

Schleifenausrollen mit nicht konstanten Grenzen in FIRM

Bachelorarbeit von

Adrian E. Lehmann

an der Fakultät für Informatik



Erstgutachter: Prof. Dr.-Ing. Gregor Snelting
Zweitgutachter: Prof. Dr. rer. net. Bernhard Beckert
Betreuende Mitarbeiter: M. Sc. Andreas Fried
Abgabedatum: 30. September 2019

Zusammenfassung

Das Schleifenausrollen ist der Prozess, bei dem der Körper einer Schleife mehrfach dupliziert wird, um die Anzahl der bedingten Sprünge zu reduzieren. Während Schleifen mit konstanten Grenzen leicht ausrollbar sind, erweist sich das Ausrollen von Schleifen mit nicht konstanten Grenzen als eine Herausforderung.

In diesem Fall ist es möglich, spekulativ eine Schleife um einen bestimmten Faktor auszurollen, sofern sichergestellt wird, dass die ausgerollte Schleife so oft wie möglich, aber nicht öfter als die ursprüngliche Schleife ausgeführt wird. Um die daraus resultierende Diskrepanz auszugleichen, wird sogenannter Fixup Code erstellt, der für die Ausführung der restlichen Iterationen verantwortlich ist. Um den oben genannten Fixup Code zu erstellen, werden zwei Ansätze verfolgt: Zum einen kann die Originalschleife dupliziert werden und zum anderen, eine generalisierte Form von Duff's Device verwendet werden.

Bei der experimentellen Bewertung des Ansatzes können wir keine Beschleunigung außerhalb des Messfehlers feststellen, obwohl die Anzahl der ausrollbaren Schleifen verdoppelt wurde (auf 10% aller Schleifen). Die Ergebnisse deuten darauf hin, dass das Schleifenausrollen, ohne weitere Optimierung, keinen (signifikanten) Nutzen zu haben scheint.

Abstract

Loop unrolling is the process of duplicating a loop's body multiple times to reduce the number of conditional jumps. While we can easily unroll loops with constant bounds, unrolling loops with non-constant bounds proves to be a different challenge.

In this case, we can speculatively unroll a loop by a given factor, while making sure that the unrolled loop runs as often as possible, but less than or equal times to the original loop. To compensate for the discrepancy, we create so-called fixup code, which is responsible for running the remaining iterations. We employ two different strategies for creating the aforementioned fixup code: One where we duplicate the original loop and another where we utilize a generalized form of Duff's device.

When experimentally evaluating the approach, even though we double the amount of unrollable loops to now unroll more than 10% of all loops, we cannot note any speed-up outside of the margin of error. The results indicate that loop unrolling does not seem to have a (significant) benefit without further optimization.

Contents

1	Introduction	7
2	Basics and related work	9
2.1	Compiler	9
2.2	Basic blocks and control-flow	9
2.3	Loops	10
2.4	Single-Static-Assignment (SSA)	10
2.5	Loop-Closed-Single-Static-Assignment (LCSSA)	11
2.6	libFIRM	11
2.7	Loop unrolling	12
2.8	Duff's device	14
2.9	Overflow detection	16
3	Design and implementation	17
3.1	Determining unrollability	18
3.2	Unrolling	19
3.3	Fixup strategies	21
3.3.1	Updating the loop condition	24
3.3.2	Generalized Duff's device	27
3.3.3	Loop duplication	30
3.4	Selecting an unroll-factor	33
4	Evaluation	37
4.1	Unrollability	37
4.2	Performance	38
4.2.1	Duff's device fixup	39
4.2.2	Loop fixup	39
5	Conclusion	47

1 Introduction

When developers craft code, there is a need to convert it from a human-readable high-level language into a machine-understandable language, called assembly. In order to do this programmers run a *compiler* that checks the code for multiple sources of errors, and, if the code is correct, converts it into an executable file. While converting the program into machine code the compiler optimizes the code. This is only beneficial to the developer, as it ensures that his/her application, in the end, runs faster and/or requires fewer system resources. A simple example of an optimization is constant folding [1], where the compiler analyzes code and precalculates all constant values, instead of letting the operations on them waste valuable runtime to calculate each time the application runs. An example can be seen below: Humans immediately see that in Figure 1.1a b is always equal to 9 and using constant folding the compiler will be able to also perform this precomputation. The result of this can be seen in Figure 1.1b. Therefore, the optimization reduces the runtime of the (admittedly small) program, by virtue of there being one less calculation required. Simplifying just one expression seems (and for a matter of fact is) quite useless, but in real-world code, an optimization like this can be applied on numerous similar expressions and hence noticeably improve the final product.

$a \leftarrow 7$	$a \leftarrow 7$
$b \leftarrow a + 2$	$b \leftarrow 9$
(a) Sample code snippet for constant folding	(b) Code with constants folded

Figure 1.1: Example of constant folding optimization

Of course there is a plethora of possibilities to optimize code. Considering that loops make up approximately 10% of code of many real-world applications¹, they are a natural point to focus optimization efforts upon. Loops can be unrolled fairly straightforward, if you know how often they are iterated, as is discussed in Chapter 2.

For example, Figure 1.2a can be easily converted to Figure 1.2b, while keeping all semantics intact. In Figure 1.3 things get trickier, since the (exact) value of N is unknown, simply unrolling a loop by copying its body a fixed number of times does not preserve the original semantics.

In Chapter 2 fundamentals for working with these loops are discussed, which are enhanced, juxtaposed and integrated in Chapter 3. Finally, in Chapter 4 we evaluate the approach experimentally to see whether it yields a tangible benefit.

¹Measured using gcc (spec2006): 8.6% of FIRM nodes are in loops

<pre><i>i</i> ← 0 while <i>i</i> < 5 do PRINT(<i>i</i>) <i>i</i> ← <i>i</i> + 1 end while</pre>	<pre>PRINT(0) PRINT(1) PRINT(2) PRINT(3) PRINT(4)</pre>
(a) Loop with constant bounds	(b) Loop with constant bounds unrolled

Figure 1.2: Unrolling a loop with constant bounds

```
i ← 0
N ← FAIRDICEROLL()
while i < N do
  PRINT(i)
  i ← i + 1
end while
```

▷ Random number in [1, 6]

Figure 1.3: Loop without constant bound

2 Basics and related work

2.1 Compiler

The primary function of a compiler is to automatically convert high-level code created by a developer into (optimized) machine code. As a compiler is an inherently large software project, an architecture needs to be chosen that allows for extensions and modifications. Modern compilers mostly follow a layered architecture style: They are each comprised of a front-, middle-, and back-end. In this architecture, the front-end converts the high-level code into an abstract intermediary representation, which is then used by the middle-end for optimizations and transformations. Lastly, the back-end is responsible for converting the optimized intermediary code into instructions for the target system architecture (e.g., RISC-V, x86, ARM, or similar).

2.2 Basic blocks and control-flow

To better handle code and give it a logical structure, most compilers divide code up into so-called *basic blocks*. (Basic) blocks are sets of consecutive operations that do not contain jumps or targets thereof, but rather only jumps connecting them. Therefore, a basic block is either executed entirely or not executed at all.

A usual way to represent this in a human-readable form is to output it as a control-flow-graph (CFG). CFGs depict basic blocks as nodes and jumps between basic blocks as edges. Furthermore, it is a convention in these graphs to have precisely one start-node and one end-node.

We note that CFGs, in general, are cyclical graphs. They are non-cyclical graphs if the original code does not contain any jumps going backward in the control-flow.

Another important concept of CFGs is dominance. To explain this concept, we define a starting block S , and assume there are two (not necessarily different) blocks, present in the CFG, N_1 , and N_2 . With this information, we define dominance as follows:

$$N_1 \text{ dominates } N_2 \iff \forall p \in \text{Paths}(S, N_2) : N_1 \in p$$

In lucid terms, this means N_1 dominates N_2 , iff to get to N_2 from S you have to visit N_1 on the way. It is important to note that a block always dominates itself.

2.3 Loops

We define a loop to be a set of blocks that are all in a cyclical control-flow structure. Formally this can be expressed as:

$$L \text{ is a loop} \iff L \neq \emptyset \wedge \forall n_1, n_2 \in L : \exists \text{Path}(n_1, n_2) \subseteq L$$

Henceforth let L be a loop.

If a loop is completely contained inside another loop, it is said to be nested.

$$L' \text{ is nested in } L \iff L' \subsetneq L \wedge L' \text{ is a loop}$$

If a loop has no nested loops inside of it, we call it an *innermost* loop.

$$L \text{ is innermost loop} \iff \nexists L' : L' \text{ is a nested loop in } L$$

Loops can furthermore have a header, which is the sole entry point into a loop [2] and defined as follows:

$$H \text{ is header of } L \iff H \in L \wedge \forall n \in L : H \text{ dominates } n$$

N.B.: Not all loops have to have a header.

If a loop has a header, its body is the set containing all blocks in the loop, except for the header.

$$B \text{ is body of } L \iff B = L \setminus \{H\}, H \text{ is header of } L$$

2.4 Single-Static-Assignment (SSA)

The *single-static-assignment* (SSA) form is a property of intermediary representations, that requires each value to only be assigned exactly once. Moreover, every value has to be assigned before it is being used [3]. This property mainly implies that the block, which declares a given value v , has to dominate all blocks that use v . This declaration point will be unambiguous across all possible usages.

The SSA form is used to simplify optimizations in the regard that one can be sure that a set point in the code currently defines a given value in use.

Figure 2.1 shows an example program in SSA form. While in the base code x is assigned twice, in the code transformed into SSA form, simply a new value was defined to make sure that each variable is only defined once.

In a loop or a conditional statement, a scenario might arise where a given value could be assigned at multiple locations. In cases like these we can use a Φ -function. A Φ -function is a theoretical construct that returns the correct value depending on the control-flow predecessor.

Figure 2.2 shows an example of the use of a Φ -function, where depending on the control-flow, either m_1 or m_2 are selected.

$x \leftarrow 1$	$x_1 \leftarrow 1$
PRINT(x)	PRINT(x_1)
$x \leftarrow 7$	$x_2 \leftarrow 7$
PRINT(x)	PRINT(x_2)

(a) An example piece of code not in SSA form (b) The same code in SSA form

Figure 2.1: An example program in SSA form

<pre> function MAX($a : \mathbb{N}, b : \mathbb{N}$) $m : \mathbb{N}$ if $a > b$ then $m \leftarrow a$ else $m \leftarrow b$ end if return m end function </pre>	<pre> function MAX($a : \mathbb{N}, b : \mathbb{N}$) $m : \mathbb{N}$ if $a > b$ then $m_1 \leftarrow a$ else $m_2 \leftarrow b$ end if $m \leftarrow \Phi(m_1, m_2)$ return m end function </pre>
--	---

(a) Non-SSA Code for a function that returns the maximum of its parameters

(b) Same Function converted into SSA

Figure 2.2: An example program transformed into SSA form

2.5 Loop-Closed-Single-Static-Assignment (LCSSA)

An extension to the aforementioned SSA form is the *loop-closed-single-static-assignment* (LCSSA) form. A CFG in LCSSA form has all properties that a CFG in SSA form has, and additionally the property that each value assigned in a given loop and used outside of this loop has to be used by a Φ -node in the first block after the loop. This form is used to reduce special casing when transforming loops [2] and is therefore utilized through all of Chapter 3.

To visualize this property, Figure 2.3 depicts its effect.

2.6 libFIRM

libFIRM is a compiler middle- and back-end that takes a graph-based intermediate representation in SSA form, optimizes it, and produces assembly code [5]. Since 1996, Karlsruhe Institute of Technology (KIT) actively develops libFIRM.

A graph in libFIRM contains information about basic blocks, the control-flow, and memory and data dependencies. Basic blocks in libFIRM contain further nodes that are responsible for the control-flow of the program. These are pointed to by (other) basic blocks that are the target of these control-flow operations. The resulting control-flow edges are represented by red edges in visualizations of libFIRM graphs. Any node operating on memory also connects to the previous node operating on

<pre> function FOO $x, y : \mathbb{N}$ repeat if CONDITION then $x_1 \leftarrow 5$ else $x_2 \leftarrow 8$ end if $x \leftarrow \Phi(x_1, x_2)$ until OTHERCONDITION $y \leftarrow x + 3$ end function </pre> <p>(a) Sample loop in SSA form</p>	<pre> function FOO $x, y : \mathbb{N}$ repeat if CONDITION then $x_1 \leftarrow 5$ else $x_2 \leftarrow 8$ end if $x_3 \leftarrow \Phi(x_1, x_2)$ until OTHERCONDITION $x \leftarrow \Phi(x_3)$ $y \leftarrow x + 3$ end function </pre> <p>(b) Same sample loop in LCSSA form</p>
---	--

Figure 2.3: An example program transformed into LCSSA Form (adapted from [4])

memory, so that the node uses the prior state of memory and then provides a new state with its changes. Memory is, like control flow, connected by edges, which are colored blue in graphical representations. Lastly, libFIRM has data dependency edges between nodes, which represent dependencies needed for calculations.

Figure 2.4 portrays an example libFIRM graph of the program initially shown in Figure 2.2. It is especially to be noted that the graph has both memory, data and control-flow edges, and is in SSA form, since it contains a Φ -node.

Another set of functionality that libFIRM provides is loop information. libFIRM will not only (if applicable) be able to map blocks to their respective loops and vice-versa, but also has information on loop nesting structure. Thus, one can quickly determine, whether a loop is an innermost loop [5]. Further, libFIRM also allows for finding the header of a loop, given that it has a header [2]. Since loop unrolling involves node duplication in the implementation that we will use [2], it is worth mentioning, that upon a node will has a field called `link` that will be set to reference the copied node and vice-versa. This allows for easy access of the duplicated and original nodes in later algorithms.

2.7 Loop unrolling

Loop unrolling is a compiler optimization that attempts to duplicate the loop body and to reduce the controlling instructions, such as the loop condition or repetitive arithmetic [6].

Figure 2.5 shows a pseudo-code example of unrolling a simple loop with a factor (the number of times the body is copied) of four. It is to be noted that the loop condition has to be checked less often, on account of each loop iteration being four times as long as in the original program.

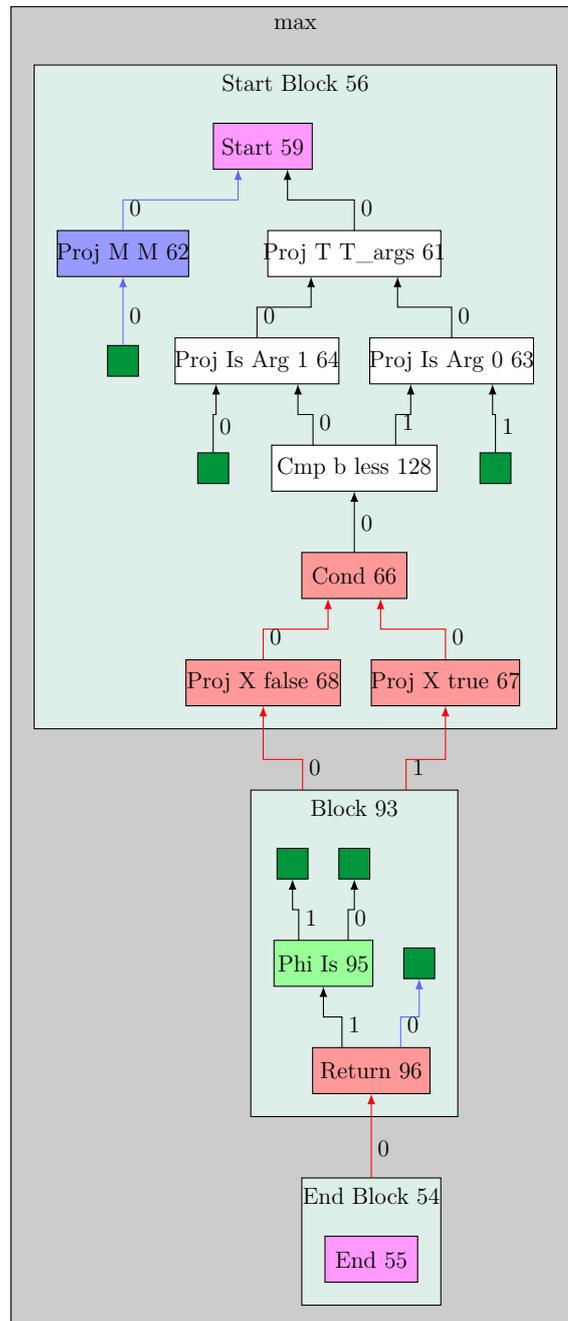


Figure 2.4: A libFIRM graph of a function that returns the larger argument

<pre>function FOO $i \leftarrow 0$ while $i < 16$ do PRINT(i) $i \leftarrow i + 1$ end while end function</pre>	<pre>function FOOUNROLLED $i \leftarrow 0$ while $i < 16$ do PRINT(i) PRINT($i + 1$) PRINT($i + 2$) PRINT($i + 3$) $i \leftarrow i + 4$ end while end function</pre>
(a) A function with a simple loop inside of it	(b) A function with the same loop unrolled

Figure 2.5: A simple loop with constant bound unrolled

Furthermore, it could also be used to vectorize the code, eliminate repeating conditions, and for many other following optimizations [7]. A negative side effect of loop unrolling is that the binary size increases and that there could be more pressure on the code cache and registers, causing more spilled values [8].

libFIRM supports a restricted form of loop unrolling for loops that have static bounds and increments [9]. This optimization was recently improved but now requires the intermediary representation to be in LCSSA form, which means libFIRM's intermediary representation has to be converted into LCSSA form prior to the optimization running [2]. Though this optimization had no preconditions and merely duplicated the loop in the hope a later optimization would remove the duplicated headers, so that it resembles an unrolled loop of our definition, in which only the body (and not the header) is duplicated. In this form the amount of conditional jumps does not decrease in most cases, since the later optimization removing headers would not trigger. Hence, the application of this method showed no improvements, which caused preconditions for the later removal of excess headers to be added [10]. libFIRM now only unrolls loops, similar to the loop in Figure 2.5, with static bounds and for which a constant bit analysis can remove excess headers. The benefits of the now changed optimization were also very slim, likely since the requirements for a loop to now be unrollable are very strict. With these restrictions in place, only approximately 5% of the innermost loops can be unrolled¹.

2.8 Duff's device

A common problem with the loop unrolling shown in Figure 2.5 is that it requires the number of iterations to be constant and divisible by the unroll-factor. A way to tackle this issue is to use a construct known as Duff's device: It preemptively unrolls a loop with a given factor and uses *fixup* code to ensure that the remaining

¹Measured in `spec2006`

<pre> function FOO($N : \mathbb{N}$) $i \leftarrow 0$ while $i < N$ do PRINT(i) $i \leftarrow i + 1$ end while end function </pre> <p>(a) An example function with a loop</p>	<pre> function FOODUFFED($N : \mathbb{N}$) $i \leftarrow 0$ switch $N \bmod 4$ do case 3 PRINT(i) $i \leftarrow i + 1$ ▷ Fall-through case 2 PRINT(i) $i \leftarrow i + 1$ ▷ Fall-through case 1 PRINT(i) $i \leftarrow i + 1$ while $i < N$ do PRINT(i) PRINT($i + 1$) PRINT($i + 2$) PRINT($i + 3$) $i \leftarrow i + 4$ end while end function </pre> <p>(b) A function with the loop unrolled using Duff's device</p>
--	---

Figure 2.6: A simple loop unrolled using Duff's device

iterations are completed [11]. Mathematically this means the construct executes the loop body $\left\lfloor \frac{M}{f} \right\rfloor \cdot f + M \bmod F = M$ times, where M is the number of total times the loop body would be executed without the transformation and where f is the unroll-factor. This is due to the fact that mod is defined as:

$$x \bmod y = x - \left\lfloor \frac{x}{y} \right\rfloor \cdot y$$

If we substitute M for x and f for y and rearrange for M , we get said form.

Figure 2.6 shows an example of unrolling a loop with a non-divisible bound using a factor of eight². Duff's device copies the loop body eight times and to ensure that the number of executions is correct, the first time around the code jumps to the corresponding instruction, depending on the need for fixup code.

Many compilers, such as GCC [12], use Duff's device for unrolling loops and improving performance while keeping code size relatively small. Further libFIRM previously utilized Duff's device for unrolling loops with static bounds, but for which no unroll-factor could be determined [9].

²The original Duff's device used special C syntax to entangle the switch statement and loop [11]

2.9 Overflow detection

When subtracting (or adding) two numbers that are integer-like the operation might cause an over- or underflow, because the integer is of a fixed bit two's complement representation. Due to the therefore inherent limitation to the range of possible values, this problem is unavoidable, yet detectable.

In the following, we will take t_{min} and t_{max} to be the lower and upper bound of an integer representation. Algorithm 2.1 [13] shows a way to detect whether an overflow or underflow occurs for an operation $x - a$, where $x, a \in \mathbb{Z} \cap [t_{min}, t_{max}]$, by checking whether the result increased or decreased relative to the bounds and comparing it to the expectation.

Algorithm 2.1 Algorithm that detects whether the operation $x - a$ will go out of the integer boundaries

```
function SUBTRACTIONWILLLEAVEBOUNDS( $x, a : \mathbb{Z} \cap [t_{min}, t_{max}]$ )  
  overflow  $\leftarrow$  HIGHESTBITSET( $x$ )  $\wedge$  ( $a > t_{max} + x$ )  
  underflow  $\leftarrow$  ( $x > 0$ )  $\wedge$  ( $a < t_{min} + x$ )  
  return overflow  $\vee$  underflow  
end function
```

3 Design and implementation

In order to unroll a loop with non-static bounds, this thesis follows a specific approach: First, we check whether we are able to unroll the loop. Section 3.1 describes the conditions necessary and how we check them. If we determine a loop to be unrollable, we will unroll it with the unrolling process covered in Section 3.2. Once this process is complete, the loop condition of the unrolled loop will be adapted to make sure it runs less than or equal times compared to the original loop. This is described in Section 3.3.1. After that, we will create the *fixup code*¹, as described in Section 3.3.

It is to be noted that in terms of actually implementing this procedure, we will create the fixup code *before* unrolling the loop. While this order seems counter-intuitive, we chose it in order to simplify the implementation of loop duplication, as described in Section 3.3.3.

Henceforth, we assume loops to be in the form of the loop shown in Figure 3.1. In the reference loop *cmp* refers to a comparison that can be one of the following: $<, >, \leq, \geq$. Further, $I \in \mathbb{Z}$ refers to the starting value, $N \in \mathbb{Z}$ to the bound, and $c \in \mathbb{Z} \setminus \{0\}$ to the increment² of such a loop. We select this form in view of the fact that most loops follow the form of using a counter or iterating over a given container, which condenses down to this form. Furthermore, this form allows for many arithmetic properties to be used, as seen in Section 3.3.

```
function FOO( $I \in \mathbb{Z}, N \in \mathbb{Z}, c \in \mathbb{Z} \setminus \{0\}, cmp \in \{<, >, \leq, \geq\}$ )  
   $i \leftarrow I$   
  while  $i$  'cmp'  $N$  do  
    DOSOMETHING  
     $i \leftarrow i + c$   
  end while  
end function
```

Figure 3.1: A general form of loop starting at I and counting in increments of c up to N

¹The term *fixup code* describes that code that has to be added to account for cases where the number of times the loop is executed modulo the unrolling factor is not equal to zero.

²N.B.: c may be negative and could hence also be a decrement

3.1 Determining unrollability

Given that the primary goal of any optimization is to conserve semantics, most optimizations are based upon assumptions. These assumptions will be assured, by checking corresponding preconditions before the optimization is applied, so that its transformed product will be semantically equivalent.

In the case of loop unrolling, we laid out the structure of the targeted loops in Figure 3.1. This section formalizes the resulting requirements and extends them, such that the further unrolling process conserves semantics.

Firstly, in view of the fact that we use the existing loop unrolling functionality as a sub-step (see Section 3.2), it needs to be ensured that the respective libFIRM-graph is in LCSSA form. We accomplish this by using the existing mechanics [2]. While it is a preliminary step, assuring LCSSA form can never be a hindrance to unrolling, since it is possible to convert any given libFIRM graph into LCSSA form. Due to the restrictions of the existing loop unrolling mechanism, a loop must also be the innermost loop, meaning it does not have any nested loops inside of it. Nested loops inherently cause larger code sizes and hardly saves jumps, since most jumps will occur in the inner loops. Therefore, the restriction will in practice most likely not harm performance. Given these conditions are met, we use the mechanics, described in Section 3.4, for determining if and how a given loop should be unrolled based on size.

Moreover, in order for loops to be in the form described in Figure 3.1, loops have to have a header, which itself controls the control flow by comparing a counter to a bound, using any of the four allowed comparison types. The header is the only point in the loop from which the loop can exit; meaning there are no conditionals in the body that allow the control flow to leave the loop. This primarily requires there not to be any `break`-like structures.

Seeing that there is an explicit entry point for the loop, the header, there are no preconditions for I , since it is therefore only evaluated once in a block dominating the header, but inherently not determining of the control flow after the initial evaluation. On the contrary, N , the bound, has to be loop-invariant, which means that it may not change through the entire evaluation of the loop, because it is checked against i in every iteration. As an example, consider a loop, such as the one in Figure 3.1, replacing `DoSomething` with $N = \text{randomNumber}$. If we now execute the body $f > 1$ times consecutively, we will effectively lose $f - 1$ checks. Assume that initially $I = 0, c = 1, f = 2, N = 2$, and assume in the first execution in the loop body N is set to 0 by chance, whereas in the second iteration it is set to 7. Now given that when unrolling the condition is removed for the entering the second body, the loop body would at least four times, which does not conserve semantics, as it should only be executed once. Concluding, only if N is loop-invariant, the bound checks can be performed less often, while keeping the original semantics intact.

If N is constant it is obviously loop-invariant, but what if it is the result of a function call or of a load from memory? For the case that N is function call, the called function must be pure (i.e., not have any side-effects), and only have

loop-invariant arguments, seeing that the call is then by definition loop-invariant itself.

In case that N is being loaded from memory, stricter conditions have to be met. All stores within the loop must be sure not to alias the memory location of N . Further, any calls must either invoke functions known at compile time and none of these may contain aliasing stores or have aliasing parameters. Otherwise, the loop cannot be unrolled with a loaded bound, due to these called functions potentially modifying N .

Lastly, the unroll-factor – meaning how often the loop body is copied inside the unrolled loop – f is selected (see Section 3.4) and hence known at compile time. We can therefore further restrict the increment c , such that $t_{min} \leq c \cdot f \leq t_{max}$, where t_{min} is the minimum value of the integer type of c and t_{max} the respective maximum. Hence, we prevent $c \cdot f$ from overflowing, which will turn out to be important in Sections 3.3.1 and 3.3.2, and further discussed there. In order to assure this property, we have to force c to be a compile-time constant (which inherently is loop invariant). Even though the restrictions on c seem comparatively tight, in real-world code (gcc, spec2006) only approximately 1.2% of loops that meet the previous conditions are not unrollable because of the restriction that $t_{min} \leq c \cdot f \leq t_{max}$.

It is worth mentioning that the unrollability with the method above is only checked if the current loop unrolling mechanism [2] determines that the current unrolling process cannot be applied. We chose this design, for the reason that statically unrolling without any further fixup code inherently simplifies the control flow and hence should yield better, or (at least) equal, performance.

3.2 Unrolling

To get started with unrolling loops that have unknown bounds, we unroll them by a given factor without considering whether the transformation is semantically invariant. Semantic equivalence, which is broken due to the failure to consider how the factor relates to the original amount of iterations, will be restored in Section 3.3.

libFIRM already provides an unrolling mechanism for unrolling a loop with a given factor f [2].³ To avoid code duplication, we will use be utilizing this solution.

Further, figures 3.2 and 3.3 show a libFIRM graph of a loop that is to be unrolled or is unrolled using a factor of two, respectively. Especially to be noted is that in Figure 3.3 we duplicate the loop header, and that hence the number of conditional jumps did not decrease through the loop unroll. With the previous usage, this was not an issue, because libFIRM would automatically remove these excess headers using its constant bit analysis. Unfortunately though in the use cases of an unknown bound, the constant bit analysis does not suffice. This is due to the fact that the additional semantics, meaning that we are sure not to have to exit the unrolled loop from its body at any time, that are implicitly affixed to

³N.B.: All following operations preserve the LCSSA property of the code.

the transformed loop, cannot be recognized by libFIRM. Therefore, the need to manually prune the graph to remove the excess headers arises. We accomplish this by using Algorithm 3.1. First, we rewire all Φ -nodes in the excess header, such that all in-loop nodes depending on any given Φ -node each get the in-loop predecessors of the Φ -node as predecessors themselves, while the Φ -node falls into desuetude and will therefore be automatically removed by a later optimization. We apply the same transformations to the descendants of the block itself.

Algorithm 3.1 Pruning excess headers after unrolling

```

function PRUNEEXCESSHEADER(copiedHeader : Block)
  for all phi  $\in$  copiedHeader do
    PRUNEPHI(phi, copiedHeader)
  end for
  for all post  $\in$  copiedHeader.descendants do
    post.predecessors  $\leftarrow$  (post.predecessors  $\setminus$  {copiedHeader})  $\cup$ 
    {b | b  $\in$  copiedHeader.predecessors, b.loop = copiedHeader.loop}
  end for
end function

function PRUNEPHI(phi :  $\Phi$ -node, copiedHeader : Block)
  for all out  $\in$  phi.descendants do
     $\triangleright$  out is ensured to be  $\Phi$ -node by the LCSSA construction algorithm [2]
    if out.block  $\neq$  copiedHeader then
      out.predecessors  $\leftarrow$  (out.predecessors  $\setminus$  {phi})  $\cup$ 
      {n | n  $\in$  phi.predecessors, n.loop = out.loop}
    end if
  end for
end function

```

Algorithm 3.2 Pseudo-code for the existing unrolling mechanism [2]

```

function UNROLLEXISTING(factor :  $\mathbb{N}_{>1}$ , loop : Loop)
  ASSURELCSSA(loop)
  for all block  $\in$  loop do
    for i  $\in$  {1..(factor - 1)} do
      DUPLICATEBLOCK(block)
    end for
  end for
  REWIREDUPLICATEDBLOCKS  $\triangleright$  Attach blocks to form unrolled structure
   $\triangleright$  loop is still in LCSSA form after unrolling
end function

```

3.3 Fixup strategies

In Section 3.2 we discussed the unrolling process. There, we did not consider the fixup code needed, but instead plainly focused on unrolling the loop. Firstly, we will now focus on making the loop run less than or equal times compared to the original loop in section Section 3.3.1. Less-than or equal is not good enough though, we want our transformed loop to run exactly as often as the original loop. Therefore, we will create fixup code, as discussed in this section and its subsections. Section 3.3.2 uses a generalized version of Duff's device to create the required fixup code, whereas in Section 3.3.3 a copy of the original loop will be used. After that, in Section 4.2, we evaluate which approach yields faster binary run-times.

To see the reason why we need fixup code and to understand what is required of it, we formally lay out conditions that need to be met in Equations (3.1) through (3.7).

Let $M \in \mathbb{N}_0$ be the number of times a loop runs before the transformation; $M_{\text{loop}} \in \mathbb{N}_0$, $M_{\text{fixup}} \in \mathbb{N}_0$ the number of times the unrolled body will run in the unrolled loop, or the fixup code respectively, after the transformation. Further, the unroll-factor will be again denoted by $f \in \mathbb{N}$, $f > 1$. Henceforth, we will assume all arithmetic operations to be integer operations for integers in the interval $[t_{\min}, t_{\max}]$. Please note that unmarked integer division will be assumed to round towards zero: For example $\frac{5}{3} \stackrel{\text{integer division}}{=} 1$. Another convention we will introduce is that any interval will be integral, meaning it will only contain integers. Additionally $x \mp y$ is

henceforth defined as
$$\begin{cases} x + y & , x < 0 \\ x - y & , x > 0 \end{cases}$$

We will now lay out properties that form the basis of further arithmetic considerations. The primary identity that is to be conserved, to retain the original semantics, is shown below in Equation (3.1). Since we know our original loop ran M times, we know that our transformed loop and the fixup code must in total also run M times.

$$M = M_{\text{loop}} + M_{\text{fixup}} \quad (3.1)$$

In order to use Duff's device, we need to restrict the amount of times the fixup code needs to run. With the requirements to preserve the semantics in mind, we will maximize M_{loop} and minimize M_{fixup} .

$$M_{\text{loop}} \stackrel{\text{integer division}}{=} \frac{M}{f} \cdot f \quad (3.2)$$

$$\stackrel{\text{integer division}}{\in}]M - f, M]$$

By construction of the unrolled loop, Equation (3.2) is always true, as the unrolled loop tries to run as often as possible, while running less than or equal times compared to the original loop.

Proof. To prove the conjecture of Equation (3.2), assume for contradiction

$$M_{\text{loop}} = M - f - b, b \in [0, f[$$

and hence

$$M_{\text{loop}} \leq M - f \Rightarrow M_{\text{loop}} \notin]M - f, f],$$

then by rerunning the unrolled again the body would be executed f times causing $M_{\text{loop}} = M - b \in]M - f, f]$, which would be a contradiction of the assumption. We then induct this pattern for $M_{\text{loop}} = M - nf - b, n \in \mathbb{N}_+, b \in [0, f[$. In these cases, the loop must merely be iterated multiple times. ■

As the loop runs as often as possible, the fixup code will always run less than f times.

$$M_{\text{fixup}} \in [0, f[\tag{3.3}$$

Proof. Conjecture: $M_{\text{fixup}} \in [0, f[$.

Assume for contradiction $M_{\text{fixup}} = f' > (f - 1)$

$$\begin{aligned} M_{\text{loop}} + M_{\text{fixup}} &\stackrel{3.2}{\geq} M - (f - 1) + f' \\ &> M - (f - 1) + (f - 1) \\ &= M \stackrel{3.1}{\neq} \end{aligned}$$

■

For the following mathematical considerations, we need to round away from zero in integer division. Lemma 1 describes how this can be accomplished.

Lemma 1. Given $Y \neq 0$: $\left\lceil \frac{X}{Y} \right\rceil = \frac{X + (Y \mp 1)}{Y}$

Proof. To prove Lemma 1, we consider the cases $X \bmod Y = 0$ and $X \bmod Y \neq 0$. Further, we assume $Y > 0$, since the proof for $Y < 0$ can be performed analogously. Consider the case that $X \bmod Y = 0$. In this case $\exists n \in \mathbb{N} : n \cdot Y = X$ and $\left\lceil \frac{X}{Y} \right\rceil = \frac{X}{Y} = n$ (\star).

$$\begin{aligned} \Rightarrow \frac{X + (Y - 1)}{Y} &= \frac{n \cdot Y + (Y - 1)}{Y} \\ &= \frac{(n + 1) \cdot Y - 1}{Y} \\ &\quad \underbrace{\hspace{1.5cm}}_{< \frac{(n+1) \cdot Y}{Y}} \\ &\stackrel{\text{integer division}}{=} \frac{n \cdot Y}{Y} \\ &= n \\ &\stackrel{\star}{=} \left\lceil \frac{X}{Y} \right\rceil \end{aligned}$$

Now consider $X \bmod Y \neq 0$. In this case $\exists n \in \mathbb{N} : n \cdot Y < X < (n + 1) \cdot Y$ and $\left\lceil \frac{X}{Y} \right\rceil \stackrel{\text{integer division}}{=} n + 1$.

$$\begin{aligned}
 &\Rightarrow n \cdot Y + (Y - 1) < X + (Y - 1) < (n + 1) \cdot Y + (Y - 1) \\
 &\Rightarrow (n + 1) \cdot Y - 1 < X + (Y - 1) < (n + 2) \cdot Y - 1 \\
 &\stackrel{\text{integers}}{\Rightarrow} (n + 1) \cdot Y \leq X + (Y - 1) < (n + 2) \cdot Y \\
 &\stackrel{Y > 0}{\Rightarrow} \frac{(n + 1) \cdot Y}{Y} \leq \frac{X + (Y - 1)}{Y} < \frac{(n + 2) \cdot Y}{Y} \\
 &\Rightarrow \frac{X + (Y - 1)}{Y} \stackrel{\text{integer division}}{=} n + 1 \\
 &= \left\lceil \frac{X}{Y} \right\rceil
 \end{aligned}$$

■

Using Lemma 1, we use the loop parameters I , N and c to calculate the total number of loop iterations.

$$M = \left\lceil \frac{N - I}{c} \right\rceil \stackrel{\text{integer division}}{=} \frac{N - I + (c \mp 1)}{c} \quad (3.4)$$

With this result, we can then calculate M accurately from the structure of the loop, since M is not directly known. As I is not the initial value for the fixup's counter, we will need to calculate the initial value based on known parameters.

$$i_{\text{post loop}} = c \cdot M_{\text{loop}} + I \quad (3.5)$$

We can then use the last two equations to calculate M_{fixup} based on only quantities that are known at either compile-time or at run-time.

$$\begin{aligned}
 M_{\text{fixup}} &\stackrel{3.1}{=} M - M_{\text{loop}} \\
 &\stackrel{3.4}{=} \frac{N - I + (c \mp 1)}{c} - M_{\text{loop}} \\
 &\stackrel{3.5}{=} \frac{N - I + (c \mp 1)}{c} - \frac{i_{\text{post loop}} + I}{c} \\
 &= \frac{N - I + I - i_{\text{post loop}} + (c \mp 1)}{c} \\
 &= \frac{N - i_{\text{post loop}} + (c \mp 1)}{c}
 \end{aligned} \quad (3.6)$$

As the above equation has a costly division operation in it, we will rearrange it, such that it never needs to be computed at runtime.

$$(3.6) \stackrel{\text{integer division}}{\iff} M_{\text{fixup}} \cdot c = N - i_{\text{post loop}} + (c \mp 1) \stackrel{3.3}{\in} \begin{cases} [0, c \cdot f[& , c > 0 \\]c \cdot f, 0] & , c < 0 \end{cases} \quad (3.7)$$

Equation (3.7) is especially significant in the construction of the generalization of Duff's device, as seen in Section 3.3.2.

3.3.1 Updating the loop condition

In the following Sections 3.3.2, and 3.3.3 we will use that $M_{\text{loop}} = \frac{M}{f} \cdot f$. Though, when unrolling (as described in Section 3.2), the original bound (N) is kept. Unfortunately, this does not guarantee M_{loop} to be correct, as made clear by an example where a loop with $I = 0, N = 3, c = 1, f = 2, cmp = <$ is unrolled. In this example, this would yield the following: $M_{\text{loop}} = 4 > M = 3$, due to the fact that after the first iteration of the unrolled loop $i = 2 < 3 = N$. To combat this, we set the bound of the unrolled loop to $\hat{N} = N - c \cdot (f - 1)$. Now we will prove the conjecture that using the bound \hat{N} , the unrolled loop runs M_{loop} times, given the operation to calculate \hat{N} will not over- or underflow.

Proof. Let M'_{loop} be the number of times then unrolled loop with bound \hat{N} runs. The proof is complete, iff $M'_{\text{loop}} = M_{\text{loop}}$. Note that $c \cdot f$ cannot overflow, as per preconditions (\star).

$$\begin{aligned} M'_{\text{loop}} &\stackrel{\text{loop construction}}{=} \left\lfloor \frac{\hat{N} - I}{c \cdot f} \right\rfloor \cdot f \\ &\stackrel{\text{integer division}}{=} \frac{\hat{N} - I + (c \cdot f \mp 1)}{c \cdot f} \cdot f \\ &= \frac{N - c \cdot (f - 1) - I + (c \cdot f \mp 1)}{c \cdot f} \cdot f \\ &= \frac{N - c \cdot f + c - I + c \cdot f \mp 1}{c \cdot f} \cdot f \\ &= \frac{N - I + c \mp 1}{c \cdot f} \cdot f \\ &\stackrel{\star}{=} \frac{N - I + c \mp 1}{f} \cdot f \\ &\stackrel{\text{integer division}}{=} \left\lfloor \frac{N - I}{c} \right\rfloor \cdot f \\ &\stackrel{3.4}{=} \frac{M}{f} \cdot f \\ &\stackrel{3.2}{=} M_{\text{loop}} \end{aligned}$$

■

Therefore we change the header condition of the unrolled loop to i ‘*cmp*’ \hat{N} . Note that even though, we are calculating the rounding to a multiple of f without the need for a slow division operation.

Figure 3.4 shows the comparison of the original condition to the changed header condition, for the loop shown in Figure 3.5. It is to be noted that the graph with bound \hat{N} can be constant folded to the same size, as the original header.

(a) The original condition with bound N (b) The changed condition with bound \hat{N}

Figure 3.4: The change of header condition for Figure 3.5. Please note that through an implicit optimization by libFIRM, the comparison has been changed to \leq and hence N to 28.

```

i ← 0
while i < 29 do
  PRINT>HelloWorld)
  i ← i + 3
end while

```

Figure 3.5: An example loop, for which the unrolling process will be explained

We know that $c \cdot f$ cannot over- or underflow, as per the preconditions laid out in Section 3.1. Since $f > 0$, $c \cdot (f - 1)$ will therefore also not overflow. Though, $N - c \cdot (f - 1)$ can still over- or underflow, due to the subtraction of $c \cdot (f - 1)$ from N , and hence there is nevertheless a possibility to construct an example where this change does not conserve semantics.

Suppose the datatype of a loop with parameters $N = 2$, $I = 0$, $c = 1$, $cmp = <$ is a 32-bit unsigned integer, and we unroll this loop by a factor of four. In this case $\hat{N} = 2 - 1 \cdot (4 - 1) = -1 \stackrel{\text{unsigned integer}}{=} t_{max}$. Thus, the loop would run $t_{max} > 2$ times.

To circumvent this problem, we use Algorithm 2.1 from Section 2.9 as a check for over- or underflows of the operation. We implement this check by placing a block between the header and its predecessors. If an under- or overflow is detected, the control flow will jump directly the fixup code. Otherwise, it will route the control flow to the header and let the loop progress as normal, given that \hat{N} now restores semantics in the header. Algorithm 3.3 shows the creation of this structure in libFIRM.

Algorithm 3.3 Algorithm that creates the check to ensure \hat{N} does not over- or underflow

```
function CREATEPREHEADER(firstFixupBlock : Block, loop : Loop)
  header  $\leftarrow$  loop.header
  pre  $\leftarrow$  NEWEMPTYBLOCK
  pre.predecessors  $\leftarrow$  {node | node  $\in$  header.predecessors, node  $\notin$  loop}
  for all phi  $\in$  header.phis do
    phi'  $\leftarrow$  NEWPHIINBLOCK(pre)
    phi'.predecessors  $\leftarrow$  {node | node  $\in$  phi.predecessors, node.block  $\notin$  loop}
    phi.predecessors  $\leftarrow$  {phi'}  $\cup$  {node | node  $\in$  phi.predecessors, node.block  $\in$ 
loop}
    for all succ  $\in$  phi.successors do
      if succ.block dominated by firstFixupBlock then
        succ.predecessors.prepend(pre)
      end if
    end for
  end for
  (trueExit, falseExit)  $\leftarrow$  CREATEOVERFLOWCONDITION  $\triangleright$  See Section 2.9
  firstFixupBlock.predecessors.prepend(falseExit)
  header.predecessors  $\leftarrow$  {trueExit}  $\cup$  {node | node  $\in$  header.predecessors  $\cap$ 
loop}
end function
```

3.3.2 Generalized Duff's device

Section 2.8 describes the original version of Duff's device. The problem with this initial approach is that it assumes $c = 1$, even though c is defined as any non-zero integer in the considered loops. Therefore, a need for generalization arises.

Using equations 3.1 through 3.7 and the general idea of Duff's device (see Section 2.8), we will create fixup code in form of a generalized Duff's device. The structure of this fixup code can be seen in Figure 3.6, which we then practically implement, using Equation (3.7), as shown in Figure 3.7.

```

switch  $M_{\text{fixup}}$  do
  case  $f - 1$ 
    BODY ▷ Fall-through
  case  $f - 2$ 
    BODY ▷ Fall-through
  ...
  case 1
    BODY

```

Figure 3.6: Generalized Duff's device fixup code based on M_{fixup}

```

switch  $N - i_{\text{post loop}} + (c \mp 1)$  do
  case  $[c \cdot (f - 1), c \cdot f[$  ▷ flip bounds for  $c < 0$ 
    BODY
  case  $[c \cdot (f - 2), c \cdot (f - 1)[$  ▷ flip bounds for  $c < 0$ 
    BODY
  ...
  case  $[c \cdot 1, c \cdot 2[$  ▷ flip bounds for  $c < 0$ 
    BODY

```

Figure 3.7: Generalized Duff's device fixup code based on variables present

For the fixup code to work correctly, it is to be ensured that $c \cdot f$ does not overflow, as otherwise, the interval we switch over, i.e., $\begin{cases} [0, c \cdot f[& , c > 0 \\]c \cdot f, 0] & , c < 0 \end{cases}$, could potentially be invalid, iff an integer over- or underflow occurs, meaning $\begin{cases} c \cdot f < 0 & , c > 0 \\ c \cdot f > 0 & , c < 0 \end{cases}$.

To avoid these problems altogether, we restricted c to being a compile-time constant, such that for integers defined from t_{\min} to t_{\max} , $c \in [\frac{t_{\min}}{f}, \frac{t_{\max}}{f}]$. Using this restriction, it can be asserted that $c \cdot f \in [t_{\min}, t_{\max}]$ and therefore does not overflow. Algorithm 3.4 details how the mechanics described above are translated into libFIRM. At first, we duplicate the loop body $f - 1$ times and add keepalive edges

to all duplicated nodes, to make sure they do not disappear through implicit premature optimizations. Then we will create the *fixup header*, meaning a block, with the calculation of $n := N - i + (c \mp 1)$. $f - 1$ newly created condition blocks will then use the calculated value by the header. In the i^{th} (counting starts at 0) condition block, it will be checked whether n is in the interval spanned by $c \cdot (f - 1 - i)$ and $c \cdot (f - i)$.⁴ After this, we will wire all duplicated blocks such that they are reachable by the conditions. Further, upon false evaluation of a condition, the following condition is evaluated, except if it is the last condition, in which case the false target is the post loop block. Additionally, except in the case of the first duplicated header, they are attached to the previous blocks as fallthrough. Lastly, the last block of the fixup code now precedes the post loop block, and the false exit for the last condition. An example of the result for creating fixup code for a loop and for $f = 2$, as seen in Figure 3.5, can be seen in Figure 3.10. Further, Figure 3.9 shows the completed unroll process with the added generalized Duff's device, given $f = 2$. Figure 3.8 shows the resulting general structure in pseudo-code.

```

function FOO( $I \in \mathbb{Z}, N \in \mathbb{Z}, c \in \mathbb{Z} \setminus \{0\}, cmp \in \{<, >, \leq, \geq\}$ )
   $i \leftarrow I$ 
  if  $\neg$ SUBTRACTIONWILLLEAVEBOUNDS( $N - c \cdot (f - 1)$ ) then
    while  $i$  ' $cmp$ ' ( $N - c \cdot (f - 1)$ ) do
      DOSOMETHING ▷  $f$  times
       $i \leftarrow i + c$  ▷  $f$  times
    end while
  end if
  switch  $N - i + (c \mp 1)$  do
    case [ $c \cdot (f - 1), c \cdot f$ ] ▷ flip bounds for  $c < 0$ 
      DOSOMETHING
       $i \leftarrow i + c$  ▷ Fall-through
    case [ $c \cdot (f - 2), c \cdot (f - 1)$ ] ▷ flip bounds for  $c < 0$ 
      DOSOMETHING
       $i \leftarrow i + c$  ▷ Fall-through
    ...
    case [ $c \cdot 1, c \cdot 2$ ] ▷ flip bounds for  $c < 0$ 
      DOSOMETHING
       $i \leftarrow i + c$ 
  end function
    
```

Figure 3.8: The general form of a loop starting at I counting in increments of c up to N transformed by loop unrolling with generalized Duff's device fixup

⁴N.B.: c being positive or not determines, which limit is the upper and which is the lower bound.

Algorithm 3.4 Algorithm to build generalized Duff’s device fixup for a given loop in libFIRM

```

function CREATEFIXUPSWITCH(loop : Loop, factor :  $\mathbb{N}_{>1}$ )
  for  $i \in \{0, \dots, (factor - 1)\}$  do
    DUPLICATEBODY(loop)
  end for
  for all node  $\in$  allCopiedNodes do
    if  $\neg$ HASKEEPALIVE(node.link) then
      ADDKEEPALIVE(node)  $\triangleright$  Prevent premature disappearance
    end if
  end for
  relation  $\leftarrow$  header.cmp.relation
  inverseRelation  $\leftarrow$  GETINVERSERELATION(relation)
  duffHeader  $\leftarrow$  NewEmptyBlock
  duffHeader.predecessors  $\leftarrow$  {loop.header}
  val  $\leftarrow$  duffHeader.addNode( $N - i + (c \begin{cases} -, & c > 0 \\ +, & c < 0 \end{cases} 1)$ )
  i  $\leftarrow$  0
  prevLast : Block
  prevCond : Block
  for all body  $\in$  duplicatedLoopBodies do
    firstBlock  $\leftarrow$  GETFIRSTBLOCKINBODY(body)
    condBlock  $\leftarrow$  NEWEMPTYBLOCK
    cond  $\leftarrow$  val ‘relation’ (factor - i)  $\wedge$  val ‘inverseRelation’ (factor - 1 - i)
    condBlock.addNode(cond)
    condBlock.predecessors  $\leftarrow$   $\begin{cases} \{duffHeader\} & , i = 0 \\ \{prevCond.falseExit, prevLast\} & , i \neq 0 \end{cases}$ 
    firstBlock.predecessors  $\leftarrow$  cond.trueExit
    prevLast  $\leftarrow$  GETLASTBLOCKINBODY(body)
    prevCond  $\leftarrow$  cond
    for all phi  $\in$  body.phis do
      phi.predecessors  $\leftarrow$   $\begin{cases} \{phi.link.predecessors\} & , i = 0 \\ \{phi.link.predecessors, prevLast.exitFor(phi)\} & , i \neq 0 \end{cases}$ 
    end for
    i  $\leftarrow$  i + 1
  end for
  for all node  $\in$  allCopiedNodes do
    if  $\neg$ HASKEEPALIVE(node.link) then
      REMOVEKEEPALIVE(node)
    end if
  end for
  postLoopBlock.predecessors  $\leftarrow$  {prevCond.falseExit, prevLast}
  REWIREPHIS  $\triangleright$  Wire just like for duplicated loop body phi’s
end function

```

```

i ← 0
if ¬SUBTRACTIONWILLLEAVEBOUNDS(29 - 6) then
  while i < 23 do
    PRINT>HelloWorld)
    i ← i + 3
    PRINT>HelloWorld)
    i ← i + 3
  end while
end if
switch 31 - i do
  case [3, 6[
    PRINT>HelloWorld)
    i ← i + 3

```

Figure 3.9: An example loop, as seen in Figure 3.5, unrolled by a factor of two, and with generalized Duff’s device fixup

3.3.3 Loop duplication

Another, perhaps simpler, way of creating fixup code is to duplicate the original loop, such that it will run M_{loop} times after the unrolled loop. Just like when using the generalized form of Duff’s device, we unroll the loop using the existing mechanics by a factor of f . Therefore, Equations 3.1 through 3.5 still hold true.

The approach now taken is to copy the original loop, change its initial value to $i_{\text{post loop}}$ and use it as fixup code, as seen in Figure 3.11.

```

i ← ipost loop
while i ‘cmp’ N do
  BODY
  i ← i + c
end while

```

Figure 3.11: The loop to run the body the remaining M_{fixup} times

Proof. To prove that this fixup code preserves semantics, first note that $M_{\text{loop}} \stackrel{3.2}{\leq} M$. Then consider two cases:

1. $M_{\text{loop}} = M$
2. $M_{\text{loop}} < M$

In the first case, $i_{\text{post loop}} \text{ ‘cmp’ } N$ must be false, as otherwise the unrolled loop would have broken semantics, and hence the new loop is never run. Therefore: $M_{\text{fixup}} = 0 \Rightarrow M_{\text{loop}} + M_{\text{fixup}} = M$

In the second case, the new loop runs until the condition is met. As the unrolled loop kept the increment semantics intact, the result is hence $M_{\text{fixup}} = M - M_{\text{loop}}$, which conserves the semantics, as per Equation (3.1). ■

Algorithm 3.5 shows how we create this structure in libFIRM. Firstly, we copy the loop, after which we rewire it, such that the fixup loop points to it, and its old successors point to the fixup loop. Once this is completed, we can unroll the original loop. Figure 3.13 shows the resulting structure of the entire process in pseudo-code. Note that, as mentioned in Chapter 3, this process occurs before we unroll the original loop. Figure 3.12 shows the result for unrolling the loop from Figure 3.5 using loop duplication fixup code and a factor of two.

```

i ← 0
if ¬SUBTRACTIONWILLLEAVEBOUNDS(29 − 6) then
  while i < 23 do
    PRINT>HelloWorld)
    i ← i + 3
    PRINT>HelloWorld)
    i ← i + 3
  end while
end if
while i < 29 do
  PRINT>HelloWorld)
  i ← i + 3
end while

```

Figure 3.12: An example loop, as seen in Figure 3.5, unrolled by a factor of two, and with loop duplication fixup

```

function FOO(I ∈ ℤ, N ∈ ℤ, c ∈ ℤ \ {0}, cmp ∈ {<, >, ≤, ≥})
  i ← I
  if ¬SUBTRACTIONWILLLEAVEBOUNDS(N − c · (f − 1)) then
    while i 'cmp' (N − c · (f − 1)) do
      DOSOMETHING                                     ▷ f times
      i ← i + c                                     ▷ f times
    end while
  end if
  while i 'cmp' N do
    DOSOMETHING
    i ← i + c
  end while
end function

```

Figure 3.13: The general form of a loop starting at I and counting in increments of c up to N transformed by the created loop unrolling with loop fixup

Algorithm 3.5 The algorithm to create a fixup loop in libFIRM

```
function CREATEFIXUPLoop(loop : Loop)
  loop' ← EXACTCOPY(loop)                                ▷ loop is not unrolled yet
  header ← loop.header
  header' ← loop'.header
  for all succ ∈ header.successors do
    if succ.loop ∉ loop then
      succ.predecessors ← succ.predecessors \ {header} ∪ {header'}
    end if
  end for
  for all node ∈ header do
    for all succ ∈ header.successors do
      if succ.loop ∉ loop then
        succ.predecessors ← succ.predecessors \ {node} ∪ {node.link}
      end if
    end for
  end for
  for all phi ∈ headerphis do
    for all pred ∈ phi.predecessors do
      if pred ∉ loop then
        phi.link.predecessors ← phi.link.predecessors \ {pred} ∪ {phi}
      end if
    end for
  end for
  header'.predecessors ← {header.falseExit}
end function
```

3.4 Selecting an unroll-factor

Previously the unroll-factor f seemed like it was chosen somewhat arbitrarily. Further, Section 2.7 describes that there are multiple factors influencing the performance of unrolled loops. Therefore, we devise a selection process.

As a convention, we will henceforth let $size$ be the number of libFIRM-nodes in a given loop. The – admittedly straightforward – algorithm tries to find an unroll-factor $f = 2^n, n \in \mathbb{N}_{>0}$, that minimizes the absolute difference between the unrolled size ($= f \cdot original\ size$), and a pre-determined maximum size. Algorithm 3.6 shows the procedure used to find these values. It is to be noted that the algorithm can also return 0 and 1, which does not fit the definition of the f described. In the case that the algorithm returns one of these two values, we will interpret it as “do not unroll”.

Algorithm 3.6 Algorithm to determine the optimal unroll-factor

```
function CALCULATEFACTOR(loop : Loop, maxSize :  $\mathbb{N}_{>0}$ )  
  loopSize  $\leftarrow$  COUNTNODES(loop)  
  factorPlain  $\leftarrow$  maxSize  $\div$  loopSize  
  factorh  $\leftarrow$  ROUNDTONEXTHIGHERPOWEROFTWO(factorPlain)  
  factorl  $\leftarrow$  factorh  $\div$  2  
  sizeh,l  $\leftarrow$  loopSize  $\cdot$  factorh,l  
  if  $\|size_h - maxSize\| < \|size_l - maxSize\|$  then  
    return factorh  
  else  
    return factorl  
  end if  
end function
```

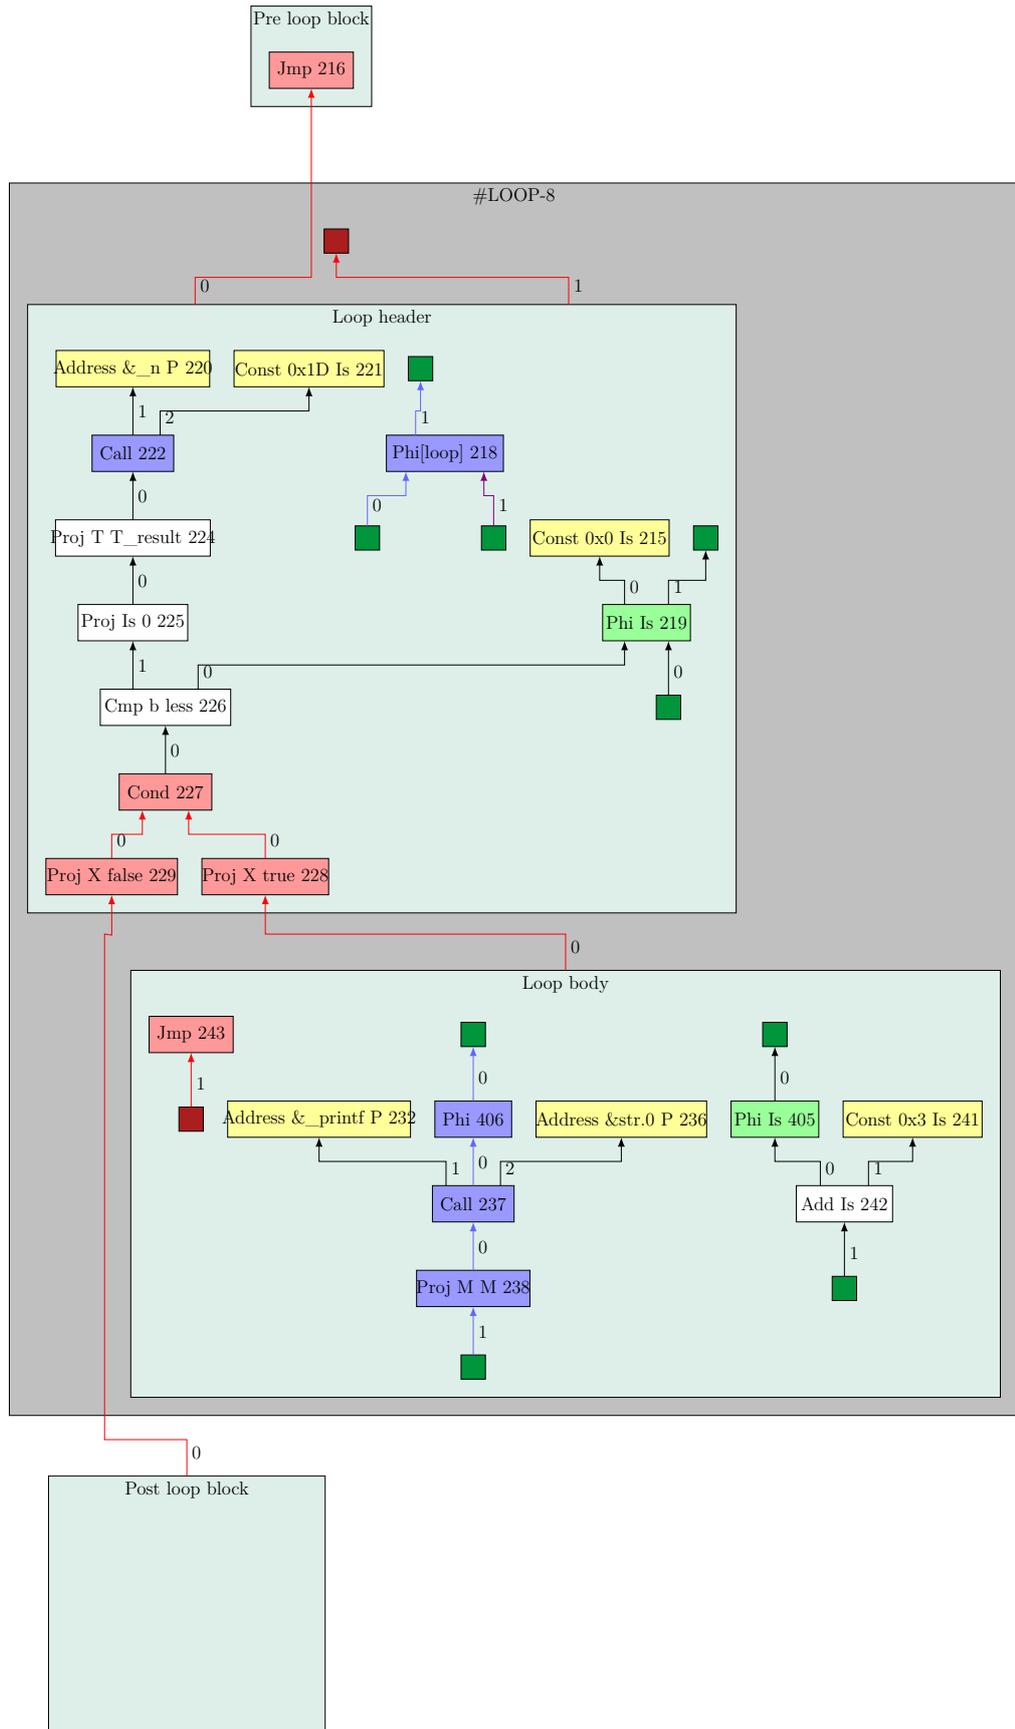


Figure 3.2: libFIRM graph of a loop with an unknown bound

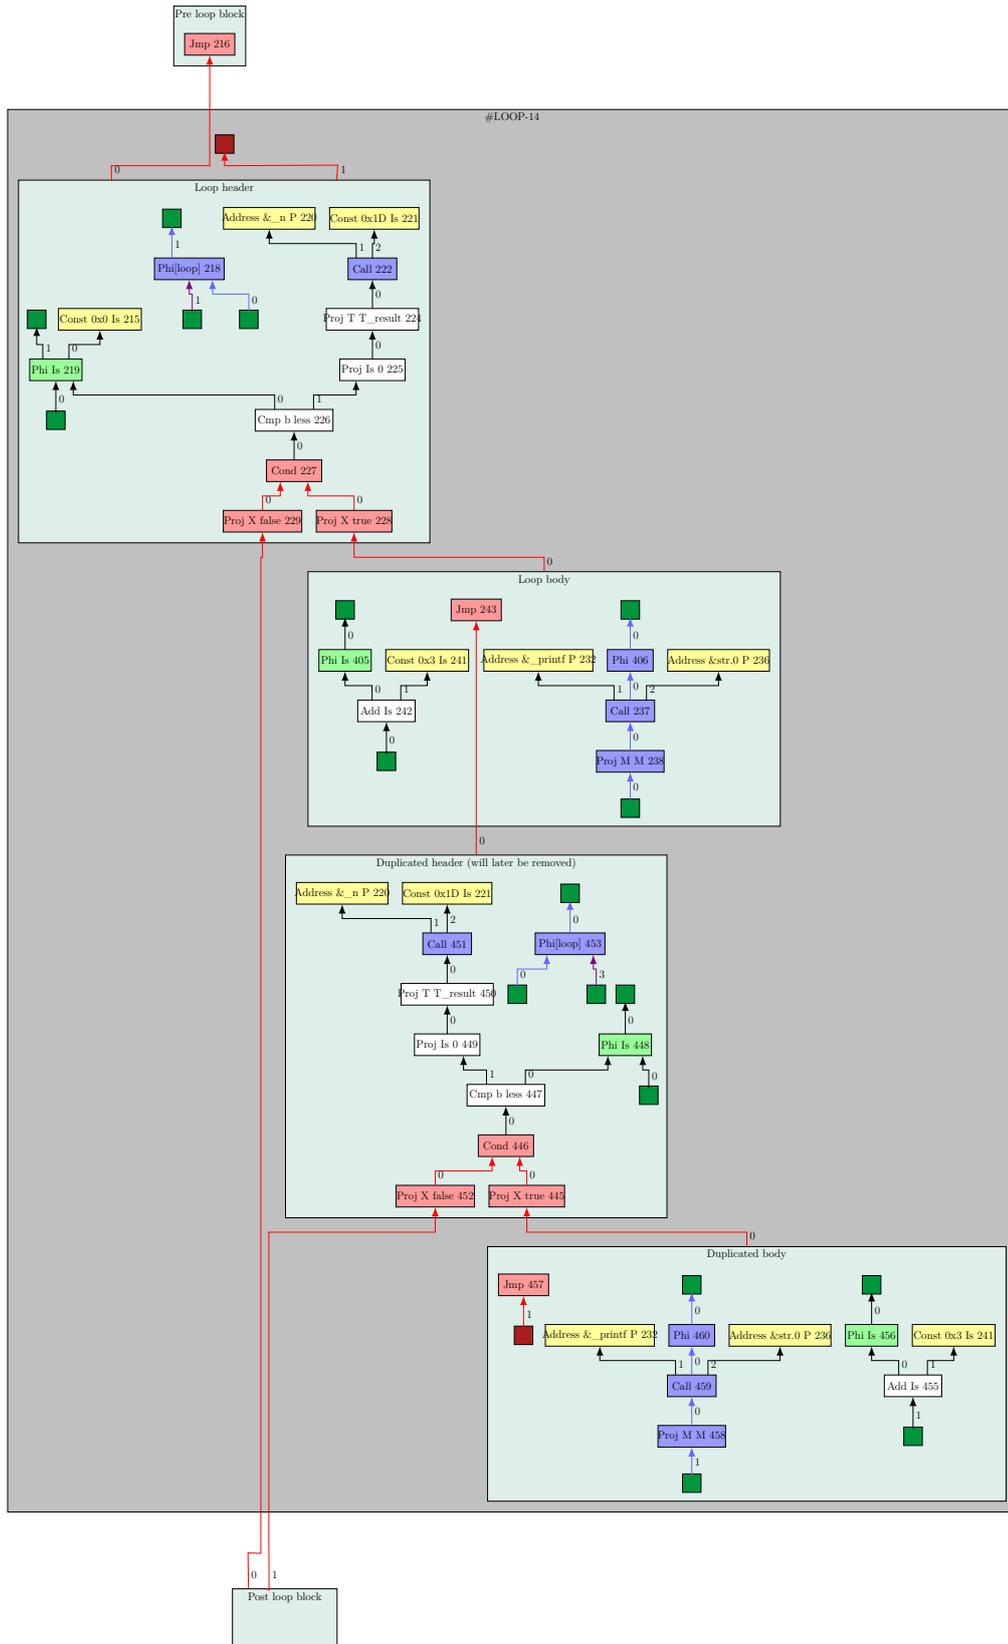


Figure 3.3: libFIRM graph of the loop shown in Figure 3.2 unrolled with a factor of two

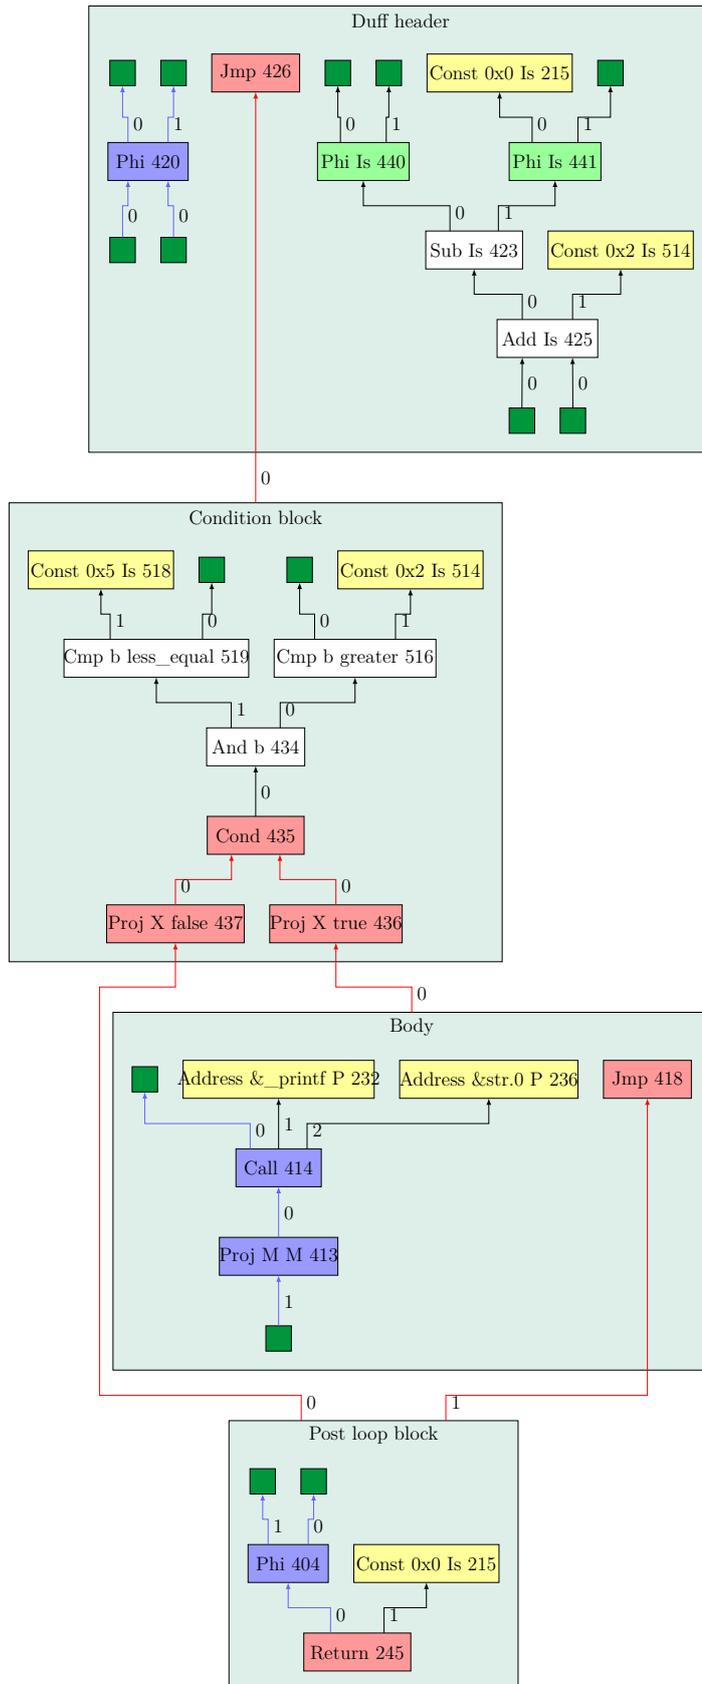


Figure 3.10: Fixup code for the loop from Figure 3.5, given $f = 2$ in libFIRM

4 Evaluation

4.1 Unrollability

One of the primary goals of this thesis was to increase the number of loops that are unrollable with libFIRM. To evaluate to what extent this goal was achieved, we ran the benchmark suite `spec2006`, and logged how many loops we encountered, how many of them were innermost loops, how many could be unrolled using the old method, and how many that were previously not unrollable can now be unrolled¹. Considering it is expected for many loops to have non-constant bounds, such as the length of a container data-structure (e.g., a list or an unbounded array), we predict the new optimization to cause a significant increase in unrollable loops. Figure 4.1 shows a table with the results. We can see, as mentioned in Section 2.7, that prior to the new optimization, 5.87% of the innermost loops could be unrolled. Now we can unroll an additional 7.37% of loops. Contrasted to the baseline of the constant bound unrolled loops this is a 125.65% increase. This means we more than doubled the number of unrollable loops using our approach. Furthermore, we note that more than 70% of loops are innermost loops. Thus, even if unrolling nested loops were advantageous – which is highly doubted – we would not miss out on many loops.

Type	Total count	Relative to loops	Relative to innermost	Relative to constant bound unrollable
Loops	23948	100%	—	—
Innermost	17014	71.05%	100%	—
Constant bound unrollable [2]	998	4.17%	5.87%	100%
Non-constant bound unrollable	1254	5.24%	7.37%	125.65%

Figure 4.1: Comparing total loops, innermost loops and the old unrolling process to the newly implemented process in terms of loops unrolled. The considered loops were all the ones present within the `spec2006` benchmark suite.

¹N.B.: The test was conducted with a max loop size of ∞

4.2 Performance

Even though a high unrollability is a noble goal, most compiler optimizations aim to improve the runtime of the binaries they produce. In order to evaluate the optimization in this regard, `spec2006` is used as a benchmark suite and run on a machine with an Intel Core i7 6700 clocked at 3.4GHz. We run the tests on the Ubuntu 16.04 operating system, with `cparser` [14] as the frontend for `libFIRM`, and the native `x86` backend of `libFIRM` in use. We use the same setup as used in the referenced work [2], such that we can get as comparable results as possible.

To evaluate the performance gain, we run the new optimization in conjunction with the old unrolling (see Section 3.1), given that it is intended as an extension. As a result of there being two approaches for the fixup, as seen in Section 3.3, we will try both of these, to see if one or the other yields better binary runtimes. Furthermore, as described in Section 3.4, the maximum unrolled size determines the scope of the optimization. Therefore, all sizes $l \in \{2^n, n \in [5, 10]\}$ are each tried for both the fixup code strategies. The reason that we chose 32 as a lower bound, is that very small loops are already more than eight nodes in size and hence wouldn't be unrolled with a maximum size that is a smaller power of two. In order to compensate for measurement uncertainties, all benchmarks run ten times, and the average (μ), as well as the standard deviation (σ), will be recorded and discussed. We will compare all results to the reference benchmark run, which itself is a run of `spec2006` without any loop unrolling turned on. These reference results can be seen in Figure 4.3.

In order to evaluate our findings in terms of performance, we should compare them to unrollability broken down by benchmark. Figure 4.2 shows the number of unrollable loops² compared to the total number of loops. Like in Figure 4.1, we assume a maximum size of infinity to collect this data. Seeing this data, we would suspect `bzip2`, `mcf` and to a lesser extent (even though it has the most unrollable loops in absolute terms) `h264ref` to have the most considerable speedup.

Benchmark	Loops	Unrollable loops	Compared to total loops
<code>perlbench</code>	3201	104	3.25%
<code>bzip2</code>	347	231	66.58%
<code>gcc</code>	10207	495	4.85%
<code>mcf</code>	74	42	56.76%
<code>gobmk</code>	3966	387	9.76%
<code>hmmer</code>	1412	239	16.93%
<code>sjeng</code>	410	43	10.49%
<code>libquantum</code>	213	27	12.68%
<code>h264ref</code>	2459	684	27.82%

Figure 4.2: The unrollability broken down by `spec2006`'s benchmarks

²Both constant and non-constant bound unrollable loops are considered together

Benchmark	μ	σ
perlbench	245.18s	0.62s
bzip2	342.68s	0.49s
gcc	181.74s	0.45s
mcf	129.16s	0.18s
gobmk	356.15s	0.29s
hmmer	603.86s	0.07s
sjeng	393.72s	0.40s
libquantum	297.28s	0.47s
h264ref	405.12s	0.27s

Figure 4.3: Results of spec2006 after running it using libfirm without any unrolling

4.2.1 Duff's device fixup

Figures 4.4 through 4.9 show the results we obtained. While for most benchmarks the results hover around the 100% mark, with no significant benefit or drawback, `h264ref` seems to profit from unrolling with maximum sizes 32 and 64, by being close to 4.5% faster. Though on account of the ratios of all the other benchmarks only diverting by three percent or less from the reference runtimes, unrolling does not seem to have a significant effect on performance.

The standard deviations are less than 1% across all tests, due to the highly controlled test environment. Though they do not entirely account for the percentage deltas, which are small, yet measurable. Further, we can expect a small percentage of systemic errors in our measurements due to system process scheduling and similar factors. Runtimes are, independent of the maximum loop size, within [99%, 101%], so we can still consider them to be within the margin of error.

4.2.2 Loop fixup

Figures 4.10 through 4.15 show the results obtained for the unrolling run with the loop fixup code. As was the case in Section 4.2.1, there does not seem to be any noticeable performance gain or loss in any benchmark, except for `h264ref`, which again sped up through unrolling by up to 5%. The other benchmarks were, compared to the reference, within the interval [99%, 103%]. It further becomes evident that there is no correlation between unrollability and performance gain, since, while `h264ref` has one of the highest unrollabilities and gains performance, `bzip2` and `mcf` have higher unrollabilities, yet see no improvement.

Benchmark	Result		Reference		Ratio to reference
	μ	σ	μ	σ	
perlbench	245.60s	0.15s	245.18s	0.62s	100.17%
bzip2	349.34s	0.50s	342.68s	0.49s	101.94%
gcc	180.70s	0.31s	181.74s	0.45s	99.43%
mcf	129.34s	0.40s	129.16s	0.18s	100.14%
gobmk	354.00s	0.42s	356.15s	0.29s	99.40%
hmmer	603.70s	0.15s	603.86s	0.07s	99.97%
sjeng	390.68s	0.23s	393.72s	0.40s	99.23%
libquantum	297.33s	0.61s	297.28s	0.47s	100.02%
h264ref	383.62s	0.73s	405.12s	0.27s	94.69%
Average					99.44%

Figure 4.4: Results of spec2006 after unrolling with maximum size 32 using the generalized Duff's device fixup strategy

Benchmark	Result		Reference		Ratio to reference
	μ	σ	μ	σ	
perlbench	243.65s	0.55s	245.18s	0.62s	99.37%
bzip2	349.74s	1.07s	342.68s	0.49s	102.06%
gcc	181.47s	0.30s	181.74s	0.45s	99.86%
mcf	129.68s	0.60s	129.16s	0.18s	100.40%
gobmk	354.12s	0.22s	356.15s	0.29s	99.43%
hmmer	603.80s	0.19s	603.86s	0.07s	99.99%
sjeng	390.77s	0.43s	393.72s	0.40s	99.25%
libquantum	297.82s	2.06s	297.28s	0.47s	100.18%
h264ref	383.16s	0.49s	405.12s	0.27s	94.58%
Average					99.46%

Figure 4.5: Results of spec2006 after unrolling with maximum size 64 using the generalized Duff's device fixup strategy

Benchmark	Result		Reference		Ratio to reference
	μ	σ	μ	σ	
perlbench	246.36s	0.26s	245.18s	0.62s	100.48%
bzip2	342.24s	0.32s	342.68s	0.49s	99.87%
gcc	181.36s	0.16s	181.74s	0.45s	99.79%
mcf	129.57s	0.47s	129.16s	0.18s	100.32%
gobmk	356.85s	0.31s	356.15s	0.29s	100.19%
hmmer	603.68s	0.22s	603.86s	0.07s	99.97%
sjeng	394.15s	0.12s	393.72s	0.40s	100.11%
libquantum	297.32s	0.40s	297.28s	0.47s	100.01%
h264ref	401.97s	0.33s	405.12s	0.27s	99.22%
Average					100.00%

Figure 4.6: Results of spec2006 after unrolling with maximum size 128 using the generalized Duff's device fixup strategy

Benchmark	Result		Reference		Ratio to reference
	μ	σ	μ	σ	
perlbench	248.95s	0.69s	245.18s	0.62s	101.54%
bzip2	348.61s	0.48s	342.68s	0.49s	101.73%
gcc	181.01s	0.22s	181.74s	0.45s	99.60%
mcf	129.62s	0.43s	129.16s	0.18s	100.36%
gobmk	355.51s	0.22s	356.15s	0.29s	99.82%
hmmer	603.34s	0.17s	603.86s	0.07s	99.91%
sjeng	396.56s	0.37s	393.72s	0.40s	100.72%
libquantum	297.14s	0.34s	297.28s	0.47s	99.95%
h264ref	402.13s	0.47s	405.12s	0.27s	99.26%
Average					100.32%

Figure 4.7: Results of spec2006 after unrolling with maximum size 256 using the generalized Duff's device fixup strategy

Benchmark	Result		Reference		Ratio to reference
	μ	σ	μ	σ	
perlbench	252.26s	0.29s	245.18s	0.62s	102.89%
bzip2	349.73s	0.34s	342.68s	0.49s	102.06%
gcc	180.59s	0.30s	181.74s	0.45s	99.37%
mcf	129.43s	0.45s	129.16s	0.18s	100.21%
gobmk	353.90s	0.25s	356.15s	0.29s	99.37%
hmmer	603.68s	0.19s	603.86s	0.07s	99.97%
sjeng	390.69s	0.22s	393.72s	0.40s	99.23%
libquantum	297.20s	0.47s	297.28s	0.47s	99.97%
h264ref	392.48s	0.71s	405.12s	0.27s	96.88%
Average					99.99%

Figure 4.8: Results of spec2006 after unrolling with maximum size 512 using the generalized Duff's device fixup strategy

Benchmark	Result		Reference		Ratio to reference
	μ	σ	μ	σ	
perlbench	252.07s	0.23s	245.18s	0.62s	102.81%
bzip2	349.04s	0.46s	342.68s	0.49s	101.86%
gcc	181.40s	0.45s	181.74s	0.45s	99.82%
mcf	129.58s	0.49s	129.16s	0.18s	100.33%
gobmk	356.92s	0.23s	356.15s	0.29s	100.22%
hmmer	603.31s	0.14s	603.86s	0.07s	99.91%
sjeng	390.46s	0.29s	393.72s	0.40s	99.17%
libquantum	297.96s	2.23s	297.28s	0.47s	100.23%
h264ref	396.44s	0.26s	405.12s	0.27s	97.86%
Average					100.24%

Figure 4.9: Results of spec2006 after unrolling with maximum size 1024 using the generalized Duff's device fixup strategy

Benchmark	Result		Reference		Ratio to reference
	μ	σ	μ	σ	
perlbench	243.77s	0.62s	245.04s	0.31s	99.48%
bzip2	339.12s	0.54s	342.69s	0.38s	98.96%
gcc	181.65s	0.18s	181.86s	0.26s	99.88%
mcf	128.88s	0.38s	129.60s	0.49s	99.44%
gobmk	357.52s	0.47s	355.54s	0.09s	100.56%
hmmer	603.34s	0.09s	603.98s	0.13s	99.89%
sjeng	393.71s	0.31s	393.89s	0.32s	99.95%
libquantum	297.27s	0.47s	297.30s	0.56s	99.99%
h264ref	402.55s	0.54s	402.67s	0.44s	99.97%
Average					99.79%

Figure 4.10: Results of spec2006 after unrolling with maximum size 32 using the loop fixup strategy

Benchmark	Result		Reference		Ratio to reference
	μ	σ	μ	σ	
perlbench	252.32s	0.63s	245.04s	0.31s	102.97%
bzip2	349.91s	1.04s	342.69s	0.38s	102.11%
gcc	180.74s	0.32s	181.86s	0.26s	99.38%
mcf	129.65s	0.62s	129.60s	0.49s	100.04%
gobmk	353.93s	0.16s	355.54s	0.09s	99.55%
hmmer	603.71s	0.24s	603.98s	0.13s	99.96%
sjeng	390.49s	0.28s	393.89s	0.32s	99.14%
libquantum	298.51s	2.76s	297.30s	0.56s	100.41%
h264ref	392.60s	0.29s	402.67s	0.44s	97.50%
Average					100.12%

Figure 4.11: Results of spec2006 after unrolling with maximum size 64 using the loop fixup strategy

Benchmark	Result		Reference		Ratio to reference
	μ	σ	μ	σ	
perlbench	251.88s	0.21s	245.04s	0.31s	102.79%
bzip2	349.04s	0.45s	342.69s	0.38s	101.85%
gcc	181.25s	0.22s	181.86s	0.26s	99.67%
mcf	129.86s	0.52s	129.60s	0.49s	100.20%
gobmk	356.96s	0.06s	355.54s	0.09s	100.40%
hmmer	603.31s	0.14s	603.98s	0.13s	99.89%
sjeng	390.44s	0.21s	393.89s	0.32s	99.12%
libquantum	297.42s	0.41s	297.30s	0.56s	100.04%
h264ref	396.57s	0.30s	402.67s	0.44s	98.49%
Average					100.27%

Figure 4.12: Results of spec2006 after unrolling with maximum size 128 using the loop fixup strategy

Benchmark	Result		Reference		Ratio to reference
	μ	σ	μ	σ	
perlbench	246.02s	0.70s	245.04s	0.31s	100.40%
bzip2	349.61s	0.38s	342.69s	0.38s	102.02%
gcc	180.83s	0.25s	181.86s	0.26s	99.43%
mcf	129.77s	0.30s	129.60s	0.49s	100.13%
gobmk	353.86s	0.22s	355.54s	0.09s	99.53%
hmmer	603.65s	0.16s	603.98s	0.13s	99.95%
sjeng	390.45s	0.32s	393.89s	0.32s	99.13%
libquantum	297.06s	0.25s	297.30s	0.56s	99.92%
h264ref	383.45s	0.18s	402.67s	0.44s	95.23%
Average					99.53%

Figure 4.13: Results of spec2006 after unrolling with maximum size 256 using the loop fixup strategy

Benchmark	Result		Reference		Ratio to reference
	μ	σ	μ	σ	
perlbench	243.47s	0.36s	245.04s	0.31s	99.36%
bzip2	349.52s	0.59s	342.69s	0.38s	101.99%
gcc	181.78s	0.53s	181.86s	0.26s	99.96%
mcf	129.75s	0.46s	129.60s	0.49s	100.11%
gobmk	354.08s	0.29s	355.54s	0.09s	99.59%
hmmer	603.63s	0.06s	603.98s	0.13s	99.94%
sjeng	390.44s	0.25s	393.89s	0.32s	99.12%
libquantum	297.29s	0.36s	297.30s	0.56s	100.00%
h264ref	383.05s	0.20s	402.67s	0.44s	95.13%
Average					99.47%

Figure 4.14: Results of spec2006 after unrolling with maximum size 512 using the loop fixup strategy

Benchmark	Result		Reference		Ratio to reference
	μ	σ	μ	σ	
perlbench	248.73s	0.53s	245.04s	0.31s	101.51%
bzip2	348.71s	0.46s	342.69s	0.38s	101.76%
gcc	181.01s	0.27s	181.86s	0.26s	99.53%
mcf	129.49s	0.51s	129.60s	0.49s	99.92%
gobmk	355.54s	0.13s	355.54s	0.09s	100.00%
hmmer	603.39s	0.13s	603.98s	0.13s	99.90%
sjeng	396.64s	0.32s	393.89s	0.32s	100.70%
libquantum	297.20s	0.36s	297.30s	0.56s	99.97%
h264ref	402.27s	0.58s	402.67s	0.44s	99.90%
Average					100.35%

Figure 4.15: Results of spec2006 after unrolling with maximum size 1024 using the loop fixup strategy

5 Conclusion

The results, discussed in Section 4.2, fall in line with the results from the ones for static bound unrolling [2]. They, therefore, suggest that there is now empirical evidence that independent of the unrolling method, and the factor chosen, loop unrolling does not yield a significant performance benefit in the current state of libFIRM. Even though more loops were able to be unrolled through the added loop optimization, the increase in unrollability only led to about one in ten loops being unrolled, which certainly is a contributing factor to the underwhelming improvements. Probably some restrictions, such as disallowing `break`-like structures, are too limited and could be dealt with through further development. Other restrictions, such as the conservative alias or call manipulation checks for the bound are unavoidable if the semantics are to be kept and forthright inherent to the task at hand. Inconsiderate of these reasons, even the benchmarks with high unrollability of their loops, did not seem to benefit (with `h264ref` being an exception). Further, it can be concluded that the choice of the fixup code strategy seems to have a negligible impact on performance. Due to the very low standard deviations across all benchmarks, the results also lead to a firm belief the obtained results are trustable and hence provide a solid foundation for empirical conclusions.

Thus, the eminent challenge seems to be the lack of performance gain through unrolling loops. Therefore, it would be a natural starting point to use the unrolled loops and optimize their bodies further. An optimization could be created that takes advantage of the implicitly added semantics as for having a specific modulus, respective to f , for each copied block. Before this potential is used, it likely would be a more lucrative endeavor, to stick to less fancy optimizations that can take advantage of the unrolled loop structures, such as automatically parallelizing non-conflicting operations.

Another factor that might have influenced the results was the method used to determine the unroll-factor. In the future, it could be evaluated, whether the performance would improve through a more sophisticated unroll-factor selection, with a multi-parameter cost function.

Once these changes are in-place, the feasibility of loop unrolling in libFIRM should be reevaluated.

Currently, the efforts of increasing unrollable loops seem to exceed the benefits. Though, if the desire for more unrollability should pick up again, it would seem a good point to look at other loop structures, such as loops with breaks, or a non-counting loop, unlike the ones examined in this thesis.

Bibliography

- [1] A. Aho, M. S. Lam, R. Sethi, and J. D. Ullman, *Compilers: principles, techniques, and tools*, 2nd ed. Pearson, 2014, pp. 632–637.
- [2] E. Aebi, “Ausrollen von schleifen für zwischensprachen in lcssa-form,” Jun. 2018.
- [3] R. Cytron, J. Ferrante, B. K. Rosen, M. N. Wegman, and F. K. Zadeck, “Efficiently computing static single assignment form and the control dependence graph,” *ACM Transactions on Programming Languages and Systems*, vol. 13, no. 4, pp. 451–490, Oct. 1991. [Online]. Available: <http://doi.acm.org/10.1145/115372.115320>
- [4] U. of Illinois at Urbana-Champaign and L. Team, “Lcssa.cpp.” [Online]. Available: <http://web.mit.edu/freebsd/head/contrib/llvm/lib/Transforms/Utils/LC SSA.cpp>
- [5] G. Lindenmaier, “libFIRM – a library for compiler optimization research implementing FIRM ,” Tech. Rep. 2002-5, Sep. 2002. [Online]. Available: http://www.info.uni-karlsruhe.de/papers/Lind_02-firm_tutorial.ps
- [6] A. V. Aho and J. D. Ullman, *Principles of compiler design*. Addison-Wesley, 1979.
- [7] A. Fog. (2018, Apr) Optimizing subroutines in assembly language. [Online]. Available: https://www.agner.org/optimize/optimizing_assembly.pdf
- [8] V. Sarkar, “Optimized unrolling of nested loops,” *International Journal of Parallel Programming*, vol. 29, no. 5, pp. 545–581, Oct 2001. [Online]. Available: <https://doi.org/10.1023/A:1012246031671>
- [9] C. Helmer, “Entwicklung von kriterien zur anwendung von schleifenoptimierungen im kontext ssa-basierter zwischensprachen,” Oct. 2010.
- [10] cparser - graph-based intermediate representation. Commit 2ba5f6e3246a1ae7905857f04b116dc880c2e040. [Online]. Available: <https://pp.ipd.kit.edu/git/libfirm/>
- [11] T. Duff. (1983, Nov) Tom duff on duff’s device. [Online]. Available: <http://www.lysator.liu.se/c/duffs-device.html>

- [12] mult. loop-unroll.c. [Online]. Available: <https://github.com/gcc-mirror/gcc/blob/2e966e2a603e7049a9ea24007f03af858407df93/gcc/loop-unroll.c>
- [13] pmg. (2009, Oct) How do i detect unsigned integer multiply overflow? [Online]. Available: <https://web.archive.org/web/20190821132916/https://stackoverflow.com/questions/199333/how-do-i-detect-unsigned-integer-multiply-overflow/1514309>
- [14] cparser - a c99-frontend. [Online]. Available: <https://pp.ipd.kit.edu/git/cparser/>

Erklärung

Hiermit erkläre ich, Adrian E. Lehmann, dass ich die vorliegende Bachelorarbeit selbstständig verfasst habe und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe, die wörtlich oder inhaltlich übernommenen Stellen als solche kenntlich gemacht und die Satzung des KIT zur Sicherung guter wissenschaftlicher Praxis beachtet habe.

Ort, Datum

Unterschrift