On the correctness of a branch displacement algorithm⋆

Jaap Boender and Claudio Sacerdoti Coen

Dipartimento di Scienze dell'Informazione,
Università degli Studi di Bologna

Abstract The branch displacement problem is a well-known problem in assembler design. It revolves around the feature, present in several processor families, of having different instructions, of different sizes, for jumps of different displacements. The problem, which is provably NP-hard, is then to select the instructions such that one ends up with the smallest possible program.

During our research with the CerCo project on formally verifying a C compiler, we have implemented and proven correct an algorithm for this problem. In this paper, we discuss the problem, possible solutions, our specific solutions and the proofs.

Keywords: formal verification, assembler, branch displacement optimisation

1 Introduction

The problem of branch displacement optimisation, also known as jump encoding, is a well-known problem in assembler design [4]. It is caused by the fact that in many architecture sets, the encoding (and therefore size) of some instructions depends on the distance to their operand (the instruction 'span'). The branch displacement optimisation problem consists of encoding these span-dependent instructions in such a way that the resulting program is as small as possible.

This problem is the subject of the present paper. After introducing the problem in more detail, we will discuss the solutions used by other compilers, present the algorithm we use in the CerCo assembler, and discuss its verification, that is the proofs of termination and correctness using the Matita proof assistant [1].

The research presented in this paper has been executed within the CerCo project which aims at formally verifying a C compiler with cost annotations. The target architecture for this project is the MCS-51, whose instruction set contains span-dependent instructions. Furthermore, its maximum addressable memory size is very small (64 Kb), which makes it important to generate programs that are as small as possible.

With this optimisation, however, comes increased complexity and hence increased possibility for error. We must make sure that the branch instructions are encoded correctly, otherwise the assembled program will behave unpredictably.

2 The branch displacement optimisation problem

In most modern instruction sets that have them, the only span-dependent instructions are branch instructions. Taking the ubiquitous x86-64 instruction set as an example, we find that it contains eleven different forms of the unconditional branch instruction, all with different ranges, instruction sizes and semantics (only six are valid in 64-bit mode, for example). Some examples are shown in Figure 1.

Instruction	Size (bytes)	Displacement range
Short jump	2	-128 to 127 bytes
Relative near jump	5	-2^{32} to $2^{32} - 1$ bytes
Absolute near jump	6	one segment (64-bit address)
Far jump	8	entire memory

Figure 1. List of x86 branch instructions

The chosen target architecture of the CerCo project is the Intel MCS-51, which features three types of branch instructions (or jump instructions; the two terms are used interchangeably), as shown in Figure 2.

Instruction	Size (bytes)	Execution time (cycles)	Displacement range
SJMP ('short jump')	2	2	-128 to 127 bytes
AJMP ('absolute jump')	2	2	one segment (11-bit address)
LJMP ('long jump')	3	3	entire memory

Figure 2. List of MCS-51 branch instructions

Conditional branch instructions are only available in short form, which means that a conditional branch outside the short address range has to be encoded using three branch instructions (for instructions whose logical negation is available, it can be done with two branch instructions, but for some instructions this is not available); the call instruction is only available in absolute and long forms.

Note that even though the MCS-51 architecture is much less advanced and simpler than the x86-64 architecture, the basic types of branch instruction remain the same: a short jump with a limited range, an intra-segment jump and a jump that can reach the entire available memory.

Generally, in code fed to the assembler as input, the only difference between branch instructions is semantics, not span. This means that a distinction is made between an unconditional branch and the several kinds of conditional branch, but not between their short, absolute or long variants.

The algorithm used by the assembler to encode these branch instructions into the different machine instructions is known as the *branch displacement algorithm*. The optimisation problem consists of finding as small an encoding as possible, thus minimising program length and execution time.

This problem is known to be NP-complete [8,11], which could make finding an optimal solution very time-consuming.

The canonical solution, as shown by Szymanski [11] or more recently by Dickson [2] for the x86 instruction set, is to use a fixed point algorithm that starts with the shortest possible encoding (all branch instruction encoded as short jumps, which is likely not a correct solution) and then iterates over the program to re-encode those branch instructions whose target is outside their range.

Adding absolute jumps

In both papers mentioned above, the encoding of a jump is only dependent on the distance between the jump and its target: below a certain value a short jump can be used; above this value the jump must be encoded as a long jump.

Here, termination of the smallest fixed point algorithm is easy to prove. All branch instructions start out encoded as short jumps, which means that the distance between any branch instruction and its target is as short as possible. If, in this situation, there is a branch instruction b whose span is not within the range for a short jump, we can be sure that we can never reach a situation where the span of j is so small that it can be encoded as a short jump. This argument continues to hold throughout the subsequent iterations of the algorithm: short jumps can change into long jumps, but not *vice versa*, as spans only increase. Hence, the algorithm either terminates early when a fixed point is reached or when all short jumps have been changed into long jumps.

Also, we can be certain that we have reached an optimal solution: a short jump is only changed into a long jump if it is absolutely necessary.

However, neither of these claims (termination nor optimality) hold when we add the absolute jump, as with absolute jumps, the encoding of a branch instruction no longer depends only on the distance between the branch instruction and its target: in order for an absolute jump to be possible, they need to be in the same segment (for the MCS-51, this means that the first 5 bytes of their addresses have to be equal). It is therefore entirely possible for two branch instructions with the same span to be encoded in different ways (absolute if the branch instruction and its target are in the same segment, long if this is not the case).

This invalidates our earlier termination argument: a branch instruction, once encoded as a long jump, can be re-encoded during a later iteration as an absolute jump. Consider the program shown in Figure 3. At the start of the first iteration,

```
        jmp X
        .
        .
        .
L₀:
        .
        .
        .
        jmp L₀
```

Figure 3. Example of a program where a long jump becomes absolute

both the branch to X and the branch to L_0 are encoded as small jumps. Let us assume that in this case, the placement of L_0 and the branch to it are such that L_0 is just outside the segment that contains this branch. Let us also assume that the distance between L_0 and the branch to it are too large for the branch instruction to be encoded as a short jump.

All this means that in the second iteration, the branch to L_0 will be encoded as a long jump. If we assume that the branch to X is encoded as a long jump as well, the size of the branch instruction will increase and L_0 will be 'propelled' into the same segment as its branch instruction, because every subsequent instruction will move one byte forward. Hence, in the third iteration, the branch to L_0 can be encoded as an absolute jump. At first glance, there is nothing that prevents us from constructing a configuration where two branch instructions interact in such a way as to iterate indefinitely between long and absolute encodings.

This situation mirrors the explanation by Szymanski [11] of why the branch displacement optimisation problem is NP-complete. In this explanation, a condition for NP-completeness is the fact that programs be allowed to contain *pathological* jumps. These are branch instructions that can normally not be encoded as a short(er) jump, but gain this property when some other branch instructions are encoded as a long(er) jump. This is exactly what happens in Figure 3. By encoding the first branch instruction as a long jump, another branch instruction switches from long to absolute (which is shorter).

In addition, our previous optimality argument no longer holds. Consider the program shown in Figure 4. Suppose that the distance between L_0 and L_1 is such that if jmp X is encoded as a short jump, there is a segment border just after L_1. Let us also assume that the three branches to L_1 are all in the same segment, but far enough away from L_1 that they cannot be encoded as short jumps.

Then, if jmp X were to be encoded as a short jump, which is clearly possible, all of the branches to L_1 would have to be encoded as long jumps. However, if jmp X were to be encoded as a long jump, and therefore increase in size, L_1 would be 'propelled' across the segment border, so that the three branches to L_1 could be encoded as absolute jumps. Depending on the relative sizes of long and absolute jumps, this solution might actually be smaller than the one reached by the smallest fixed point algorithm.

```
L₀: jmp X
X:
    ⋮
L₁:
    ⋮
    jmp L₁
    ⋮
    jmp L₁
    ⋮
    jmp L₁
    ⋮
```

Figure 4. Example of a program where the fixed-point algorithm is not optimal

3 Our algorithm

3.1 Design decisions

Given the NP-completeness of the problem, to arrive at an optimal solution (using, for example, a constraint solver) will potentially take a great amount of time.

The SDCC compiler [9], which has a backend targetting the MCS-51 instruction set, simply encodes every branch instruction as a long jump without taking the distance into account. While certainly correct (the long jump can reach any destination in memory) and a very fast solution to compute, it results in a less than optimal solution.

On the other hand, the gcc compiler suite [3], while compiling C on the x86 architecture, uses a greatest fix point algorithm. In other words, it starts with all branch instructions encoded as the largest jumps available, and then tries to reduce the size of branch instructions as much as possible.

Such an algorithm has the advantage that any intermediate result it returns is correct: the solution where every branch instruction is encoded as a large jump is always possible, and the algorithm only reduces those branch instructions whose destination address is in range for a shorter jump. The algorithm can thus be stopped after a determined number of steps without sacrificing correctness.

The result, however, is not necessarily optimal. Even if the algorithm is run until it terminates naturally, the fixed point reached is the *greatest* fixed point, not the least fixed point. Furthermore, gcc (at least for the x86 architecture) only uses short and long jumps. This makes the algorithm more efficient, as shown in the previous section, but also results in a less optimal solution.

In the CerCo assembler, we opted at first for a least fixed point algorithm, taking absolute jumps into account.

Here, we ran into a problem with proving termination, as explained in the previous section: if we only take short and long jumps into account, the jump encoding can only switch from short to long, but never in the other direction. When we add absolute jumps, however, it is theoretically possible for a branch instruction to switch from absolute to long and back, as previously explained.

Proving termination then becomes difficult, because there is nothing that precludes a branch instruction from oscillating back and forth between absolute and long jumps indefinitely.

In order to keep the algorithm in the same complexity class and more easily prove termination, we decided to explicitly enforce the 'branch instructions must always grow longer' requirement: if a branch instruction is encoded as a long jump in one iteration, it will also be encoded as a long jump in all the following iterations. This means that the encoding of any branch instruction can change at most two times: once from short to absolute (or long), and once from absolute to long.

There is one complicating factor. Suppose that a branch instruction is encoded in step n as an absolute jump, but in step $n + 1$ it is determined that (because of changes elsewhere) it can now be encoded as a short jump. Due to the requirement that the branch instructions must always grow longer, this means that the branch encoding will be encoded as an absolute jump in step $n + 1$ as well.

This is not necessarily correct. A branch instruction that can be encoded as a short jump cannot always also be encoded as an absolute jump, as a short jump can bridge segments, whereas an absolute jump cannot. Therefore, in this situation we have decided to encode the branch instruction as a long jump, which is always correct.

The resulting algorithm, while not optimal, is at least as good as the ones from SDCC and `gcc`, and potentially better. Its complexity remains the same (there are at most $2n$ iterations, where n is the number of branch instructions in the program).

3.2 The algorithm in detail

The branch displacement algorithm forms part of the translation from pseudo-code to assembler. More specifically, it is used by the function that translates pseudo-addresses (natural numbers indicating the position of the instruction in the program) to actual addresses in memory.

Our original intention was to have two different functions, one function `policy` : $\mathbb{N} \to \{$`short_jump`, `absolute_jump`, `long_jump`$\}$ to associate jumps to their intended encoding, and a function $\sigma : \mathbb{N} \to$ `Word` to associate pseudo-addresses to machine addresses. σ would use `policy` to determine the size of jump instructions.

This turned out to be suboptimal from the algorithmic point of view and impossible to prove correct.

From the algorithmic point of view, in order to create the `policy` function, we must necessarily have a translation from pseudo-addresses to machine ad-

dresses (i.e. a σ function): in order to judge the distance between a jump and its destination, we must know their memory locations. Conversely, in order to create the σ function, we need to have the `policy` function, otherwise we do not know the sizes of the jump instructions in the program.

Much the same problem appears when we try to prove the algorithm correct: the correctness of `policy` depends on the correctness of σ, and the correctness of σ depends on the correctness of `policy`.

We solved this problem by integrating the `policy` and σ algorithms. We now have a function $\sigma : \mathbb{N} \rightarrow$ `Word` $\times$ `bool` which associates a pseudo-address to a machine address. The boolean denotes a forced long jump; as noted in the previous section, if during the fixed point computation an absolute jump changes to be potentially re-encoded as a short jump, the result is actually a long jump. It might therefore be the case that jumps are encoded as long jumps without this actually being necessary, and this information needs to be passed to the code generating function.

The assembler function encodes the jumps by checking the distance between source and destination according to σ, so it could select an absolute jump in a situation where there should be a long jump. The boolean is there to prevent this from happening by indicating the locations where a long jump should be encoded, even if a shorter jump is possible. This has no effect on correctness, since a long jump is applicable in any situation.

> **function** F(*labels*,*old_sigma*,*instr*,*ppc*,*acc*)
> $\langle added, pc, sigma \rangle \leftarrow acc$
>   **if** *instr* is a backward jump to j **then**
> $length \leftarrow \text{jump\_size}(pc, sigma_1(labels(j)))$
>   **else if** *instr* is a forward jump to j **then**
> $length \leftarrow \text{jump\_size}(pc, old\_sigma_1(labels(j)) + added)$
>   **else**
> $length \leftarrow$ `short_jump`
>   **end if**
> $old\_length \leftarrow \text{old\_sigma}_1(ppc)$
> $new\_length \leftarrow \max(old\_length, length)$
> $old\_size \leftarrow \text{old\_sigma}_2(ppc)$
> $new\_size \leftarrow \text{instruction\_size}(instr, new\_length)$
> $new\_added \leftarrow added + (new\_size - old\_size)$
> $new\_sigma_1(ppc + 1) \leftarrow pc + new\_size$
> $new\_sigma_2(ppc) \leftarrow new\_length$
> **return** $\langle new\_added, pc + new\_size, new\_sigma \rangle$
> **end function**

Figure 5. The heart of the algorithm

The algorithm, shown in Figure 5, works by folding the function F over the entire program, thus gradually constructing *sigma*. This constitutes one step in

the fixed point calculation; successive steps repeat the fold until a fixed point is reached.

Parameters of the function F are:

- a function *labels* that associates a label to its pseudo-address;
- *old_sigma*, the σ function returned by the previous iteration of the fixed point calculation;
- *instr*, the instruction currently under consideration;
- *ppc*, the pseudo-address of *instr*;
- *acc*, the fold accumulator, which contains *pc* (the highest memory address reached so far), *added* (the number of bytes added to the program size with respect to the previous iteration), and of course *sigma*, the σ function under construction.

The first two are parameters that remain the same through one iteration, the final three are standard parameters for a fold function (including *ppc*, which is simply the number of instructions of the program already processed).

The σ functions used by F are not of the same type as the final σ function: they are of type $\sigma : \mathbb{N} \to \mathbb{N} \times \{\texttt{short\_jump}, \texttt{absolute\_jump}, \texttt{long\_jump}\}$; a function that associates a pseudo-address with a memory address and a jump length. We do this to be able to ease the comparison of jump lengths between iterations. In the algorithm, we use the notation $sigma_1(x)$ to denote the memory address corresponding to x, and $sigma_2(x)$ to denote the jump length corresponding to x.

Note that the σ function used for label lookup varies depending on whether the label is behind our current position or ahead of it. For backward branches, where the label is behind our current position, we can use *sigma* for lookup, since its memory address is already known. However, for forward branches, the memory address of the address of the label is not yet known, so we must use *old_sigma*.

We cannot use *old_sigma* without change: it might be the case that we have already increased the size of some branch instructions before, making the program longer and moving every instruction forward. We must compensate for this by adding the size increase of the program to the label's memory address according to *old_sigma*, so that branch instruction spans do not get compromised.

Note also that we add the pc to *sigma* at location $ppc + 1$, whereas we add the jump length at location *ppc*. We do this so that $sigma(ppc)$ will always return a pair with the start address of the instruction at *ppc* and the length of its branch instruction (if any); the end address of the program can be found at $sigma(n + 1)$, where n is the number of instructions in the program.

4 The proof

In this section, we present the correctness proof for the algorithm in more detail. The main correctness statement is as follows (slightly simplified, here):

```
definition sigma_policy_specification :=
 λprogram: pseudo_assembly_program.
 λsigma: Word → Word.
 λpolicy: Word → bool.
  sigma (zero ...) = zero ... ∧
 ∀ppc: Word.∀ppc_ok.
 let ⟨preamble, instr_list⟩ := program in
 let pc := sigma ppc in
 let instruction :=
  \fst (fetch_pseudo_instruction instr_list ppc ppc_ok) in
 let next_pc := \fst (sigma (add ? ppc (bitvector_of_nat ? 1))) in
  (nat_of_bitvector ... ppc ≤ |instr_list| →
   next_pc = add ? pc (bitvector_of_nat ...
   (instruction_size ... sigma policy ppc instruction)))
  ∧
  ((nat_of_bitvector ... ppc < |instr_list| →
   nat_of_bitvector ... pc < nat_of_bitvector ... next_pc)
  ∨ (nat_of_bitvector ... ppc = |instr_list| → next_pc = (zero ...))).
```

Informally, this means that when fetching a pseudo-instruction at ppc, the translation by σ of $ppc + 1$ is the same as $\sigma(ppc)$ plus the size of the instruction at ppc. That is, an instruction is placed consecutively after the previous one, and there are no overlaps.

Instructions are also stocked in order: the memory address of the instruction at ppc should be smaller than the memory address of the instruction at $ppc + 1$. There is one exeception to this rule: the instruction at the very end of the program, whose successor address can be zero (this is the case where the program size is exactly equal to the amount of memory).

Finally, we enforce that the program starts at address 0, i.e. $\sigma(0) = 0$.

Since our computation is a least fixed point computation, we must prove termination in order to prove correctness: if the algorithm is halted after a number of steps without reaching a fixed point, the solution is not guaranteed to be correct. More specifically, branch instructions might be encoded which do not coincide with the span between their location and their destination.

Proof of termination rests on the fact that the encoding of branch instructions can only grow larger, which means that we must reach a fixed point after at most $2n$ iterations, with n the number of branch instructions in the program. This worst case is reached if at every iteration, we change the encoding of exactly one branch instruction; since the encoding of any branch instructions can change first from short to absolute and then from absolute to long, there can be at most $2n$ changes.

The proof has been carried out using the "Russell" style from [10]. We have proven some invariants of the F function from the previous section; these invariants are then used to prove properties that hold for every iteration of the fixed point computation; and finally, we can prove some properties of the fixed point.

4.1 Fold invariants

These are the invariants that hold during the fold of F over the program, and that will later on be used to prove the properties of the iteration.

Note that during the fixed point computation, the σ function is implemented as a trie for ease of access; computing $\sigma(x)$ is achieved by looking up the value of x in the trie. Actually, during the fold, the value we pass along is a pair $\mathbb{N} \times \text{ppc}_p\text{c}_m\text{ap}$. The first component is the number of bytes added to the program so far with respect to the previous iteration, and the second component, ppc_pc_map, is a pair consisting of the current size of the program and our σ function.

```
definition out_of_program_none :=
 λprefix:list labelled_instruction.λsigma:ppc_pc_map.
 ∀i.i < 2^16 → (i > |prefix| ↔
  bvt_lookup_opt ... (bitvector_of_nat ? i) (\snd sigma) = None ?).
```

This invariant states that any pseudo-address not yet examined is not present in the lookup trie.

```
definition not_jump_default :=
 λprefix:list labelled_instruction.λsigma:ppc_pc_map.
 ∀i.i < |prefix| →
  ¬is_jump (\snd (nth i ? prefix ⟨None ?, Comment []⟩)) →
  \snd (bvt_lookup ... (bitvector_of_nat ? i) (\snd sigma)
   ⟨0,short_jump⟩) = short_jump.
```

This invariant states that when we try to look up the jump length of a pseudo-address where there is no branch instruction, we will get the default value, a short jump.

```
definition jump_increase :=
 λprefix:list labelled_instruction.λop:ppc_pc_map.λp:ppc_pc_map.
 ∀i.i ≤ |prefix| →
 let ⟨opc,oj⟩ :=
  bvt_lookup ... (bitvector_of_nat ? i) (\snd op) ⟨0,short_jump⟩ in
 let ⟨pc,j⟩ :=
  bvt_lookup ... (bitvector_of_nat ? i) (\snd p) ⟨0,short_jump⟩ in
  jmpleq oj j.
```

This invariant states that between iterations (with *op* being the previous iteration, and *p* the current one), jump lengths either remain equal or increase. It is needed for proving termination.

```
definition sigma_compact_unsafe :=
 λprogram:list labelled_instruction.λlabels:label_map.λsigma:ppc_pc_map.
 ∀n.n < |program| →
  match bvt_lookup_opt ... (bitvector_of_nat ? n) (\snd sigma) with
  [ None ⇒ False
  | Some x ⇒ let ⟨pc,j⟩ := x in
    match bvt_lookup_opt ... (bitvector_of_nat ? (S n)) (\snd sigma) with
    [ None ⇒ False
    | Some x1 ⇒ let ⟨pc1,j1⟩ := x1 in
      pc1 = pc + instruction_size_jmplen j
        (\snd (nth n ? program ⟨None ?, Comment []⟩))))
    ]
  ].
```

This is a temporary formulation of the main property
(sigma_policy_specification); its main difference from the final version is
that it uses instruction_size_jmplen to compute the instruction size. This
function uses j to compute the span of branch instructions (i.e. it uses the σ
function under construction), instead of looking at the distance between source
and destination. This is because σ is still under construction; later on we will
prove that after the final iteration, sigma_compact_unsafe is equivalent to the
main property.

```
definition sigma_safe :=
 λprefix:list labelled_instruction.λlabels:label_map.λadded:ℕ.
 λold_sigma:ppc_pc_map.λsigma:ppc_pc_map.
 ∀i.i < |prefix| → let ⟨pc,j⟩ :=
  bvt_lookup ... (bitvector_of_nat ? i) (\snd sigma) ⟨0,short_jump⟩ in
  let pc_plus_jmp_length := bitvector_of_nat ? (\fst (bvt_lookup ...
   (bitvector_of_nat ? (S i)) (\snd sigma) ⟨0,short_jump⟩)) in
  let ⟨label,instr⟩ := nth i ? prefix ⟨None ?, Comment [ ]⟩ in
   ∀dest.is_jump_to instr dest →
    let paddr := lookup_def ... labels dest 0 in
    let addr := bitvector_of_nat ? (if leb i paddr (* forward jump *)
    then \fst (bvt_lookup ... (bitvector_of_nat ? paddr) (\snd old_sigma)
     ⟨0,short_jump⟩) + added
    else \fst (bvt_lookup ... (bitvector_of_nat ? paddr) (\snd sigma)
     ⟨0,short_jump⟩)) in
    match j with
    [ short_jump ⇒ ¬is_call instr ∧
      \fst (short_jump_cond pc_plus_jmp_length addr) = true
    | absolute_jump ⇒ ¬is_relative_jump instr ∧
      \fst (absolute_jump_cond pc_plus_jmp_length addr) = true ∧
      \fst (short_jump_cond pc_plus_jmp_length addr) = false
    | long_jump ⇒ \fst (short_jump_cond pc_plus_jmp_length addr) = false
      ∧ \fst (absolute_jump_cond pc_plus_jmp_length addr) = false
    ].
```

This is a more direct safety property: it states that branch instructions are encoded properly, and that no wrong branch instructions are chosen.

Note that we compute the distance using the memory address of the instruction plus its size: this follows the behaviour of the MCS-51 microprocessor, which increases the program counter directly after fetching, and only then executes the branch instruction (by changing the program counter again).

```
\fst (bvt_lookup ... (bitvector_of_nat ? 0) (\snd policy)
 ⟨0,short_jump⟩) = 0)
\fst policy = \fst (bvt_lookup ...
 (bitvector_of_nat ? (|prefix|)) (\snd policy) ⟨0,short_jump⟩)
```

These two properties give the values of σ for the start and end of the program; $\sigma(0) = 0$ and $\sigma(n)$, where n is the number of instructions up until now, is equal to the maximum memory address so far.

```
(added = 0 → policy_pc_equal prefix old_sigma policy))
(policy_jump_equal prefix old_sigma policy → added = 0))
```

And finally, two properties that deal with what happens when the previous iteration does not change with respect to the current one. *added* is a variable that keeps track of the number of bytes we have added to the program size by changing the encoding of branch instructions. If *added* is 0, the program has not changed and vice versa.

We need to use two different formulations, because the fact that *added* is 0 does not guarantee that no branch instructions have changed. For instance, it is possible that we have replaced a short jump with an absolute jump, which does not change the size of the branch instruction.

Therefore `policy_pc_equal` states that $old\_sigma_1(x) = sigma_1(x)$, whereas `policy_jump_equal` states that $old\_sigma_2(x) = sigma_2(x)$. This formulation is sufficient to prove termination and compactness.

Proving these invariants is simple, usually by induction on the prefix length.

4.2 Iteration invariants

These are invariants that hold after the completion of an iteration. The main difference between these invariants and the fold invariants is that after the completion of the fold, we check whether the program size does not supersede 64 Kb, the maximum memory size the MCS-51 can address.

The type of an iteration therefore becomes an option type: `None` in case the program becomes larger than 64 Kb, or `Some` σ otherwise. We also no longer use a natural number to pass along the number of bytes added to the program size, but a boolean that indicates whether we have changed something during the iteration or not.

If an iteration returns `None`, we have the following invariant:

```
definition nec_plus_ultra :=
 λprogram:list labelled_instruction.λp:ppc_pc_map.
 ¬(∀i.i < |program| →
  is_jump (\snd (nth i ? program ⟨None ?, Comment []⟩)) →
  \snd (bvt_lookup ... (bitvector_of_nat 16 i) (\snd p) ⟨0,short_jump⟩) =
   long_jump).
```

This invariant is applied to *old_sigma*; if our program becomes too large for memory, the previous iteration cannot have every branch instruction encoded as a long jump. This is needed later in the proof of termination.

If the iteration returns Some σ, the invariants out_of_program_none, not_jump_default, jump_increase, and the two invariants that deal with $\sigma(0)$ and $\sigma(n)$ are retained without change.

Instead of using sigma_compact_unsafe, we can now use the proper invariant:

```
definition sigma_compact :=
 λprogram:list labelled_instruction.λlabels:label_map.λsigma:ppc_pc_map.
 ∀n.n < |program| →
  match bvt_lookup_opt ... (bitvector_of_nat ? n) (\snd sigma) with
  [ None ⇒ False
  | Some x ⇒ let ⟨pc,j⟩ := x in
    match bvt_lookup_opt ... (bitvector_of_nat ? (S n)) (\snd sigma) with
    [ None ⇒ False
    | Some x1 ⇒ let ⟨pc1,j1⟩ := x1 in
      pc1 = pc + instruction_size
        (λid.bitvector_of_nat ? (lookup_def ?? labels id 0))
        (λppc.bitvector_of_nat ?
         (\fst (bvt_lookup ... ppc (\snd sigma) ⟨0,short_jump⟩)))
        (λppc.jmpeqb long_jump (\snd (bvt_lookup ... ppc
         (\snd sigma) ⟨0,short_jump⟩))) (bitvector_of_nat ? n)
        (\snd (nth n ? program ⟨None ?, Comment []⟩))
    ]
  ].
```

This is almost the same invariant as sigma_compact_unsafe, but differs in that it computes the sizes of branch instructions by looking at the distance between position and destination using σ.

In actual use, the invariant is qualified: σ is compact if there have been no changes (i.e. the boolean passed along is true). This is to reflect the fact that we are doing a least fixed point computation: the result is only correct when we have reached the fixed point.

There is another, trivial, invariant if the iteration returns Some σ:

```
\fst p < 2^16
```

The invariants that are taken directly from the fold invariants are trivial to prove.

The proof of nec_plus_ultra works as follows: if we return None, then the program size must be greater than 64 Kb. However, since the previous iteration

did not return `None` (because otherwise we would terminate immediately), the program size in the previous iteration must have been smaller than 64 Kb.

Suppose that all the branch instructions in the previous iteration are encoded as long jumps. This means that all branch instructions in this iteration are long jumps as well, and therefore that both iterations are equal in the encoding of their branch instructions. Per the invariant, this means that $added = 0$, and therefore that all addresses in both iterations are equal. But if all addresses are equal, the program sizes must be equal too, which means that the program size in the current iteration must be smaller than 64 Kb. This contradicts the earlier hypothesis, hence not all branch instructions in the previous iteration are encoded as long jumps.

The proof of `sigma_compact` follows from `sigma_compact_unsafe` and the fact that we have reached a fixed point, i.e. the previous iteration and the current iteration are the same. This means that the results of `instruction_size_jmplen` and `instruction_size` are the same.

4.3 Final properties

These are the invariants that hold after $2n$ iterations, where n is the program size (we use the program size for convenience; we could also use the number of branch instructions, but this is more complex). Here, we only need `out_of_program_none`, `sigma_compact` and the fact that $\sigma(0) = 0$.

Termination can now be proved using the fact that there is a $k \leq 2n$, with n the length of the program, such that iteration k is equal to iteration $k+1$. There are two possibilities: either there is a $k < 2n$ such that this property holds, or every iteration up to $2n$ is different. In the latter case, since the only changes between the iterations can be from shorter jumps to longer jumps, in iteration $2n$ every branch instruction must be encoded as a long jump. In this case, iteration $2n$ is equal to iteration $2n + 1$ and the fixpoint is reached.

5 Conclusion

In the previous sections we have discussed the branch displacement optimisation problem, presented an optimised solution, and discussed the proof of termination and correctness for this algorithm, as formalised in Matita.

The algorithm we have presented is fast and correct, but not optimal; a true optimal solution would need techniques like constraint solvers. While outside the scope of the present research, it would be interesting to see if enough heuristics could be found to make such a solution practical for implementing in an existing compiler; this would be especially useful for embedded systems, where it is important to have as small solution as possible.

In itself the algorithm is already useful, as it results in a smaller solution than the simple 'every branch instruction is long' used up until now—and with only 64 Kb of memory, every byte counts. It also results in a smaller solution than

the greatest fixed point algorithm that `gcc` uses. It does this without sacrificing speed or correctness.

This algorithm is part of a greater whole, the CerCo project, which aims to completely formalise and verify a concrete cost preserving compiler for a large subset of the C programming language. More information on the formalisation of the assembler, of which the present work is a part, can be found in a companion publication [7].

5.1 Related work

As far as we are aware, this is the first formal discussion of the branch displacement optimisation algorithm.

The CompCert project is another verified compiler project. Their backend [5] generates assembly code for (amongst others) subsets of the PowerPC and x86 (32-bit) architectures. At the assembly code stage, there is no distinction between the span-dependent jump instructions, so a branch displacement optimisation algorithm is not needed.

An offshoot of the CompCert project is the CompCertTSO project, who add thread concurrency and synchronisation to the CompCert compiler [12]. This compiler also generates assembly code and therefore does not include a branch displacement algorithm.

Finally, there is also the Piton stack [6], which not only includes the formal verification of a compiler, but also of the machine architecture targeted by that compiler, a bespoke microprocessor called the FM9001. However, this architecture does not have different jump sizes (branching is simulated by assigning values to the program counter), so the branch displacement problem is irrelevant.

5.2 Formal development

All Matita files related to this development can be found on the CerCo website, `http://cerco.cs.unibo.it`. The specific part that contains the branch displacement algorithm is in the `ASM` subdirectory, in the files `PolicyFront.ma`, `PolicyStep.ma` and `Policy.ma`.

References

1. Asperti, A., Sacerdoti Coen, C., Tassi, E., Zacchiroli, S.: User interaction with the Matita proof assistant. Automated Reasoning 39, 109–139 (2007)
2. Dickson, N.G.: A simple, linear-time algorithm for x86 jump encoding. CoRR abs/0812.4973 (2008)
3. Gnu compiler collection 4.7.0. `http://gcc.gnu.org/` (2012)
4. Hyde, R.: Branch displacement optimisation. `http://groups.google.com/group/alt.lang.asm/msg/d31192d442accad3` (2006)
5. Leroy, X.: A formally verified compiler back-end. Journal of Automated Reasoning 43, 363–446 (2009), `http://dx.doi.org/10.1007/s10817-009-9155-4`, 10.1007/s10817-009-9155-4
6. Moore, J.S.: Piton: A mechanically verified assembly language, Automated Reasoning Series, vol. 3. Springer (1996)
7. Mulligan, D.P., Sacerdoti Coen, C.: On the correctness of an optimising assembler for the Intel MCS-51 microprocessor (2012), submitted
8. Robertson, E.L.: Code generation and storage allocation for machines with span-dependent instructions. ACM Trans. Program. Lang. Syst. 1(1), 71–83 (Jan 1979), `http://doi.acm.org/10.1145/357062.357067`
9. Small device C compiler 3.1.0. `http://sdcc.sourceforge.net/` (2011)
10. Sozeau, M.: Subset coercions in Coq. In: TYPES. pp. 237–252 (2006)
11. Szymanski, T.G.: Assembling code for machines with span-dependent instructions. Commun. ACM 21(4), 300–308 (Apr 1978), `http://doi.acm.org/10.1145/359460.359474`
12. Ševčik, J., Vafeiadis, V., Zappa Nardelli, F., Jagannathan, S., Sewell, P.: Relaxed-memory concurrency and verified compilation. SIGPLAN Not. 46(1), 43–54 (Jan 2011), `http://doi.acm.org/10.1145/1925844.1926393`