \lambda tp. tp = ([], \langle \text{code contra} \rangle) \\
Q_2 \stackrel{\text{def}}{=} \lambda tp. tp = ([Bk], \langle (\text{code contra}, \text{code contra}) \rangle) \\
Q_3 \stackrel{\text{def}}{=} \lambda tp. \exists k. tp = (Bk^k, \langle I \rangle)
\end{array}
\quad
\frac{\frac{\{Q_1\} \text{copy } \{Q_2\} \quad \{Q_2\} H \{Q_3\}}{\{Q_1\} \text{copy} \oplus H \{Q_3\}} \quad \{Q_3\} \text{dither } \{Q_3\}}{\{Q_1\} \text{contra } \{Q_3\}}$$

This time the Hoare-triple states that *contra* terminates with the “output” $\langle I \rangle$. In both case we come to an contradiction, which means we have to abandon our assumption that there exists a Turing machine *H* which can in general decide whether Turing machines terminate.

3 Abacus Machines

Boolos et al [2] use abacus machines as a stepping stone for making it less laborious to write Turing machine programs. Abacus machines operate over a set of registers R_0, R_1, \dots each being able to hold an arbitrary large natural number. We use natural numbers to refer to registers; we also use a natural number to represent a program counter and to represent jumping “addresses”. An abacus program is a list of *instructions* defined by the datatype:

$$\begin{array}{ll}
i ::= \text{Inc } R & \text{increment register } R \text{ by one} \\
| \text{Dec } R \ i & \text{if content of } R \text{ is non-zero, then decrement it by one} \\
& \text{otherwise jump to instruction } i \\
| \text{Goto } i & \text{jump to instruction } i
\end{array}$$

For example the program clearing the register R (that is setting it to 0) can be defined as follows:

$$\text{clear } R \ i \stackrel{\text{def}}{=} [\text{Dec } R \ i, \text{Goto } 0]$$

Running such a program means we start with the first instruction then execute one instructions after the other, unless there is a jump. For example the second instruction *Goto 0* in *clear* means we jump back to the first instruction thereby closing the loop. Like with our Turing machines, we fetch instructions from an abacus program such that a jump out of “range” behaves like a *Nop*-action. In this way it is again easy to define a function *steps* that executes n instructions of an abacus program. A *configuration* of an abacus machine is the current program counter together with a snapshot of all registers. By convention the value calculated by an abacus program is stored in the last register (the register with the highest index).

The main point of abacus programs is to be able to translate them to Turing machine programs. Registers and their content are represented by standard tapes. Because of the

jumps in abacus programs, it seems difficult to build a Turing machine programs out of components using our \oplus -operation shown in the previous section. To overcome this difficulty, we calculate a *layout* of an abacus program as follows

$$\begin{aligned} \text{layout } [] & \stackrel{\text{def}}{=} [] \\ \text{layout } (\text{Inc } R::os) & \stackrel{\text{def}}{=} 2 * R + 9::\text{layout } os \\ \text{layout } (\text{Dec } R i::os) & \stackrel{\text{def}}{=} 2 * R + 16::\text{layout } os \\ \text{layout } (\text{Goto } i::os) & \stackrel{\text{def}}{=} 1::\text{layout } os \end{aligned}$$

This gives us a list of natural numbers specifying how many states are needed to translate each abacus instruction. The *Goto* instruction is easiest to translate requiring only one state in the corresponding Turing machine:

$$[(\text{Nop}, i), (\text{Nop}, i)]$$

where i is the state in the Turing machine program to jump to. For translating the instruction *Inc*, one has to remember that the content of the registers are encoded in the Turing machine as standard tape. Therefore the translated Turing machine needs to first find the number corresponding to the register R . This needs a machine with $2 * R$ states and can be constructed as follows:

$$\begin{aligned} \text{find_nth } 0 & \stackrel{\text{def}}{=} [] \\ \text{find_nth } (\text{Suc } n) & \stackrel{\text{def}}{=} \\ & \text{find_nth } n @ [(W_{Oc}, 2 * n + 1), (R, 2 * n + 2), (R, 2 * n + 3), (R, 2 * n + 2)] \end{aligned}$$

Then we need to increase the “number” on the tape by one, and adjust the following “registers”. By adjusting we only need to replace the first Oc of each number by Bk and the last one from Bk to Oc . Finally, we need to transition the head of the Turing machine back into the standard position. This requires a Turing machine with 9 states (we omit the details). Similarly for the translation of *Dec*, where the translated Turing machine needs to first check whether the content of the corresponding register is 0. For this we have a Turing machine program with 16 states (again details are omitted).

Finally, having a Turing machine for each abacus instruction we need to “stitch” the Turing machines together into one so that each Turing machine component transitions to next one, just like in the abacus programs. One last problem to overcome is that an abacus program is assumed to calculate a value stored in the last register. That means we have to append a Turing machine that “mops up” the tape (cleaning all Ocs) except for the Ocs of the last number represented on the tape.

While generating the Turing machine program for an abacus program is not too difficult to formalise, the problem is that it contains *Gotos* all over the place. The unfortunate result is that we cannot use our Hoare-rules for reasoning about sequentially composed programs (for this the programs need to be independent). Instead we have to treat the translated Turing machine as one “big block” and prove as invariant that it performs the same operations as the abacus program. For this we have to show that for each configuration of an abacus machine the *step*-function is simulated by zero or more steps in our translated Turing machine. This leads to a rather large “monolithic”

correctness proof (4600 loc and 380 sublemmas) that on the conceptual level is difficult to break down into smaller components.

4 Recursive Functions and a Universal Turing Machine

5 Conclusion

We have formalised the main results from three chapters in the textbook by Boolos et al [2]. Following in the footsteps of another paper [5] formalising the results from a textbook, we could have titled our paper “Boolos et al are (almost) Right”. We have not attempted to formalise everything precisely as Boolos et al present it, but find definitions that make mechanised proofs manageable. We have found a small inconsistency in the usage of definitions of ... Our interest in this subject arose from correctness proofs about algorithms where we were unable to formalise arguments about decidability.

The most closely related work is by Norrish [6], and Asperti and Ricciotti [1]. Norrish bases his approach on λ -terms. For this he introduced a clever rewriting technology based on combinators and de-Bruijn indices for rewriting modulo β -equivalence (to keep it manageable)

Later on we need to consider specific Turing machines that start with a tape in standard form and halt the computation in standard form. To define a tape in standard form, it is useful to have an operation that translates lists of natural numbers into tapes.

This means the Turing machine starts with a tape containing n *Ocs* and the head pointing to the first one; the Turing machine halts with a tape consisting of some *Bks*, followed by a “cluster” of *Ocs* and after that by some *Bks*. The head in the output is pointing again at the first *Oc*. The intuitive meaning of this definition is to start the Turing machine with a tape corresponding to a value n and producing a new tape corresponding to the value l (the number of *Ocs* clustered on the output tape).

Magnus: invariants – Section 5.4.5 on page 75.

There is a tantalising connection with recent work [4] about verifying X86 assembly code. They observed It remains to be seen whether their specification logic for assembly code can make it easier to reason about our Turing programs.

References

1. A. Asperti and W. Ricciotti. Formalizing Turing Machines. In *Proc. of the 19th International Workshop on Logic, Language, Information and Computation (WoLLIC)*, volume 7456 of *LNCS*, pages 1–25, 2012.
2. G. Boolos, J. P. Burgess, and R. C. Jeffrey. *Computability and Logic (5th ed.)*. Cambridge University Press, 2007.
3. E. W. Dijkstra. Go to Statement Considered Harmful. *Communications of the ACM*, 11(3):147–148, 1968.
4. J. B. Jensen, N. Benton, and A. Kennedy. High-Level Separation Logic for Low-Level Code. In *Proc. of the 40th Symposium on Principles of Programming Languages (POPL)*, pages 301–314, 2013.

5. T. Nipkow. Winskel is (almost) Right: Towards a Mechanized Semantics Textbook. *Formal Aspects of Computing*, 10:171–186, 1998.
6. M. Norrish. Mechanised Computability Theory. In *Proc. of the 2nd Conference on Interactive Theorem Proving (ITP)*, volume 6898 of *LNCS*, pages 297–311, 2011.
7. E. Post. Finite Combinatory Processes-Formulation 1. *Journal of Symbolic Logic*, 1(3):103–105, 1936.
8. R. M. Robinson. Undecidability and Nonperiodicity for Tilings of the Plane. *Inventiones Mathematicae*, 12:177–209, 1971.
9. C. Urban, J. Cheney, and S. Berghofer. Mechanizing the Metatheory of LF. *ACM Transactions on Computational Logic*, 12:15:1–15:42, 2011.
10. C. Wu, X. Zhang, and C. Urban. A Formal Model and Correctness Proof for an Access Control Policy Framework. Submitted, 2013.