Supporting On-Stack Replacement in Unstructured Languages by Loop Reconstruction and Extraction*

Raphael Mosaner
Johannes Kepler University Linz
Austria
raphael.mosaner@jku.at

David Leopoldseder
Johannes Kepler University Linz
Austria
david.leopoldseder@jku.at

Manuel Rigger
ETH Zürich
Switzerland
manuel.rigger@inf.ethz.ch

Roland Schatz
Oracle Labs
Austria
roland.schatz@oracle.com

Hanspeter Mössenböck
Johannes Kepler University Linz
Austria
hanspeter.moessenboeck@jku.at

Abstract

On-stack replacement (OSR) is a common technique employed by dynamic compilers to reduce program warm-up time. OSR allows switching from interpreted to compiled code during the execution of this code. The main targets are long running loops, which need to be represented explicitly, with dedicated information about condition and body, to be optimized at run time. Bytecode interpreters, however, represent control flow implicitly via unstructured jumps and thus do not exhibit the required high-level loop representation. To enable OSR also for jump-based—often called unstructured—languages, we propose the partial reconstruction of loops in order to explicitly represent them in a bytecode interpreter. Besides an outline of the general idea, we implemented our approach in Sulong, a bytecode interpreter for LLVM bitcode, which allows the execution of C/C++. We conducted an evaluation with a set of C benchmarks, which showed speed-ups in warm-up of up to 9x for certain benchmarks. This facilitates execution of programs with long-running loops in rarely called functions, which would yield significant slowdown without OSR. While shown with a prototype implementation, the overall idea of our approach is generalizable for all bytecode interpreters.

CCS Concepts • **Software and its engineering → Just-in-time compilers**; **Dynamic compilers**; *Virtual machines*.

Keywords On-stack Replacement, Truffle, Sulong

1 Introduction

On-stack replacement [9] is a technique employed by dynamic compilers for reducing program warm-up time [1]. Based on Barrett's work [1], we define program warm-up as the time it takes a dynamic compiler to identify and compile the hot parts of a program to reach a steady state of peak performance. OSR usually works by detecting hot but not yet compiled code and performing a switch from the interpreted to the compiled version of this code while it is being executed. This is most useful for long-running loops, which can take up most of the execution time and should be compiled as soon as possible. However, method-based compilation systems do not directly support OSR, because methods are the most fine-grained compilation units. For instance, a computation intensive loop in a *main*-function would never be compiled, as the *main*-function is only called once and thus never considered hot. To make OSR work in such systems, it would be possible to extract loops into separate functions which can be independently compiled by a method-based compiler even when it lacks support for OSR. While for structured languages, loop bodies could be extracted to separate functions to allow for their quick compilation, unstructured languages lack high-level loop information. To tackle this issue and support OSR for unstructured languages in method-based compilation systems we propose an approach in which we (1) reconstruct loop information from unstructured, basic block-based program representations and (2) explicitly represent loops as separate functions to enable OSR.

We implemented our approach in Sulong [21, 22], an interpreter for LLVM IR, which suffers from long warm-up time [22]. LLVM IR is an unstructured language that does

not explicitly represent loops. Sulong is used in the multilingual GraalVM [7, 18, 25] to implement native function interfaces of dynamic languages such as TruffleRuby [2]. All these language implementations are based on a common language implementation framework, called Truffle [31]. It uses dynamic-compilation based on profiling information gathered during interpretation of abstract syntax trees (ASTs) for efficient execution of the implemented language. Truffle does not directly support OSR. However, it provides language implementers with a framework to enable loop-level OSR which requires extracting loops to form separate compilation units as outlined above. We demonstrate that reconstructing loops and extracting them to separate units gives significant speed-ups. Specifically, our evaluation with a set of well-known benchmarks shows significant reductions by up to 9x in program warm-up given that OSR is applicable. Note that our approach can be used by other Truffle bytecode-based implementations lacking high-level loop information, including GraalSqueak [19], Truffle Java [11] and Truffle CIL. Furthermore, it is applicable in any compiler with a background system which provides means for establishing mappings between extracted control flow.

In summary, this paper contributes the following:

- a novel multi-tier approach for employing OSR for unstructured languages
  - detection of loops from unstructured control flow
  - reconstruction of high-level loops
  - extraction of loops into separate functions to model loop-level OSR using method-based compilation
- a prototype implementation in Sulong,
- extensive experiments suggesting significant improvements in warm-up performance.

2 Background

This section provides the necessary background information to understand our approach for supporting OSR in Truffle-based implementations of unstructured languages. We first give an overview of the Truffle language implementation framework and its OSR mechanism for structured languages. We then discuss bytecode-interpreter-based Truffle languages and the problems they cause for OSR. We implemented a prototype of our approach in Sulong, an interpreter for LLVM IR. Thus, we also give background information about Sulong and LLVM.

2.1 Truffle and Graal

The Truffle language implementation framework [32] provides language implementers with means for creating efficient Abstract Syntax Tree (AST) interpreters. In an AST interpreter, each operation is implemented by a node that computes its value by potentially evaluating its children,

to then return this value to its parent. Truffle uses the dynamic Graal compiler [32] to efficiently compile frequently-executed functions to machine code. When invoked, it recursively inlines all node execution methods of the AST (which is a form of partial evaluation [10]), to then further optimize it.

To achieve optimal performance, language implementers have to use various constructs of the Truffle framework. Graal, when used as a standard Java compiler, is capable of performing OSR on loop-level. Truffle, however, provides an interface that needs to be implemented by guest language nodes to support OSR for structured languages. This RepeatingNode interface assumes that the interpreted language has a concept of high-level loops in order to enable OSR.

Listing 1 shows a typical implementation of a RepeatingNode. The executeRepeating() method executes only a single loop iteration and returns either true if the loop is to be executed again or false if it was the loop's last iteration. The Truffle framework wraps the RepeatingNode with a LoopNode that executes the loop. Its implementation is transparent to the language implementer. Since each loop iteration is performed individually, the LoopNode can trigger compilation after any loop iteration and switch from interpreted to compiled execution. Currently, Truffle uses a constant iteration threshold before OSR-compiling a loop. Note that this technique is generally applicable to compilation systems for adding OSR support.

```
class LLVMRepeatingNode implements RepeatingNode{
  public boolean executeRepeating(Frame frame) {
    if ((boolean)conditionNode.execute(frame)) {
      try{
        bodyNode.execute(frame);
      } catch (BreakException e) {
        return false;
      } catch (ContinueException e) {
        return true;
      }
      return true;
    }
    return false;
  }
}
```

Listing 1. A typical implementation of a RepeatingNode for structured languages. Both conditionNode(loop condition) and bodyNode(loopBody) are child nodes of the RepeatingNode. Execution of the body can be interrupted by break- or continue-statements, which throw an exception.

2.2 Sulong and LLVM IR

Sulong [21, 22] is a bytecode interpreter for LLVM IR built on top of Truffle. LLVM IR is the intermediate representation of the LLVM compilation framework [16], which provides front ends for various languages such as C/C++ and Fortran, allowing Sulong to execute these languages via LLVM IR. Unlike structured languages, LLVM IR has no notion of loops;

instead, it uses basic blocks [29], which consist of sequential instructions that end with a terminating instruction transferring control flow to another basic block. Since these terminating instructions are conditional or unconditional jumps, LLVM IR is considered a unstructured language [5]. Furthermore, it can exhibit irreducible control flow [8], where loops have more than one header (for example by having a jump from outside into the loop). This prevents loops from being representable by a single AST node without further processing them, to handle the irreducibility.

Figure 1c shows the LLVM IR of the main function from Figure 1a with basic block 1 forming the loop. The control-flow graph of basic blocks can be inferred from Figure 1d. Basic block 0 is the function start which unconditionally jumps to basic block 1, which is the loop header. This block is then terminated with a conditional branch instruction that either transfers control back to itself or exits the loop by jumping to basic block 2, which has only the ret instruction for returning from the function. In this simple example, the loop consists of only one block, which holds both, condition and body.

Sulong's execution mode is different from classical AST-based Truffle language implementations, such as TruffleRuby or GraalPython. It uses a BlockDispatchNode that dispatches the control flow between the individual LLVM IR BasicBlockNodes [21], whereas the high-level control flow is unknown to Sulong (see Figure 1d). This makes Sulong's structure similar to a bytecode interpreter (see Listing 2). For each basic block, Sulong builds a conventional AST with a BasicBlockNode as its root. The BasicBlockNode executes all its instructions sequentially. The last instruction of the basic block returns an index that is used to fetch the next basic block to be executed. To return from a function, −1 is returned as a special value. GraalSqueak [19] for example, is implemented in a similar way.

2.3 Partial Evaluation of Bytecode Interpreters

Truffle optimizes functions at run time by partially evaluating their ASTs. Partial evaluation (PE) recursively inlines the execution code of children AST nodes into their parent node until one final compilation unit is derived. However, Sulong represents loops with its bytecode-interpreter-like basic block dispatch node (see Listing 2). In the case of loops, the standard version of PE would infinitely inline successor blocks of already inlined blocks. Thus, the partial evaluation had to be adapted [21] to support bytecode interpreters as described below.

Sulong's basic block dispatch loop is treated in a special way by Graal. As for AST-based Truffle interpreters, Graal starts to partially evaluate the interpreter, starting with the first basic block. The basic blocks and the indices of their possible successor blocks are constant during run time. Thus, Graal can continue to recursively inline all successor blocks. In contrast to AST-based approaches, however, Graal keeps

track of paths that have already been expanded before, which it can determine based on their index in the constant basic block array. If it detects an already partially evaluated block, Graal connects the path from the currently processed block with the previously expanded successor block. Effectively, it reconstructs guest-language-level loops. The merging of already expanded basic blocks also effectively limits loop expansion to the number of original blocks in the LLVM IR program. Even though bytecode might exhibit irreducible control flow, Graal does not support it. Reducible control flow yields easier optimization properties and faster algorithms. Thus, when encountering irreducible control flow, Graal, wraps the irreducible blocks in a dispatch loop during partial evaluation, as it is also done in Sulong.

```
int blockIndex = 0
while(blockIndex!=−1){
    blockIndex = basicBlocks[blockIndex].execute();
}
```

Listing 2. Block Dispatch Loop; compacted for brevity

2.4 Problem of OSR for Bytecode Interpreter Languages

The difference between conventional AST- and bytecode-interpreted languages also manifests itself in how OSR can be applied. In order to switch from an interpreted to a compiled loop, a mapping between the frame state of the interpreter and the compiled code must be established to continue execution in the same state after the transition [9]. The frame state describes all dynamic information of the current program execution including local variables and the program counter. Figure 1a depicts a source program with structured control flow, which is translated into a Truffle AST as shown in Figure 1b. The RepeatingNode wraps one loop iteration and is turned into a CallTarget, which is equivalent to creating a special function for executing the loop body. Thus, when OSR is triggered, the mapping of the program state between interpreted and compiled loop is trivial, as only the frame state at the loop's CallTarget invocation has to be mapped to the state at the beginning of the compiled loop. Hence, the frame state between two loop iterations is used as a parameter for the loop CallTarget, which suffices to continue execution in the compiled loop.

However, in bytecode interpreter languages, like LLVM IR, there is no concept of loops because not all bytecodes are directly derived from structured languages. Figure 1c shows the structure of basic blocks derived after compiling the main function from Figure 1a to LLVM IR. The BlockDispatchNode, which is the AST root node for functions in Sulong, stores the array of basic blocks without any knowledge of the high-level control flow. This is visualized in Figure 1d. Therefore, even with the compiled versions of the two programs (structured C program and LLVM IR code)

```
int main(int argc, char** argv) {
int i = 0;
do {
processRequest();
i++;
} while (i < 1000000);

return 0;
}
```
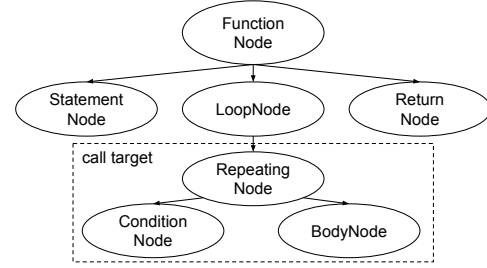
(a) Source level function

```
define i32 @main(i32 %argc, i8** %argv) #0 {
br label %1  (basic block 0)

; <label>:1  (basic block 1)
%i.0 = phi i32 [ 0, %0 ], [ %2, %1 ]
call void @processRequest()
%2 = add nsw i32 %i.0 , 1
%3 = icmp slt i32 %2, 1000000
br i1 %3, label %1, label %4

; <label>:4  (basic block 2)
ret i32 0
}
```
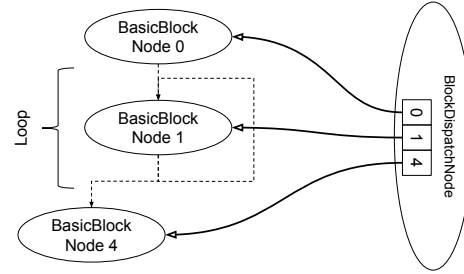
(c) Function in LLVM IR

(b) Truffle AST of program with high-level loop

(d) Bytecode interpreter block dispatch; control flow is depicted as a dashed line

Figure 1. Test program: Source code; LLVM IR, Truffle AST and basic block interpreter.

being identical after partial evaluation, the interpreted loop cannot be mapped to the compiled loop, preventing Truffle's OSR mechanism to be applied for bytecode interpreter languages.

3   OSR for Bytecode-based Truffle Interpreters

In this section, we present our approach to enable OSR for bytecode-based Truffle interpreters. The idea is general enough to be applicable to any bytecode-based language. For this paper we use our prototype implementation in Sulong as an example. To re-use Truffle's built-in mechanism for OSR, we reconstruct loops from the potentially unstructured control flow and create new nodes for representing high-level loops according to the Truffle interface.

Below we list all necessary steps to support OSR in Sulong:

1. The language implementation's parser needs to detect loops in the function before executing it the first time.
2. After the loops are identified, the parser also needs to determine loop relations to handle nested loops separately.
3. Then, the Truffle implementation needs to create `LoopNodes`, and integrate them in the block dispatch loop such that execution of the next `BasicBlockNode` after a `LoopNode` works seamlessly.

4. Finally, during execution, the `LoopNode` needs to communicate the successor of the loop to the enclosing loop or function.

We expect language implementations for other unstructured languages to use the same steps to implement support for OSR. The remainder of this section describes each step in detail.

3.1   Loop Detection

When parsing the basic blocks of the LLVM IR, we initially build a control-flow graph (CFG) by identifying predecessors and successors of all blocks, which is needed for the loop detection algorithm. Our loop-detection algorithm is based on Graal's depth-first algorithm [30] for detecting loops in bytecode is shown in Algorithm 1.

We now discuss Algorithm 1 in detail. It takes basic block b as parameter for the next depth-first step. A block can be marked as *visited*, which means that it was already seen by the algorithm. When a block is marked *active*, the algorithm has started but not finished processing it and the block is still further up in the current recursive traversal.

The algorithm checks whether b is already visited (line 1) and if it is still marked as active (line 2). If both conditions hold, b must have been reached via a backedge and thus a loop with header b is detected. A new loop is created (done by function makeLoop in line 3) with b as header. This loop is denoted *mainLoop* of b. The algorithm always returns the

---

Algorithm 1: FINDLOOPS detects loops in a CFG of basic blocks

Input: The basic block b where to start a depth-first loop detection

Output: A set of loops found in the graph

```
 1 if b.visited then
 2  │   if b.active then
 3  │   │   makeLoop(b)      // new loop with header b
    │   │                    // b.isLoopHeader = true
 4  │   │   return b.loops   // loops which contain b
 5  │   else if b.isLoopHeader then
 6  │   │   return setSubtraction(b.loops, b.mainLoop)
    │   │     // return all loops except the loop
    │   │       // which is associated to the header
 7  │   else
 8  │   │   return b.loops   // loops which contain b

 9
10 b.visited ← true, b.active ← true
11
12 foreach block s ∈ b.successors do
13  │   b.loops.add(findLoops(s))      // add loops
    │                   // propagated from successors
14 foreach loop l ∈ b.loops do
15  │   l.add(b)       // add b to containing loops
16
17 b.active ← false
18
19 if b.isLoopHeader then
20  │   return setSubtraction(b.loops, b.mainLoop)
    │     // return all loops except the loop
    │         // which is associated to the header
21 else
22  │   return b.loops      // loops which contain b
```

---

set of loops which contain b's predecessor in the traversal which we denote as back-propagation (lines 4, 6, 8, 20 and 22). There are three different cases:

- b is visited and active and thus is detected to be a loop header (line 4): Then, all of b's loops contain b's predecessor in the recursive traversal (source of the backedge). This includes the newly created loop.
- b is visited but not active anymore and is marked as loop header (line 5): This means that b is not reached via a backedge. Therefore all of b's loops except its *mainLoop* contain b's predecessor in this traversal. In the algorithm we remove the *mainLoop* from b's loops using a set subtraction (lines 6 and 20).

- b is visited but not active anymore and no loop header (line 7): Therefore, all of b's loops contain b's predecessor in this traversal (lines 8 and 22).

In case a block b is not yet visited, it is marked both visited and active (line 10) and the algorithm is called recursively for each successor (lines 12 and 13). The back-propagated loops are added to b's loops (line 13) and vice versa (lines 14 and 15). After the recursive traversal of all successors is finished, b is set to inactive (line 17) and depending on its loop header flag the set of loops which contain b's predecessor is returned (line 20 or 22). Eventually, the first block of the function is processed and at that point the set of loops has to be empty to ensure that the context of all loops has been closed correctly. If this set is not empty, there is a path into some loop which does not pass the initially identified header yielding multiple loop entries and thus, irreducibility.

We believe that irreducible control flow in unstructured languages is rare; in fact, we did not encounter it in the benchmarks used in the evaluation (see Section 4.1). Thus, we omit the handling of irreducible control flow and bail out (i.e., we execute the function without trying to reconstruct its loops) for functions that contain it.

Figure 2 shows two examples of loop detection applications, one with a simple, reducible loop and the other with a slightly changed, yet irreducible control flow (having two loop entry points). In the reducible example, in Step 3 the backedge points to an active block which is marked as loop header and triggers back-propagation of the newly created loop ID. After Step 4 the left most block is processed and then in Step 5 backtracking arrives at the start block. Note, that some intermediate steps in the irreducible example are omitted as up to Step 3 everything is identical to the example with reducible control flow. In Step 3, the second entry to the loop (bypassing the header) is processed and the loop ID is thus propagated upwards to the start block. The irreducible loop is detected, because the function start block, which can never be part of a loop in LLVM IR, got a loop ID assigned.

3.2 Identifying Loop Relations

We want to support OSR not only for the outermost loops, but also for all inner ones. To achieve this, we need to model a loop nesting based on which we build a loop hierarchy in our node structure later on. Since the loop detection is unaware of the nesting level, we have to manually determine a contains-relationship between the loops. This is done by recursively checking if a loop contains the header of another loop, which is then added to its set of inner loops. We use the depth-first approach to additionally have the loops sorted, that is, an outer loop is processed after its contained inner loops. Hence, when creating LoopNodes we can use the established order to have all inner loops resolved by the time a dependent outer loop is created.
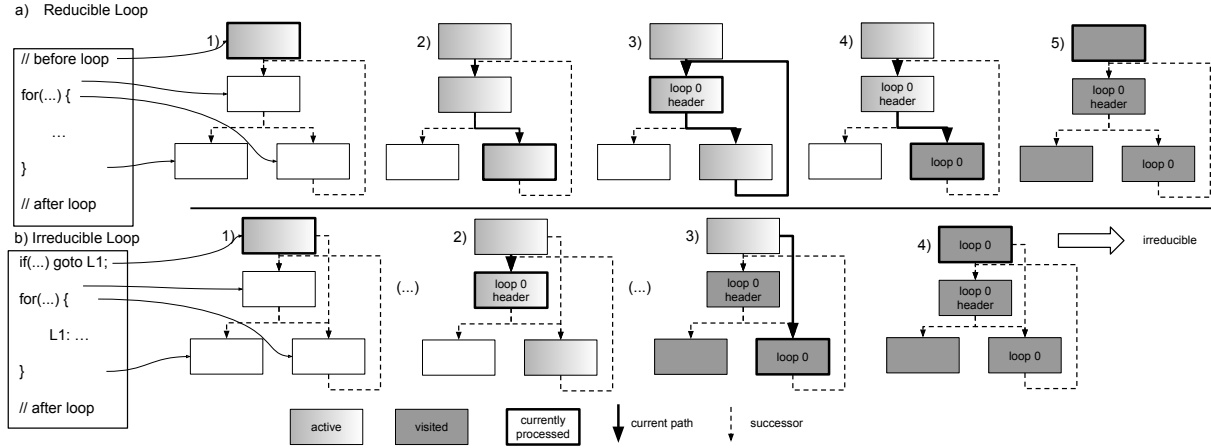
Figure 2. Visualization of two Loop Detection applications. Subfigure a) shows the detection of a reducible loop while subfigure b) shows an irreducible loop being detected as such, denoted by a loop ID set for the first block

Currently, we do not identify loop relations for irreducible loops. While they could be partially supported by adapting our approach (see Section 3.6), it is still necessary to determine which loop is an inner loop and which is an outer loop. If this distinction is impossible, we are not able to establish the node hierarchy because inner loops have to be processed before the outer loop is created. Thus, we check for a bi-directional contains-relationships and bail out for such constructs.

3.3 Node Creation and Integration

To use Truffle's OSR mechanism, we create nodes for identified loops and embed their execution within the dispatch loops of basic blocks. This includes implementing the Truffle `LoopNode` interface for modeling loops, adopting the block dispatch to also support high-level loop structures and providing a way for executing the loop.

The function parser creates `LoopNodes` together with all other Truffle nodes when parsing a method. After the control flow analysis, it wraps the `BasicBlockNodes` forming a loop into a `LoopDispatchNode`, which models the loop body and is responsible for executing the loop. The `LoopDispatchNode` is then used to create `RepeatingNodes` and `LoopNodes` conforming to the Truffle interface. Each `LoopNode` replaces the `BasicBlockNode` which denotes the header of the very same loop in the array of `BasicBlockNodes`. Thus, already created `LoopNodes` can be found in the body of outer loops. To establish this node nesting, the function parser has to process loops iteratively from inside out, so that outer loops are created after their inner loops. However, the parser automatically ensures this by processing the loops according to the established order as described in Section 3.2. Finally, the function node is created by using the array of `BasicBlocks` and `LoopNodes` which

are not already nested in another loop. This leads to a semi-hierarchical node structure, in comparison to the initially flat `BasicBlockNode` array. Figure 3 depicts this internal structure based on the function in Figure 1a.
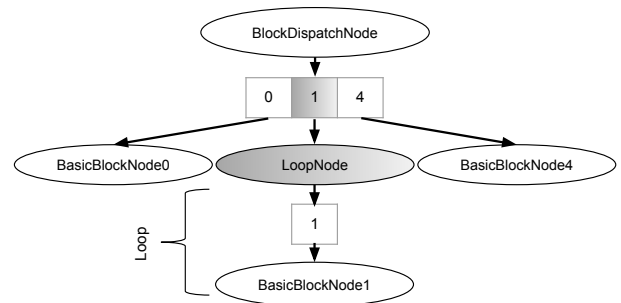

Figure 3. Node hierarchy; compacted for brevity (no `RepeatingNode` and `LoopDispatchNode` modeled)

3.4 Loop Execution

There are two ways to execute implicit unstructured LLVM IR loops as explicit loops in Sulong. Either by a full reconstruction the loop structure (body and condition nodes are not distinguished during loop detection), or by wrapping a block dispatch logic for the loop into the `LoopNode`. The latter approach can be justified by taking a look at it from the compiler's perspective. In that sense, an OSR loop is a special call target which is optimized after several calls (i.e., iterations). Therefore, the new `LoopDispatchNode` is closely related to a `BlockDispatchNode` for functions, yet with some differences. First, less terminating instructions have to be handled, because a `return` block, for example, can never be part of a loop as there is no path back to the loop header. Second, no function arguments have to be handled. Those two points make a `LoopDispatchNode` simpler than

its counterpart for functions. However, a `LoopDispatchNode` has to return after each loop iteration or if the loop is exited. This check for ending an iteration is more costly than for ending a function. Rather than dispatching to the next block until no successor is available (successor = -1; end of function), the loop dispatch has to both check if the successor is the loop header which indicates that the iteration has ended or if the loop is exited by dispatching to one of its successors. However, Graal optimizes this additional checks away.

While a simple loop condition is mapped to only one LLVM block and can thus be used directly as a condition for executing one loop iteration, more complex boolean conditions also consist of multiple blocks. Therefore, a wrapper node would be needed, including another dispatch solely for loop conditions. In order to maintain simplicity and due to the fact that it should be transparent whether a loop is exited because the condition failed or because another exit path was taken (e.g. `return`, `break`, `goto`, ...), the condition is pulled into the loop body. In addition to having solved the problem of handling multi-condition loops, the distinction between condition and body nodes is made superfluous.

3.5 Successor Determination

The final task is to continue function execution at the correct successor block after a loop is executed. However, in contrast to structured control flow, our reconstructed loops can have arbitrary successors reached by jumps from within the loop. Therefore, when entering a loop, it is not known in advance where to continue the execution after the `LoopNode` is evaluated. While the `LoopDispatchNode` could return the successor at run time, the enclosing `RepeatingNode` has its predefined interface (see Listing 1) hindering passing the successor to the caller as return value. However, we can use the run time stack for storing the successor value in a predefined, constant frame slot `LoopDispatchNode`. While this process might seem costly, performance of tight loops is only affected in interpreted mode, due to frame virtualization in compiled programs.

In order to generate code that can be easily optimized by Graal, the block dispatch loop (see Listing 2) and all block successors have to be constant (see Section 2.3). Thus, when reading the loop successor from a run-time-written frame slot, the interpreter loop cannot be unrolled. We solve this issue by storing the set of constant successors in each `LoopNode` and use a lookup loop for finding the constant successor with the same value as in the frame slot. Then, the constant value is used to determine the successor block. As the compiler can determine the set of possible successors for each loop, it can unroll the dispatch loop accordingly. This approach is visualized in Figure 4, where the value in the frame slot is shown to be used as a lookup into the set of constant successors.
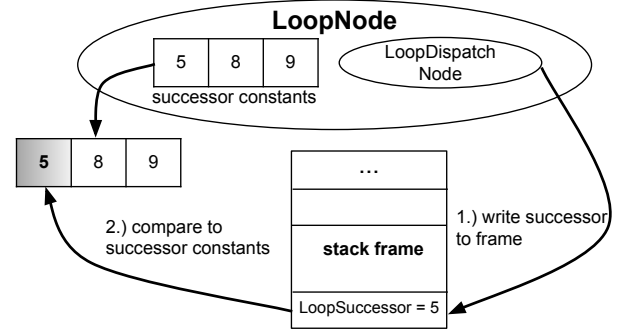

Figure 4. Successor Determination. The successor discovered at run time is written to a stack frame and used as a lookup into the set of constant successors.

3.6 Irreducible Control Flow

Graal can not handle irreducible control flow (ICF), but rather uses a dispatch loop for such patterns. A similar idea is used for our approach implemented in Truffle based languages, which is a way to have irreducible control flow reduced to one AST node. However, support for ICF could be added as outlined below. Irreducible loops are characterized by having multiple entries. Thus, a loop cannot be collapsed to one black box block, as the entry point might differ from execution to execution. There are approaches for transforming irreducible into reducible control flow, which mostly center around code duplication [8].

We initially experimented with the Tarjan algorithm [27] for detecting strongly connected components in order to find such irreducible loops. However, due to additional complexity for resolving nested loops, code growth for duplicated nodes and the fact that ICF was never encountered in our benchmarks this approach was discarded in favor of bailing out in case of ICF.

4 Evaluation

We evaluated our approach which we prototyped in Sulong with a set of C benchmarks to investigate improvements in warm-up and impacts on peak performance.

Hypothesis We hypothesize that enabling OSR reduces the program warm-up significantly while not influencing the peak performance.

4.1 Setup

Benchmarks We primarily selected benchmarks from the *Computer Language Benchmarks Game (CLBG)*[1]. They are micro-benchmarks and do not reflect the behavior of real-world applications. However, their small size and their structure—benchmarks like fannkuch have long-running and computation-intensive loops—makes them suitable to study
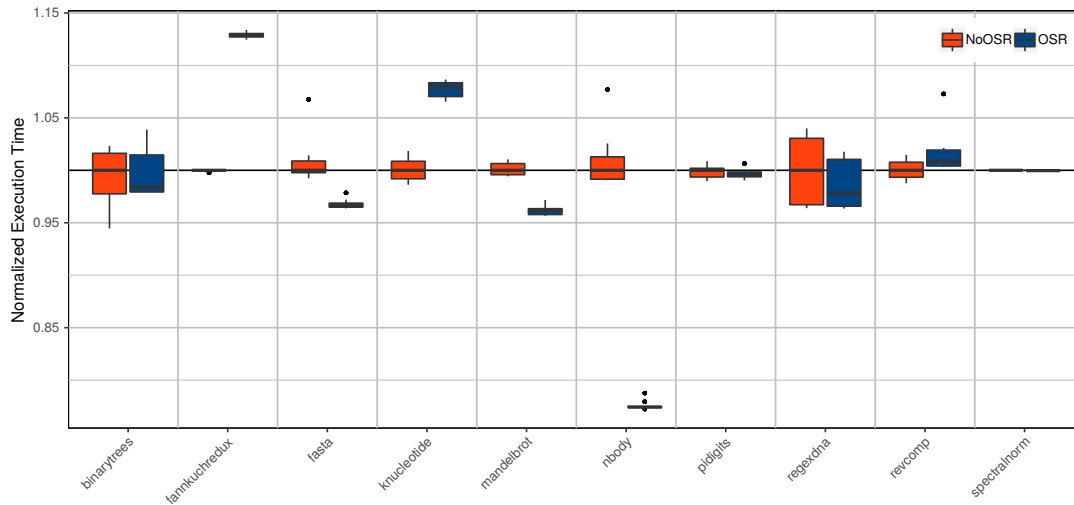
---

[1]https://benchmarksgame-team.pages.debian.net/benchmarksgame/

Figure 5. Normalized execution times after reaching peak performance of Sulong with and without OSR; lower is better.
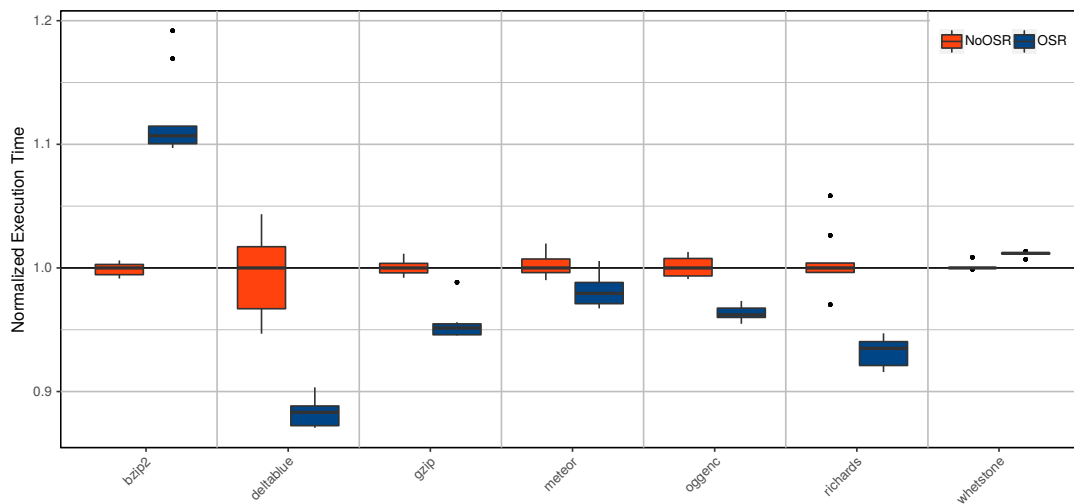

Figure 6. Normalized execution times after reaching peak performance of Sulong with and without OSR; lower is better.

the benefits of OSR. The problem sizes were chosen to allow for multiple benchmark runs. In addition, we used other popular benchmarks like *whetstone*[2], *deltablue*[3] and *richards*[4]. Furthermore, we evaluated our approach with three larger applications, *bzip2*, *gzip* and *oggenc*, which are part of the *Large single compilation-unit C programs*[5], to better investigate the impact on real world programs.

Harness   We used a configurable harness for evaluating the warm-up behavior and peak performance of the benchmarks and conducted ten out-of-process benchmark executions with at least 50 in-process iterations. We analyzed warm-up behavior by analyzing the first ten in-process iterations

of the respective benchmark. We verified that within these ten iterations, execution of the benchmarks reached peak performance. For peak performance, we evaluated only the last ten in-process iterations of each benchmark.

Environment   Our benchmarking machines are equipped with Intel Xeon CPU E5-2699 processors with 72 cores at 2.30GHz each along with 256GB of RAM. For compiling the benchmarks to LLVM IR we used the LLVM front end Clang 3.8 at optimization level O3. Sulong is executed on top of GraalVM, including the default Truffle OSR threshold of 100 000 iterations. We compare Sulong with the implemented OSR approach against its version without any OSR related changes.
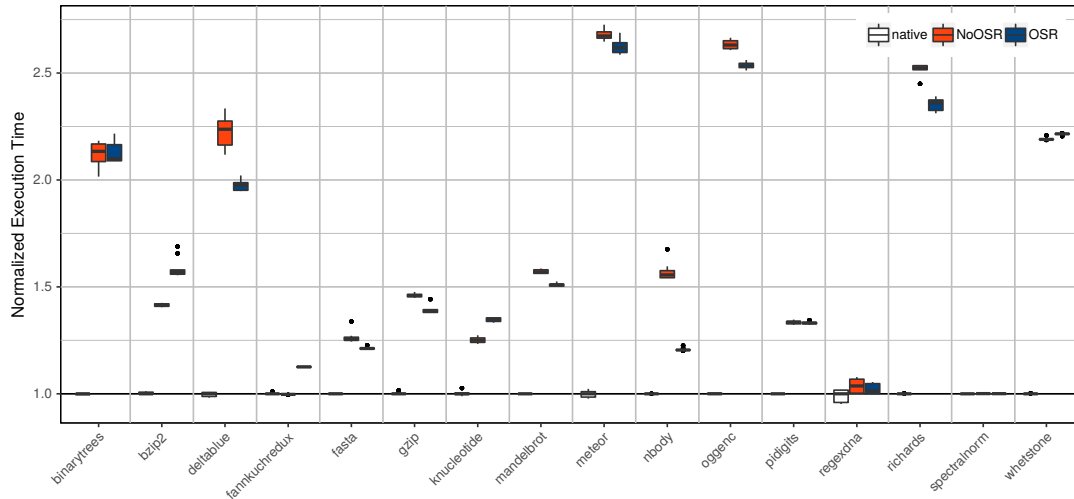
---

[2]www.netlib.org/benchmark/whetstone.c

[3]https://constraints.cs.washington.edu/deltablue/

[4]https://www.cl.cam.ac.uk/~mr10/Bench.html

[5]http://people.csail.mit.edu/smcc/projects/single-file-programs/

Figure 7. Normalized execution times after reaching peak performance of Sulong without OSR-related changes compared to native benchmark execution. Lower is better.

4.2 Results

Native Execution In order to show that Sulong is a mature execution system, we compared its performance against state of the art compiled native executables. Figure 7 gives a peak performance comparison of the benchmarks executed natively and using Sulong with and without OSR related changes. Warm-up is of course not present in the native executions and thus omitted to be shown. Although no contribution proposed in this paper is visible in these benchmark runs, we want to provide a context for the peak performance deviations imposed by our approach. Most benchmarks are, when executed on Sulong, slower by factors up to 2.8x with *revcomp* being about 6x slower. Thus, *revcomp* is not visualized to keep the plot readable. Table 1 contains a condensed depiction of the performance data including the median and the standard error for each benchmark.

Peak-performance Impact While OSR should conceptually not impact a program's peak performance, our experiments disproved this initial hypothesis. Figure 5 and Figure 6 show deviations with up to 13% overhead (median) for *fannkuch* down to 23% speed-up (median) for *nbody*. Only for three benchmarks (*fannkuch*, *knucleotide* and *bzip2*) the peak performance is significantly worse (median at least 5% worse). Three benchmarks (*deltablue*, *nbody* and *richards*) show significantly higher peak performance (median at least 5% better). The majority of our benchmarks, that are 11 out of 17, differ only slightly from the non-OSR version of Sulong, however with a tendency to increased peak performance. We found that due to the changes in the Truffle node structure, optimizations are applied differently by the Graal compiler, resulting in the observed deviations. It is difficult to attribute the peak performance changes to individual optimizations,

Table 1. Overview of peak performance for benchmarks normalized to their native execution and the standard error (SE). Lower is better.

	native		no OSR		OSR	
	median	SE	median	SE	median	SE
binarytrees	1.0	2.449e-04	2.134	4.914e-03	2.099	4.832e-03
bzip2	1.0	4.645e-04	1.417	6.567e-04	1.568	4.435e-03
deltablue	1.0	9.799e-04	2.237	6.649e-03	1.976	2.379e-03
fannkuchredux	1.0	3.272e-04	0.997	8.543e-05	1.126	2.459e-04
fasta	1.0	1.354e-04	1.254	2.618e-03	1.212	5.316e-04
gzip	1.0	5.093e-04	1.460	9.248e-04	1.389	1.745e-03
knucleotide	1.0	9.804e-04	1.251	1.340e-03	1.350	9.148e-04
mandelbrot	1.0	3.973e-05	1.570	8.501e-04	1.510	6.628e-04
meteor	1.0	1.512e-03	2.673	2.196e-03	2.618	3.150e-03
nbody	1.0	1.187e-05	1.556	3.897e-03	1.205	6.787e-04
oggenc	1.0	2.314e-04	2.631	1.943e-03	2.531	1.446e-03
pidigits	1.0	8.932e-04	1.336	7.600e-04	1.331	5.900e-04
regexdna	1.0	2.610e-03	1.037	3.110e-03	1.014	2.335e-03
revcomp	1.0	6.581e-04	6.090	4.951e-03	6.143	1.219e-02
richards	1.0	7.472e-05	2.525	5.535e-03	2.360	2.816e-03
spectralnorm	1.0	1.559e-05	1.001	6.512e-06	1.001	9.955e-06
whetstone	1.0	8.373e-05	2.190	5.841e-04	2.216	3.595e-04

as multiple overlapping compilation paradigms produce the measured results.

Program Warm-up We identified two categories of behaviors in the investigated benchmarks regarding program warm-up, which can be seen in Figure 8 and Figure 9. Firstly, for applications with long-running loops, which trigger OSR, speed-ups by factors of 1.5x (*knucleotide*) up to 9x (*whetstone*) are encountered for the first one to four iterations. This can be seen for benchmarks *fannkuch*, *fasta*, *knucleotide*, *mandelbrot*, *nbody* and *revcomp (reverse-complement)* in the CLBG suite in Figure 8 and also for *gzip*, *meteor* and *whetstone* in Figure 9. The warm-up of *bzip2* shown in Figure 9 behaves differently as the second iteration shows a peak in execution time, which is linked to a de-optimization issue tied to this
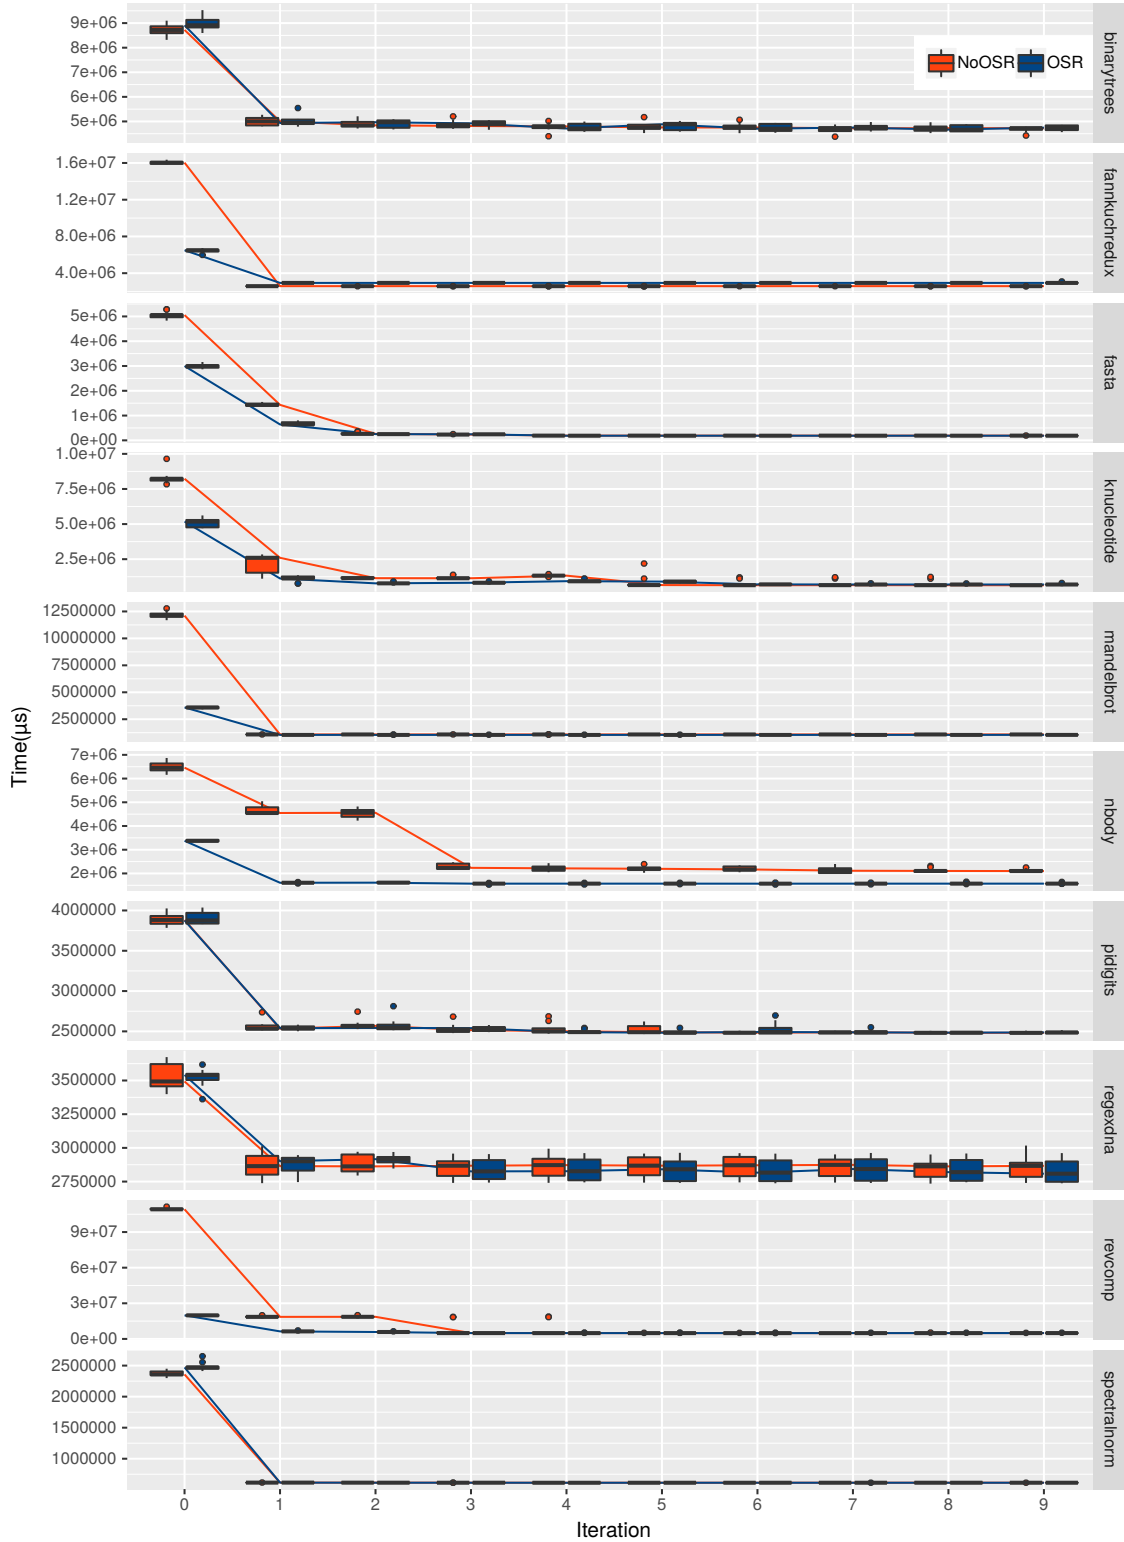
Figure 8. Warm-up of Sulong with and without OSR; the x-axis shows the sequence number in the series of consecutive in-process benchmark executions, while the y-axis shows the execution time of a benchmark execution. Lower is better.

Figure 9. Warm-up of Sulong with and without OSR; the x-axis shows the sequence number in the series of consecutive in-process benchmark executions, while the y-axis shows the execution time of a benchmark execution. Lower is better.
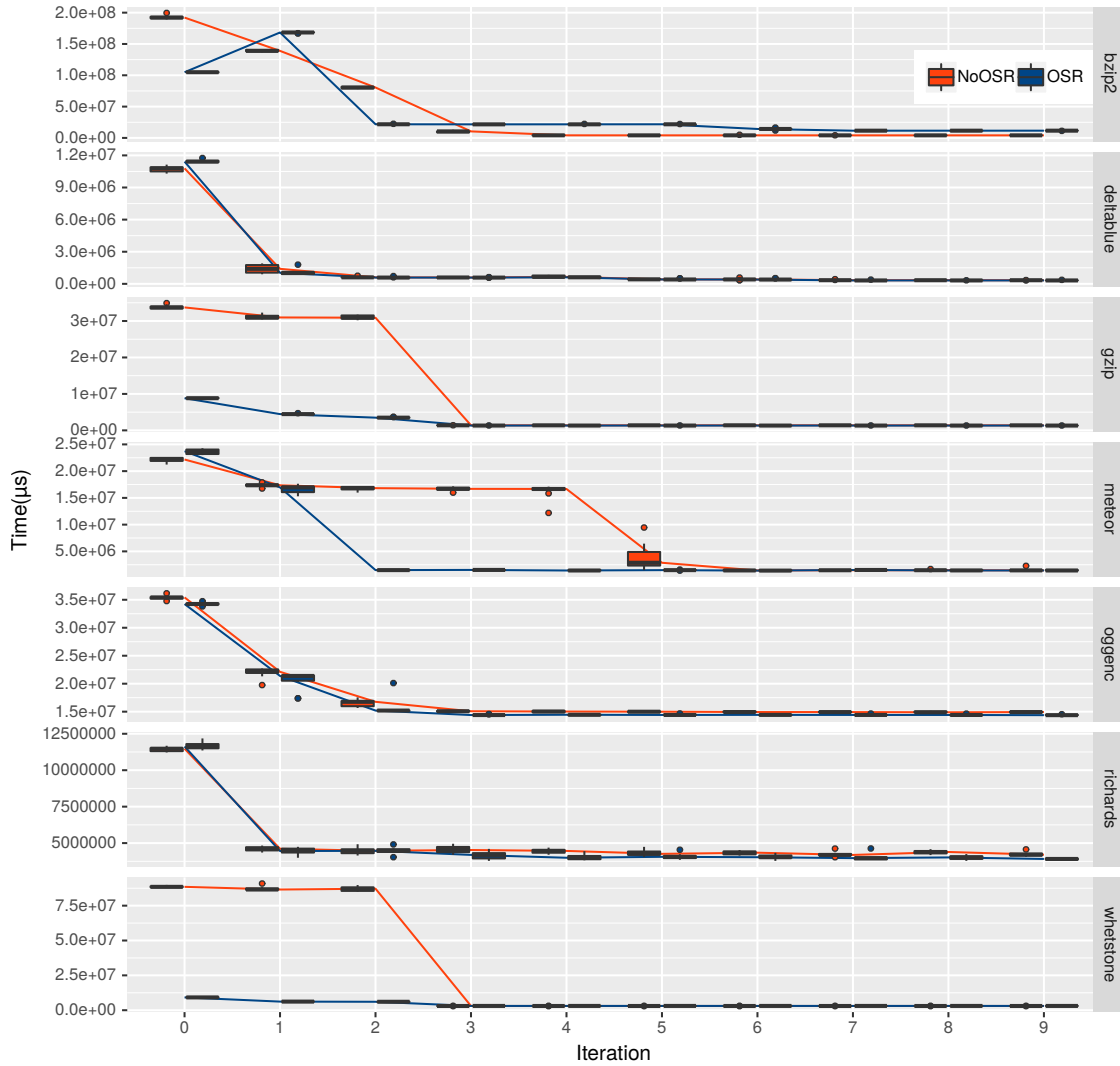
benchmark. Secondly, there are benchmarks where the program warm-up is not affected by OSR at all. This is either the case in the absence of loops or if loops do not reach the iteration threshold for OSR compilation. Minor deviations in warm-up result from the changes in the node structure.

4.3 Discussion

The results suggest that our OSR approach can significantly reduce warm-up time. The actual improvement is highly dependent on the problem size of the benchmark, because for larger problem sizes, loops are often running longer in interpreted mode. For example, in *fannkuch*, the problem size is the length of a permutation array, which determines the number of loop iterations needed to cover all possible permutations. With a problem size of 11, *fannkuch* did not finish within a day without OSR, but in less than 30 minutes with our approach enabled. For evaluation feasibility

fannkuch's problem size was reduced to 9. The evaluation also demonstrated that peak performance is affected by our approach. Due to changes in the Truffle node structure, a different program is produced which explains the deviations. Some benchmarks show peak performance regressions, others speed up, however with a tendency to the latter.

5 Related Work

To the best of our knowledge, we have presented the first approach for providing OSR in a Truffle implementation for an unstructured language. Below, we consider the wider context of related work.

OSR On-stack replacement was first researched by Hölzle and Ungar [14] as a strategy to switch between different versions of an executed method in the context of (re-)compilation and dynamic de-optimization [13] to support

debugging of optimized code. OSR is supported by most of the popular method-based dynamic compilers. For example, the HotSpotVM [15], which is one of the most popular VMs for Java bytecode, supports OSR by incrementing counters at loop backedges that trigger compilation when a counter overflows. When compilation finishes, the next executed backedge will transition to the compiled version of the loop. Values in the interpreter frame are mapped to an OSR buffer, from which the compiled code extracts the values. To support compilers in applying OSR for unstructured languages we proposed reconstructing and extracting loops. D'Elia and Demetrescu [3] conducted a case study on how to implement function-level on-stack replacement for LLVM at arbitrary points, by using glue code to facilitate a smooth transition between the two versions of a function. In their recent work they focused on a more abstract view on OSR in the context of code transformation [4]. In [4] a more general way to apply OSR is proposed by enabling transitions at any code location.

Truffle Bytecode vs. AST Interpreters Niephaus et al. [19] compared an AST-based with a bytecode-interpreter-based Truffle implementation for Squeak/Smalltalk. For the AST-based approach, they had to decompile the generated bytecode to construct high-level Truffle AST nodes. They also report significant performance gains when using Truffle's `LoopNode` interface, which required additions in the decompilation process, to successfully reconstruct loop condition and body. They observed a warm-up period for the AST-based interpreter, but not for the bytecode interpreter approach. However, they remarked that the results might not be generalizable, since they evaluated their implementation with only two micro-benchmarks. We speculate that the bytecode-based version of Squeak/Smalltalk, would also benefit from enabling OSR like in Sulong.

Reconstruction of loops Loop reconstruction is a well researched topic, which dates back into the 1970s with Tarjan [27, 28] formulating his interval analysis algorithm capable of identifying loops in reducible control flow graphs. The algorithm creates a depth-first tree of the CFG and identifies loops in a bottom up traversal from inside out, by collapsing inner loops into single vertices [20]. This depth-first nature is found in many loop reconstruction algorithms being developed over the years—including the one used in this paper. Havlak refined Tarjan's algorithm to work with irreducible control flow too [12], however, in quadratic time [20]. But Havlak's algorithm is still used in more recent work [23]. Ramalingam then extended Havlak's algorithm to identify both reducible and irreducible loops in almost linear time [20]. Other algorithms build on the work of Tarjan and add support for irreducible control flow [24, 26], which slightly differ in the returned irreducible loops [20]. For our work, as irreducible control flow is not supported, the simple depth-first

reconstruction of reducible loops in combination with the detection of irreducible loops is sufficient.

Reconstruction of high-level control flow Leopoldseder et al. [17] described an approach to reconstruct high-level language constructs from compiler IR. They compiled Java bytecode to JavaScript by first translating it to Graal IR [6] and then reconstructing high-level constructs. They faced similar problems like we did, for example, the problem of handling different successor blocks of loop exits. However, their approach merged loop exits whereas we use the dynamically identified loop successor as lookup for the constant successors, known at compile time.

Zakai [33] introduced the *Emscripten* compiler, which translates LLVM IR to JavaScript. In its architecture, a module for reconstructing high-level JavaScript loops from LLVM IR is presented, called *The Relooper*. They also point out that due to extensive use of `goto` statements no meaningful high-level structure might be re-established.

6 Conclusion and Future Work

OSR for loops in unstructured languages is problematic due to the lack of high-level representations on which optimizations can be performed. In this paper, we have presented an approach for enabling OSR in Truffle-based interpreters for unstructured languages by partially reconstructing high-level loops from basic blocks. Unlike traditional OSR, we wrap loops into AST nodes to enable Truffle to extract the *LoopNodes* into separate *CallTargets*, which are function equivalents and can be OSR-compiled after each loop iteration. We implemented this approach in Sulong, and demonstrated that it can significantly reduce warm-up time. Our approach is applicable to any other Truffle based bytecode interpreter, but also other languages can implement similar approaches building on our multi-tier system.

As part of future work, instead of bailing out on whole functions when irreducible control flow is detected, reducible loops could be handled correctly while irreducible loops are not detected as loops at all. This would be useful for functions with many loops where few isolated irreducible loops would then not prevent OSR compilation for the others as well. Alternatively, support for irreducible loops could be implemented for Sulong as suggested in Section 3.6. For a more detailed evaluation, the used benchmarks could be analyzed and compared in terms of their code structure (number of loops, call sites, etc.).

References

[1] Edd Barrett, Carl Friedrich Bolz-Tereick, Rebecca Killick, Sarah Mount, and Laurence Tratt. 2017. Virtual Machine Warmup Blows Hot and Cold. *Proc. ACM Program. Lang.* 1, OOPSLA, Article 52 (Oct. 2017), 27 pages. https://doi.org/10.1145/3133876

[2] Benoit Daloze, Stefan Marr, Daniele Bonetta, and Hanspeter Mössenböck. 2016. Efficient and Thread-Safe Objects for Dynamically-Typed Languages. In *Proceedings of the 2016 ACM International Conference*

on Object Oriented Programming Systems Languages & Applications (OOPSLA'16). 642–659.

[3] Daniele Cono D'Elia and Camil Demetrescu. 2016. Flexible On-stack Replacement in LLVM. In *Proceedings of the 2016 International Symposium on Code Generation and Optimization (CGO '16).* ACM, New York, NY, USA, 250–260. https://doi.org/10.1145/2854038.2854061

[4] Daniele Cono D'Elia and Camil Demetrescu. 2018. On-stack Replacement, Distilled. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2018).* ACM, New York, NY, USA, 166–180. https://doi.org/10.1145/3192366.3192396

[5] Edsger W. Dijkstra. 1972. Structured Programming. Academic Press Ltd., London, UK, UK, Chapter Chapter I: Notes on Structured Programming, 1–82. http://dl.acm.org/citation.cfm?id=1243380.1243381

[6] Gilles Duboscq, Lukas Stadler, Thomas Würthinger, Doug Simon, Christian Wimmer, and Hanspeter Mössenböck. 2013. Graal IR: An Extensible Declarative Intermediate Representation. In *Proceedings of the Asia-Pacific Programming Languages and Compilers Workshop.*

[7] Gilles Duboscq, Thomas Würthinger, Lukas Stadler, Christian Wimmer, Doug Simon, and Hanspeter Mössenböck. 2013. An Intermediate Representation for Speculative Optimizations in a Dynamic Compiler. In *Proceedings of the 7th ACM Workshop on Virtual Machines and Intermediate Languages (VMIL '13).* ACM, New York, NY, USA, 1–10. https://doi.org/10.1145/2542142.2542143

[8] Ana M. Erosa and Laurie J. Hendren. 1994. Taming Control Flow: A Structured Approach to Eliminating Goto Statements. In *ICCL.* https://doi.org/10.1109/ICCL.1994.288377

[9] Stephen J. Fink and Feng Qian. 2003. Design, Implementation and Evaluation of Adaptive Recompilation with On-stack Replacement. In *Proceedings of the International Symposium on Code Generation and Optimization: Feedback-directed and Runtime Optimization (CGO '03).* IEEE Computer Society, Washington, DC, USA, 241–252.

[10] Yoshihiko Futamura. 1999. Partial Evaluation of Computation Process–An Approach to a Compiler-Compiler. *Higher-Order and Symbolic Computation* 12, 4 (01 Dec 1999), 381–391. https://doi.org/10.1023/A:1010095604496

[11] Matthias Grimmer, Stefan Marr, Mario Kahlhofer, Christian Wimmer, Thomas Würthinger, and Hanspeter Mössenböck. 2017. Applying Optimizations for Dynamically-typed Languages to Java. In *Proceedings of the 14th International Conference on Managed Languages and Runtimes (ManLang 2017).* ACM, New York, NY, USA, 12–22. https://doi.org/10.1145/3132190.3132202

[12] Paul Havlak. 1997. Nesting of Reducible and Irreducible Loops. *ACM Trans. Program. Lang. Syst.* 19, 4 (July 1997), 557–567. https://doi.org/10.1145/262004.262005

[13] Urs Hölzle, Craig Chambers, and David Ungar. 1992. Debugging Optimized Code with Dynamic Deoptimization. In *Proceedings of the ACM SIGPLAN 1992 Conference on Programming Language Design and Implementation (PLDI '92).* ACM, New York, NY, USA, 32–43. https://doi.org/10.1145/143095.143114

[14] Urs Hölzle and David Ungar. 1994. A Third-generation SELF Implementation: Reconciling Responsiveness with Performance. In *Proceedings of the Ninth Annual Conference on Object-oriented Programming Systems, Language, and Applications (OOPSLA '94).* ACM, New York, NY, USA, 229–243. https://doi.org/10.1145/191080.191116

[15] HotSpot JVM 2019. Java Version History (J2SE 1.3). http://en.wikipedia.org/wiki/Java_version_history

[16] Chris Lattner and Vikram Adve. 2004. LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In *Proceedings of the International Symposium on Code Generation and Optimization: Feedback-directed and Runtime Optimization (CGO '04).* IEEE Computer Society, Washington, DC, USA, 75–.

[17] David Leopoldseder, Lukas Stadler, Christian Wimmer, and Hanspeter Mössenböck. 2015. Java-to-JavaScript Translation via Structured Control Flow Reconstruction of Compiler IR. In *Proceedings of the 11th*

Symposium on Dynamic Languages (DLS 2015). ACM, New York, NY, USA, 91–103. https://doi.org/10.1145/2816707.2816715

[18] David Leopoldseder, Lukas Stadler, Thomas Würthinger, Josef Eisl, Doug Simon, and Hanspeter Mössenböck. 2018. Dominance-based Duplication Simulation (DBDS): Code Duplication to Enable Compiler Optimizations. In *Proceedings of the 2018 International Symposium on Code Generation and Optimization (CGO 2018).* ACM, New York, NY, USA, 126–137. https://doi.org/10.1145/3168811

[19] Fabio Niephaus, Tim Felgentreff, and Robert Hirschfeld. 2018. GraalSqueak: A Fast Smalltalk Bytecode Interpreter Written in an AST Interpreter Framework. In *Proceedings of the 13th Workshop on Implementation, Compilation, Optimization of Object-Oriented Languages, Programs and Systems (ICOOOLPS '18).* ACM.

[20] G. Ramalingam. 1999. Identifying Loops in Almost Linear Time. *ACM Trans. Program. Lang. Syst.* 21, 2 (March 1999), 175–188. https://doi.org/10.1145/316686.316687

[21] Manuel Rigger, Matthias Grimmer, Christian Wimmer, Thomas Würthinger, and Hanspeter Mössenböck. 2016. Bringing Low-level Languages to the JVM: Efficient Execution of LLVM IR on Truffle. In *Proceedings of the 8th International Workshop on Virtual Machines and Intermediate Languages (VMIL 2016).* ACM, New York, NY, USA, 6–15. https://doi.org/10.1145/2998415.2998416

[22] Manuel Rigger, Roland Schatz, René Mayrhofer, Matthias Grimmer, and Hanspeter Mössenböck. 2018. Sulong, and Thanks for All the Bugs: Finding Errors in C Programs by Abstracting from the Native Execution Model. In *Proceedings of the Twenty-Third International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '18).* ACM, New York, NY, USA, 377–391. https://doi.org/10.1145/3173162.3173174

[23] Yukinori Sato, Yasushi Inoguchi, and Tadao Nakamura. 2011. On-the-fly Detection of Precise Loop Nests Across Procedures on a Dynamic Binary Translation System. In *Proceedings of the 8th ACM International Conference on Computing Frontiers (CF '11).* ACM, New York, NY, USA, Article 25, 10 pages. https://doi.org/10.1145/2016604.2016634

[24] Vugranam C. Sreedhar, Guang R. Gao, and Yong-Fong Lee. 1996. Identifying Loops Using DJ Graphs. *ACM Trans. Program. Lang. Syst.* 18, 6 (Nov. 1996), 649–658. https://doi.org/10.1145/236114.236115

[25] Lukas Stadler, Thomas Würthinger, and Hanspeter Mössenböck. 2014. Partial Escape Analysis and Scalar Replacement for Java. In *Proceedings of Annual IEEE/ACM International Symposium on Code Generation and Optimization (CGO '14).* ACM, New York, NY, USA, Article 165, 10 pages. https://doi.org/10.1145/2544137.2544157

[26] Bjarne Steensgaard. 1993. Sequentializing Program Dependence Graphs for Irreducible Programs.

[27] Robert Tarjan. 1971. Depth-first search and linear graph algorithms. In *12th Annual Symposium on Switching and Automata Theory (swat 1971).* 114–121. https://doi.org/10.1109/SWAT.1971.10

[28] Robert Tarjan. 1973. Testing Flow Graph Reducibility. In *Proceedings of the Fifth Annual ACM Symposium on Theory of Computing (STOC '73).* ACM, New York, NY, USA, 96–107. https://doi.org/10.1145/800125.804040

[29] Linda Torczon and Keith Cooper. 2011. *Engineering A Compiler* (2nd ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

[30] Christian Wimmer. 2004. Linear Scan Register Allocation for the Java HotSpot™ Client Compiler. Institute for System Software, Johannes Kepler University Linz. Master's Thesis.

[31] Thomas Würthinger, Christian Wimmer, Christian Humer, Andreas Wöß, Lukas Stadler, Chris Seaton, Gilles Duboscq, Doug Simon, and Matthias Grimmer. 2017. Practical Partial Evaluation for High-performance Dynamic Language Runtimes. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2017).* ACM, New York, NY, USA, 662–676. https://doi.org/10.1145/3062341.3062381

[32] Thomas Würthinger, Christian Wimmer, Andreas Wöß, Lukas Stadler, Gilles Duboscq, Christian Humer, Gregor Richards, Doug Simon, and Mario Wolczko. 2013. One VM to Rule Them All. In *Proceedings of the 2013 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming & Software (Onward! 2013)*. ACM, New York, NY, USA, 187–204. https://doi.org/10.1145/2509578.2509581

[33] Alon Zakai. 2011. Emscripten: an LLVM-to-JavaScript compiler. In *Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and Applications*. 301–312.