Leveraging LLVM's ScalarEvolution for Symbolic Data Cache Analysis

Valentin Touzeau
Saarland University
Saarland Informatics Campus
Saarbrücken, Germany
valentin.touzeau@cs.uni-saarland.de

Jan Reineke
Saarland University
Saarland Informatics Campus
Saarbrücken, Germany
reineke@cs.uni-saarland.de

Abstract—While instruction cache analysis is essentially a solved problem, data cache analysis is more challenging. In contrast to instruction fetches, the data accesses generated by a memory instruction may vary with the program's inputs and across dynamic occurrences of the same instruction in loops.

We observe that the plain control-flow graph (CFG) abstraction employed in classical cache analyses is inadequate to capture the dynamic behavior of memory instructions. On top of plain CFGs, accurate analysis of the underlying program's cache behavior is impossible.

Thus, our first contribution is the definition of a more expressive program abstraction coined symbolic control-flow graphs, which can be obtained from LLVM's ScalarEvolution analysis. To exploit this richer abstraction, our main contribution is the development of *symbolic data cache analysis*, a smooth generalization of classical LRU must analysis from plain to symbolic control-flow graphs.

The experimental evaluation demonstrates that symbolic data cache analysis consistently outperforms classical LRU must analysis both in terms of accuracy and analysis runtime.

Index Terms—cache analysis, chains of recurrences, data caches, symbolic analysis

I. INTRODUCTION

Due to technological developments, the latency of accesses to DRAM-based main memory is much higher than the latency of arithmetic and logic computations on processor cores. This "memory gap" is commonly tackled by a hierarchy of caches between the processor cores and main memory.

In the presence of caches, the latency of a memory access may vary widely depending on the level of the memory hierarchy that is able to serve the access. Hits to the first-level cache take just a few processor cycles, while accesses that miss in all cache levels and thus need to be served by main memory can take hundreds of cycles.

This variability is a challenge in the context of real-time systems, where it is necessary to bound a program's worst-case execution time (WCET) [1] to guarantee that safety-critical applications meet all of their deadlines. For accurate WCET analysis it is thus imperative to take caches into account. The timing variability induced by caches also introduces security challenges. Implementations of cryptographic algorithms have been shown to be vulnerable to cache timing attacks [2] and cache analysis [3], [4], [5] may help to uncover such vulnerabilities or prove their absence.

Cache analysis aims to statically characterize a program's cache behavior by classifying memory accesses in the program as guaranteed cache hits or misses. One perspective on cache analysis is that it is the composition of two phases:
1) A transformation of the program under analysis into a simpler program abstraction: a control-flow graph (CFG) whose edges are decorated with memory accesses.
2) An analysis of this decorated CFG that classifies accesses as "always hit", "always miss", or "unknown".

For instruction cache analysis this two-phase approach works well, as CFGs accurately captures most programs' instruction fetch sequences. For data cache analysis, however, a plain CFG abstraction can be highly inaccurate. Consider for example the following simple loop:

```
for (int x = 0; x < 100; x++)
    sum += A[x]
```

In each iteration of the loop a different address is accessed, and so the corresponding edge in the CFG needs to be conservatively decorated with all possible addresses. The order in which the array elements are accessed is lost and it becomes impossible to make accurate predictions about the program's cache behavior. A program abstraction that more precisely captures a program's memory access behavior is thus needed.

Our first contribution is the definition of symbolic control-flow graphs in Section IV, which is our formalization of the output of LLVM's ScalarEvolution analysis [6], [7]. Symbolic CFGs accurately capture the link between loop iterations and accessed memory blocks via chains of recurrences [8], [9] in a manner that is amenable to static analysis.

To exploit this more expressive program abstraction our main contribution is the development of *symbolic data cache analysis* in Section V, a smooth generalization of Ferdinand's classical LRU must analysis [10], [11] from plain to symbolic control-flow graphs. To fully realize the potential of symbolic data cache analysis we further introduce a context-sensitive analysis combining loop peeling and unrolling in Section VI and various implementation tricks in Section VII.

The experimental evaluation on the PolyBench benchmark suite in Section VIII demonstrates that symbolic cache analysis compares favorably to classical LRU must analysis both in terms of accuracy and analysis runtime.

II. BACKGROUND

A. Caches

Caches are fast but small memories that buffer parts of the large but slow main memory in order to bridge the speed gap between the processor and main memory. Caches consist of *cache lines*, which store data at the granularity of memory blocks $b \in \mathcal{B}$. Memory blocks usually comprise a power-of-two number of bytes BS, e.g. 64 bytes, so that the block $block(a)$ that address a maps to is determined by truncating the least significant bits of a, i.e., $block(a) = \lfloor a/BS \rfloor$. In order to facilitate an efficient cache lookup, the cache is organized in *sets* such that each memory block maps to a unique cache set $set(b) = b \bmod NS$, where NS is the number of sets. The number of cache lines k in each cache set is called the *associativity* of the cache.

If an accessed block resides in the cache, the access *hits* the cache. Upon a cache *miss*, the block is loaded from the next level of the hierarchy. Then, another memory block has to be evicted due to the limited size of the cache. The block to evict is determined by the *replacement policy*. In this paper, we assume the least-recently-used (LRU) policy, which replaces the block that has been accessed least recently. A memory block b hits in an LRU cache of associativity k if b has been accessed previously and less than k distinct blocks in the same cache set have been accessed since the last access to b. LRU is generally considered to be the most predictable replacement policy [12].

In this paper, we refer to the *age* of block b as the number of distinct blocks in the same cache set that have been accessed since the last access to b. Thus, an access to block b hits the cache if and only if its age is less than the associativity k.

B. Control-Flow Graphs as a Program Representation

Control-flow graphs (CFGs) are a program representation commonly employed in compilers and static analysis tools. A CFG is a directed graph $\mathcal{G} = (V, E, v_0)$, whose vertices V correspond to control locations in the program including the initial control location $v_0 \in V$, and whose edges E represent the possible control flow between the graph's vertices.

For the purpose of cache analysis, CFGs are used to represent the possible sequences of memory accesses generated by the underlying program. To this end, each edge of the CFG is decorated with the set of memory addresses that may be accessed when control passes along that edge.

As defined above, CFGs over-approximate the behavior of the program they represent as they do not capture the functional semantics of the instructions. In particular, all paths through the graph are assumed to be feasible even if, in reality, some are not. Also, and this is particularly problematic for data cache analysis, the CFG representation does not capture the dependence of the accessed memory addresses on the loop iterations. We will see in Section III how this may lead to gross overapproximations of the number of cache misses. To overcome this issue, we introduce symbolic control-flow graphs in Section IV.

C. Ferdinand's May and Must Cache Analysis

The aim of Ferdinand's may and must cache analyses [10], [11] is to classify memory accesses in a CFG as definite hits or definite misses. As noted before, under LRU replacement, an access results in a cache hit if and only if the age of the accessed block is less than the cache's associativity.

Instead of computing all reachable concrete cache states, must and may analysis operates on abstract cache states, which maintain upper and lower bounds on the age of each memory block. Each block's bounds hold independently of the ages of other blocks. This allows for a compact representation of large sets of concrete cache states. For example, the abstract must cache state $\lambda b.\infty$ that maps every block to age bound ∞ compactly represents all possible concrete cache states. As the correlation between the ages of different blocks is lost, the resulting analysis is not exact. However, recent work [13], [14], [15] has shown that the loss in precision due to this abstraction is small in practice. Our symbolic data cache analysis introduced in Section V, can be seen as a smooth generalization of Ferdinand's must analysis to symbolic control-flow graphs.

III. ILLUSTRATIVE EXAMPLE

As an illustrative example of the drawbacks of cache analysis performed on plain CFGs, consider the simple program in Figure 1a. The first loop of our example program iterates across array A in the forward direction, while the second loop iterates across the same array in the opposite direction.

A. Intuitive Cache Analysis

Let us intuitively analyze the program's cache behavior. For this analysis, we will assume a tiny set-associative cache with LRU replacement consisting of 2 cache sets, an associativity of 4, and cache lines of size 8 bytes. Thus, the cache has a capacity of $2 \cdot 4 \cdot 8 = 64$ bytes. Also assume that integers are of size 4 bytes, and so the cache can hold 16 array cells.

Assuming an initially empty cache, the first loop does not exhibit any temporal locality, as each array cell is only touched once. However, it does exhibit spatial locality, as pairs of adjacent array cells may reside in the same memory blocks. Thus, every other iteration of the first loop will result in a cache hit.

The second loop accesses the same array cells as the first. Now it depends on the cache geometry whether and to what extent this temporal locality can be exploited. Under our assumptions, the cache will contain array cells $A[84], A[85], \ldots, A[99]$ after the first loop has terminated. Thus, the first 16 iterations of the second loop hit the cache. The remaining iterations profit only from spatial locality as the first loop did, hitting in every other iteration.

B. Traditional Cache Analysis

Under Ferdinand's must cache analysis [10], [11] and recent exact analyses [13], [14] the program is abstracted via its CFG and the CFG's edges are annotated with the sets of memory blocks that may be accessed while executing the corresponding part of the program as discussed in Section II-B. Figure 1b shows the plain CFG abstraction for our example

```
int A[100];

for (int x = 0; x < 100; x++)
    sum += A[x]
for (int y = 99; y >= 0; y--)
    sum -= A[y]
```

(a) Iteration-dependent data accesses.   (b) Plain control-flow-graph abstraction.   (c) Symbolic control-flow-graph abstraction.
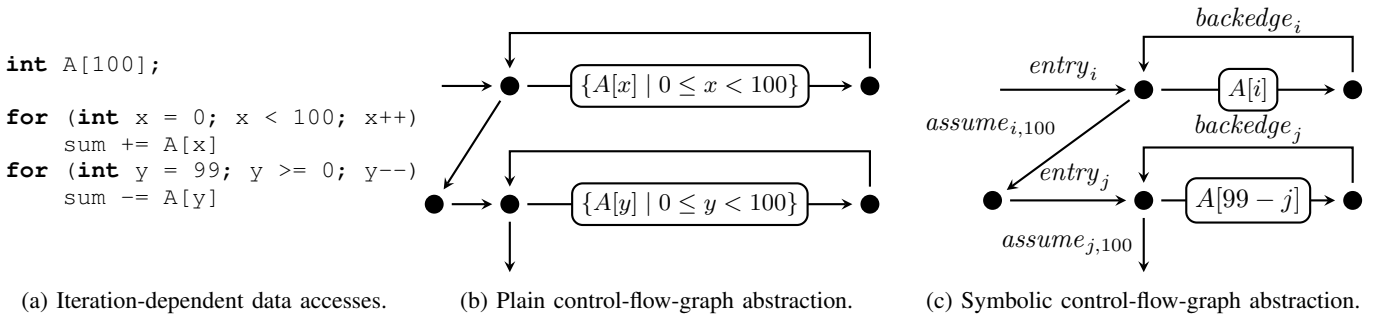
Fig. 1: Simple program and its plain and symbolic control-flow-graph abstractions.

program. While this abstraction is adequate for instruction cache analysis as the same instructions are accessed in each loop iteration, it is inadequate for data cache analysis, as the link between the loop iteration and the accessed address is lost. As a consequence, it is impossible to predict any of the memory accesses in the program to be cache hits or misses.

If the entire set of memory blocks that can potentially be accessed fits into the cache, then persistence analysis [16], [17], [18], [19], [15] may deduce that each of these blocks results in at most one cache miss. However, in our example, the array A does not fully fit into the cache, and so persistence analysis is of no use here.

C. Symbolic Control-Flow Graphs and Cache Analysis

a) Symbolic Control-Flow Graphs: We have seen that the plain CFG abstraction is inadequate for data cache analysis, because the link between loop iterations and accessed memory blocks is lost. Thus, our first step towards accurate data cache analysis is to employ what we coin *symbolic CFGs*, a simple yet powerful program representation that concisely captures the link between loop iterations and accessed data. Symbolic CFGs are our formalization of the output of LLVM's ScalarEvolution Analysis [6], [7].

Figure 1c shows a symbolic CFG for our example program. In a symbolic CFG—where possible—the addresses of memory accesses are expressed in terms of the loop iterations of their enclosing loops. To this end, symbolic CFGs make it explicit when a loop is entered and when a new loop iteration begins. These transitions are indicated by annotating edges with $entry_i$ and $backedge_i$, where i is the identifier of a loop. Consider the edge annotated with $A[99 - j]$. This is to be interpreted as follows: In an execution of the program, let $\sigma(j)$ be the number of times that $backedge_j$ has been traversed since the last time $entry_j$ has been taken. Then, the accessed address is $A[99 - \sigma(j)]$.

For some loops, ScalarEvolution is also able to derive the exact number of times that a loop's back edges are taken from entry to exit. To express such information, symbolic CFGs may contain $assume_{i,e}$ statements, where e is an expression that may refer to loop variables other than i itself. An edge annotated with $assume_{i,e}$ can only be taken if the value of i is equal to the value of expression e. In our example, the back edges of both loops are taken exactly 100 times, and so the exit edges of both loops are annotated accordingly with assume statements.

We define symbolic CFGs in Section IV. There we also discuss multivariate chains of recurrences [8], [9], [20], which are used to represent access expressions and loop bounds.

b) Symbolic Cache Analysis: Symbolic CFGs are useful for data cache analysis as they capture a program's memory access behavior more precisely than plain CFGs. In fact, in our example, the symbolic CFG perfectly captures the sequence of memory accesses generated by the program.

It remains to define a static analysis that can efficiently exploit this information. Simply applying Ferdinand's must analysis would not be fruitful as the underlying abstraction does not capture the relation between loop iterations and cache states. A relatively straightforward approach would be to virtually unroll the loops for the sake of the analysis, resulting in an exploded plain CFG in which each edge could once more be annotated with a concrete memory access. Ferdinand's must analysis could then be employed successfully on this exploded plain CFG. However, this approach would be very costly, in particular for programs with large loop bounds. We are thus seeking a precise analysis whose runtime is independent of the loop bounds of the program.

To this end, our first basic idea are symbolic cache states that capture how cache states depend on the loop iteration. To motivate symbolic cache states, consider Figures 2a and 2b, which show the concrete cache states at the ends of iterations 15 and 17 of the first loop from our example program. As we assume cache lines of size 8 bytes, each line contains two cells of the array. We represent each memory block by the first array cell mapping to that block. Our idea is to represent memory blocks *symbolically* in terms of the values of loop variables. For example, $A[14]$ can be expressed as $A[i-1]$ if i's value is 15. If we represent the states from Figures 2a and 2b in this way we arrive at the symbolic cache state depicted in Figure 2c. Furthermore, the *same* symbolic state will be reached at the end of each odd loop iteration, starting from iteration 15.

Like Ferdinand's must analysis our symbolic data cache analysis determines upper bounds on the ages of memory blocks. However, instead of associating bounds with concrete memory blocks, it associates these bounds with symbolic memory blocks. A peculiar consequence of this abstraction is

that symbolic cache states also need to be updated when the value of a loop variable changes. For example, if the back edge of the first loop is taken to move from iteration 15 ($17, 19, \dots$) to iteration 16 ($18, \dots$), then the symbolic cache state needs to be updated to account for incrementing i. The resulting symbolic cache state is depicted in Figure 2d. We show how to lift Ferdinand's analysis to symbolic cache analysis in Section V.

In our example, one can observe that the symbolic cache states "stabilize" in odd and even loop iterations after the cache has been filled in the first 16 iterations. Thus the analysis needs to distinguish the first 16 loop iterations from the rest, and odd from even loop iterations in the remainder of the execution. This can be achieved by context-sensitive analysis [21], [22], [23]. In Section VI we introduce a context-sensitive analysis that can be configured to virtually peel and unroll the loops appropriately for a given cache configuration.

IV. SYMBOLIC CONTROL-FLOW GRAPHS

We have seen the intuition behind symbolic control-flow graphs in Section III-C. One aspect that has been left undefined there is the shape of expressions used to represent memory accesses and loop bounds. We fill this gap in Section IV-A, which is then used in the formal definition of symbolic control-flow graphs in Section IV-B. In Section IV-C we provide a semantics for symbolic CFGs, which will allow us to make formal correctness statements about the symbolic data cache analysis introduced in Section V.

A. Multivariate Chains of Recurrences

We employ *multivariate chains of recurrences* [8], [9], [20] (short: MCRs) as the formalism for expressions. Given a subset of a program's loop variables $S \subseteq \mathit{LoopVar}$, the set $M(S)$ of MCRs over S is given by the following grammar:

$$e := n \in \mathbb{Z}$$
$$\mid e_1 \; bop \; e_2 \; \text{where} \; bop \in \{+, -, \cdot\}, \text{and} \; e_1, e_2 \in M(S)$$
$$\mid \{e_1, +, e_2\}_i \; \text{where} \; i \in S, e_1 \in M(S \setminus \{i\}), e_2 \in M(S)$$

Thus, expressions can (i) be constants; (ii) they can be formed from subexpressions via addition, subtraction, and multiplication; and (iii) they can be *add recurrences* of the form $\{e_1, +, e_2\}_i$.

Given an environment $\sigma : \mathit{LoopVar} \to \mathbb{N}$ assigning loop variables to their values, an MCR can be evaluated as follows:

$$\llbracket n \rrbracket_\sigma := n$$
$$\llbracket e_1 \; bop \; e_2 \rrbracket_\sigma := \llbracket e_1 \rrbracket_\sigma \; bop \; \llbracket e_2 \rrbracket_\sigma$$
$$\llbracket \{e_1, +, e_2\}_i \rrbracket_\sigma := \llbracket e_1 \rrbracket_\sigma + \sum_{k=0}^{\sigma(i)-1} \llbracket e_2 \rrbracket_{\sigma[i \mapsto k]}$$

By $\sigma[i \mapsto v]$ we denote the function that maps i to v and otherwise is the same as σ. Thus, in an add recurrence e_1 can be seen as the initial value, and e_2 as the increment. For example:

$$\llbracket \{23, +, 4\}_i \rrbracket_\sigma = \llbracket 23 \rrbracket_\sigma + \sum_{k=0}^{\sigma(i)-1} \llbracket 4 \rrbracket_{\sigma[i \mapsto k]} = 23 + 4 \cdot \sigma(i)$$

Thus, the array access $A[i]$ from our example can be expressed as $\{A, +, 4\}_i$, assuming $A \in \mathbb{N}$ is the base address of the array and each element of the array is of size 4. Similarly, the array access $A[99-j]$ can be expressed as $\{A+396, +, -4\}_j$.

Nested add recurrences can represent arbitrary polynomial functions, e.g. $\llbracket \{0, +, \{5, +, 1\}_i\}_j \rrbracket_\sigma = 5 \cdot \sigma(j) + \sigma(i) \cdot \sigma(j)$ and $\llbracket \{0, +, \{0, +, 2\}_i\}_i \rrbracket_\sigma = (\sigma(i) - 1) \cdot \sigma(i)$.

In order to update symbolic cache states upon incrementing loop variables, we need a shift operation on MCRs that adapts an expression to account for the increment of a variable. Such an operation should thus satisfy the following equality: $\llbracket Sh(e, i) \rrbracket_{\sigma[i \mapsto \sigma(i)+1]} = \llbracket e \rrbracket_\sigma$. For example, $Sh(\{A, +, 1\}_i, i)$ could be $\{A-1, +, 1\}_i$.

To implement such a shift operation we need an initialization operation that satisfies $\llbracket Init(e, i) \rrbracket_\sigma = \llbracket e \rrbracket_{\sigma[i \mapsto 0]}$, which can be implemented as follows:

$$Init(n, i) := n$$
$$Init(e_1 \; bop \; e_2, i) := Init(e_1, i) \; bop \; Init(e_2, i)$$
$$Init(\{e_1, +, e_2\}_j, i) := \begin{cases} e_1 & : i = j \\ \{Init(e_1, i), +, Init(e_2, i)\}_j & : i \neq j \end{cases}$$

This allows us to implement the $Sh(e, i)$ operation:

$$Sh(n, i) := n$$
$$Sh(e_1 \; bop \; e_2, i) := Sh(e_1, i) \; bop \; Sh(e_2, i)$$
$$Sh(\{e_1, +, e_2\}_i, i) := \{e_1 - Init(Sh(e_2, i), i), +, Sh(e_2, i)\}_i$$
$$Sh(\{e_1, +, e_2\}_j, i) := \{Sh(e_1, i), +, Sh(e_2, i)\}_j$$

The correctness of the $Init(e, i)$ and $Sh(e, i)$ can be shown by structural induction, which we omit here for brevity.

To take into account loop bounds when exiting a loop provided by assume statements in our symbolic control-flow graphs, we rely on a substitution operation with the following semantics: $\llbracket Sub(e, i, expr) \rrbracket_\sigma = \llbracket e \rrbracket_{\sigma[i \mapsto \llbracket expr \rrbracket_\sigma]}$. A heuristic implementation of $Sub(e, i, expr)$, which may fail on some inputs, can be realized as follows:

$$Sub(e, i, expr) :=$$
$$\begin{cases} e & \text{if } e \in \mathbb{Z} \\ s_1 \; bop \; s_2 & \text{if } e = e_1 \; bop \; e_2 \wedge s_1 \neq \mathit{fail} \wedge s_2 \neq \mathit{fail} \\ \{s_1, +, s_2\}_j & \text{if } e = \{e_1, +, e_2\}_j \wedge j \neq i \\ & \quad \wedge \; s_1 \neq \mathit{fail} \wedge s_2 \neq \mathit{fail} \\ e_1 + e_2 \cdot expr & \text{if } e = \{e_1, +, e_2\}_i \wedge i \notin e_2 \\ \mathit{fail} & \text{otherwise} \end{cases}$$

where $s_1 = Sub(e_1, i, expr)$ and $s_2 = Sub(e_2, i, expr)$.

Engelen [9], [20] provides a set of rewrite rules for MCRs that are proven to be confluent and terminating. We rely on these rewrite rules to bring MCRs into a normal form.

In general, not all accesses generated in a program can be accurately captured by an MCR. As an example, consider the accesses generated by a loop traversing a dynamic heap data structures, such as a linked list. To soundly represent such accesses we introduce *unknown accesses*, denoted $\mathbf{X}$, which are interpreted to take any possible value. Fortunately, $\mathbf{X}$ is

Set 0	Set 1
$A[12]$	$A[14]$
$A[8]$	$A[10]$
$A[4]$	$A[6]$
$A[0]$	$A[2]$

(a) Cache state at the end of iteration 15.

Set 0	Set 1
$A[16]$	$A[14]$
$A[12]$	$A[10]$
$A[8]$	$A[6]$
$A[4]$	$A[2]$

(b) Cache state at the end of iteration 17.

Set a	Set b
$A[i-3]$	$A[i-1]$
$A[i-7]$	$A[i-5]$
$A[i-11]$	$A[i-9]$
$A[i-15]$	$A[i-13]$

(c) Symbolic cache state at the end of iterations $15, 17, 19, \ldots$

Set a	Set b
$A[i-4]$	$A[i-2]$
$A[i-8]$	$A[i-6]$
$A[i-12]$	$A[i-10]$
$A[i-16]$	$A[i-14]$

(d) Symbolic cache state at the start of iterations $16, 18, 20, \ldots$

Fig. 2: Cache states that arise during the execution of the first loop.

only rarely needed in the analysis of real-time applications, in which dynamic data structures are uncommon.

B. Symbolic Control-Flow Graphs

A *symbolic CFG* is a tuple $\mathcal{G} = (V, E, LoopVar, v_0)$, where V is a set of vertices and $E \subseteq V \times \mathcal{D} \times V$ is a set of edges, $LoopVar$ is a set of loop variables, and $v_0 \in V$ is a vertex with no incoming edges marking the program entry.

Edges are decorated with *accesses* $\mathcal{A}$ and *statements* $\mathcal{S}$, i.e., $\mathcal{D} = \mathcal{S} \cup \mathcal{A}$:

- Accesses are MCRs or unknowns:
  $\mathcal{A} := M(LoopVar) \cup \{\mathbf{X}\}$
- Statements either mark the entry to a loop ($entry_i$), a back edge of a loop ($backedge_i$), or an assumption on the value of a loop variable ($assume_{i,e}$):

$$\mathcal{S} := \{entry_i, backedge_i \mid i \in LoopVar\}$$
$$\cup \{assume_{i,e} \mid i \in LoopVar, e \in M(LoopVar \setminus \{i\})\}$$

C. Semantics of Symbolic Control-Flow Graphs

The state of an execution of a symbolic control-flow graph consists of two parts: The program state $\sigma_p \in \Sigma_p$ and the cache state $\sigma_c \in \Sigma_c$.

We represent the program state by a map σ_p that maps loop variables to their values. Each loop variable counts the number of times that the loop back edge has been taken since last entering the loop. The program semantics of a symbolic CFG is then captured by a transformer $update_\mathcal{S}$ that captures the effects of statements on program states.

$$update_\mathcal{S}(\sigma_p, s) :=$$
$$\begin{cases} \sigma_p[i \mapsto 0] & \text{if } s = entry_i \\ \sigma_p[i \mapsto \sigma_p(i) + 1] & \text{if } s = backedge_i \\ \bot_p & \text{if } s = assume_{i,expr} \wedge \sigma_p(i) \neq \llbracket expr \rrbracket_{\sigma_p} \\ \sigma_p & \text{if } s = assume_{i,expr} \wedge \sigma_p(i) = \llbracket expr \rrbracket_{\sigma_p} \end{cases}$$

Note that we use the special value $\bot_p$ to represent unreachable program states, i.e. those not satisfying an assume statement.

We represent cache states as maps σ_c from memory blocks to ages, i.e. σ_c tracks the age of each memory block in its cache set: $\sigma_c \in \mathcal{B} \to \mathbb{N}$. The LRU replacement policy is then captured by the following transformer:

$$update_{LRU}(\sigma_c, b) := \lambda b' \in \mathcal{B}.$$
$$\begin{cases} 0 & \text{if } b = b' \\ \sigma_c(b') & \text{else if } set(b) \neq set(b') \\ \sigma_c(b') & \text{else if } \sigma_c(b) \leq \sigma_c(b') \\ \sigma_c(b') + 1 & \text{otherwise} \end{cases}$$

To paraphrase the above definition: (i) The accessed block b attains age 0. (ii) The ages of blocks in other cache sets ($set(b) \neq set(b')$) do not change. (iii) If the accessed block b is younger than block b', then b has already been accounted for in the age of b', and thus the age of b' should not increase. (iv) Otherwise, b maps to the same cache set as b' and is older than b' and thus the access increases the age of b'.

The complete state of the system is a pair (σ_p, σ_c) and we can capture its evolution upon arbitrary CFG decorations by combining the previous transformers into a single one and accounting for unknown accesses:

$$update((\sigma_p, \sigma_c), d) :=$$
$$\begin{cases} \{(update_\mathcal{S}(\sigma_p, d), \sigma_c)\} & \text{if } d \in \mathcal{S} \\ \{(\sigma_p, update_{LRU}(\sigma_c, block(\llbracket d \rrbracket_{\sigma_p})))\} & \text{if } d \in \mathcal{A} \setminus \{\mathbf{X}\} \\ \{(\sigma_p, update_{LRU}(\sigma_c, b)) \mid b \in \mathcal{B}\} & \text{if } d = \mathbf{X} \end{cases}$$

where $block$ maps addresses to the corresponding memory blocks (see Section II-A). Note that $update((\sigma_p, \sigma_c), d)$ maps to sets of states to capture the non-determinism introduced by unknown accesses. We lift $update$ to sets of states as follows:

$$update(S, d) := \{(\sigma'_p, \sigma'_c) \mid (\sigma_p, \sigma_c) \in S$$
$$\wedge (\sigma'_p, \sigma'_c) \in update((\sigma_p, \sigma_c), d) \wedge \sigma'_p \neq \bot_p\}$$

We drop unreachable states (where $\sigma'_p = \bot_p$) here.

We define the set of reachable states at each control location $R^C : V \to \mathcal{P}(\Sigma_p \times \Sigma_c)$ as the least solution to the following set of equations:

$$R^C(v_0) = \{(\lambda i.0, \sigma_c) \mid \sigma_c \in \Sigma_c\} \tag{1}$$

$$\forall v \in V \setminus \{v_0\} : R^C(v) = \bigcup_{(u,d,v) \in E} update(R^C(u), d) \tag{2}$$

Equation (1) captures that initially all loop variables are zero, while the initial cache state can be arbitrary. Equation (2)

captures that the reachable states at node v are determined by the reachable states at v's predecessor nodes u updated according to the CFG decoration between u and v. In keeping with abstract interpretation literature [24], we refer to R^C as the *collecting semantics*.

V. SYMBOLIC DATA CACHE ANALYSIS

Explicitly computing the collecting semantics R^C would be very costly and only possible at all if all loops were bounded. In this section, we lift Ferdinand's must analysis to symbolic control-flow graphs to obtain a tractable analysis.

A. Abstract Domain

As described earlier, Ferdinand's must analysis maps memory blocks to an upper bound on their maximum age in order to classify memory accesses as hits. Our analysis relies on a similar map, except that it maps symbolic blocks, represented via MCRs, to such age bounds. Our abstract domain is thus

$$\widehat{\sigma} \in \widehat{SymCache} = M(LoopVar) \hookrightarrow \{0, \dots, k-1, \infty\},$$

where $\hookrightarrow$ indicates that symbolic cache states are partial functions. We refer to the domain of a cache state $\widehat{\sigma}$, i.e., the set of MCRs for which $\widehat{\sigma}$ provides an age bound, as $dom(\widehat{\sigma})$.

If our analysis maps an MCR e to age x at program point v, it means that the memory block containing the address given by $[\![e]\!]_{\sigma_p}$ has age at most x for any program state σ_p reachable at v. This set of program and cache states associated with an abstract state $\widehat{\sigma}$ is captured by the concretization function γ:

$$\gamma(\widehat{\sigma}) := \{(\sigma_p, \sigma_c) \mid \forall e \in dom(\widehat{\sigma}) : \sigma_c(block([\![e]\!]_{\sigma_p})) \leq \widehat{\sigma}(e)\} \tag{3}$$

Similarly to the definition of the collecting semantics (see Equations (1) and (2)), which uses set unions to capture all possible behaviors of the program, we need a join operator on the abstract domain to summarize states from several incoming CFG edges. This join operator $\sqcup$ conservatively keeps, for each MCR, the maximum of the two upper bounds provided by the joined states: $\widehat{\sigma}_1 \sqcup \widehat{\sigma}_2 = \lambda e \in dom(\widehat{\sigma}_1) \cap dom(\widehat{\sigma}_2). \max\{\widehat{\sigma}_1(e), \widehat{\sigma}_2(e)\}$. This join operator is correct with respect to the concretization function:

Lemma 1 (Join Correctness). *For all $\widehat{\sigma}_1, \widehat{\sigma}_2 \in \widehat{SymCache}$:*

$$\gamma(\widehat{\sigma}_1) \cup \gamma(\widehat{\sigma}_2) \subseteq \gamma(\widehat{\sigma}_1 \sqcup \widehat{\sigma}_2)$$

The proofs of all lemmas and theorems can be found in the appendix.

B. Abstract Transformers

To reflect the cache updates upon memory accesses, we provide two abstract transformers: $\widehat{update}_{\mathcal{A}\setminus\{\mathbf{X}\}}$, for accesses to MCRs, and $\widehat{update}_{\mathbf{X}}$, for unknown accesses.

Unknown accesses can potentially increase the age of any block in the cache. Thus:

$$\widehat{update}_{\mathbf{X}}(\widehat{\sigma}) := \lambda e' \in dom(\widehat{\sigma}). \begin{cases} \widehat{\sigma}(e') + 1 & \text{if } \widehat{\sigma}(e') + 1 < k \\ \infty & \text{otherwise} \end{cases}$$

It is easy to prove that this transformer is correct:

Fig. 3: Lattice of alias relations.

Lemma 2 (Unknown Access Transformer Correctness). *For all $\widehat{\sigma} \in \widehat{SymCache}$, we have:*

$$update(\gamma(\widehat{\sigma}), \mathbf{X}) \subseteq \gamma(\widehat{update}_{\mathbf{X}}(\widehat{\sigma}))$$

The $\widehat{update}_{\mathcal{A}\setminus\{\mathbf{X}\}}$ transformer is similar to the one used by Ferdinand's must analysis; it rejuvenates the accessed symbolic block, and increases the ages of blocks in the same cache set that are younger than the accessed block.

The main difference lies in the fact that contrary to concrete memory blocks, which have a fixed address, it is not always obvious whether two symbolic blocks map to the same cache set or even to the same block. We thus rely on an auxiliary function $alias$, which, given two symbolic blocks, determines their alias relation.

There are six possible alias relations between two MCRs:
1) "Same block" sb: they map to the same memory block.
2) "Same set" ss: they map to the same cache set.
3) "Different set" ds: they map to different cache sets.
4) "Different block" db: they map to different blocks.
5) "Same set, diff. block" $ssdb$: conjunction of ss and db.
6) "Same block or different set" $sb+ds$: disjunction of ds and sb; can also be seen as the complement of $ssdb$.

As shown in [25], these relations form a lattice, whose Hasse diagram is shown in Figure 3. The alias relation of two MCRs e_1 and e_2 can be determined as follows, where BS is the size of memory blocks (in bytes) and NS is the number of cache sets:

$$alias(e_1, e_2) :=$$
$$\begin{cases} sb & \text{if } e_1 - e_2 = n \in \mathbb{Z} \wedge n = 0 \\ ds & \text{else if } e_1 - e_2 = n \in \mathbb{Z} \wedge \\ & \quad BS \leq n \bmod (NS \cdot BS) \leq (BS \cdot NS) - BS \\ sb+ds & \text{else if } e_1 - e_2 = n \in \mathbb{Z} \wedge -BS < n < BS \\ \top & \text{otherwise} \end{cases}$$

We assume a modulo operation based on *floored division*, i.e., $a \bmod n := a - n \cdot \lfloor a/n \rfloor$, so that $0 \leq a \bmod n < n$ for $n > 0$.

The alias relation between e_1 and e_2 is determined by computing the difference n of the two expressions. If the difference between e_1 and e_2 is not a constant expression, then no relation is established (last case). Otherwise, different relations can be deduced depending on the value of n:

(i) If n is 0, we can deduce sb.
(ii) Addresses whose difference is a multiple of the way size ($NS \cdot BS$) are guaranteed to be in the same cache set.

Conversely, if the difference between e_1 and e_2 is more than BS "away" from being a multiple of the way size, then e_1 and e_2 must map to different sets.

(iii) If e_1 and e_2 are close, i.e., less than a block size apart, they either map to the same block or to different sets.

Other aliasing relations, such as $ssdb$ and db could also be deduced, but are not useful in the following.

Using $alias$ to deduce the relation between symbolic blocks, we can formally define the transformer $\widehat{update}_{\mathcal{A}\setminus\{\mathbf{X}\}}$ to apply when performing the memory access associated with MCR e.

$$\widehat{update}_{\mathcal{A}\setminus\{\mathbf{X}\}}(\widehat{\sigma}, e) := \lambda e' \in dom(\widehat{\sigma}) \cup \{e\}.$$

$$\begin{cases} 0 & \text{if } alias(e, e') \sqsubseteq sb \\ \widehat{\sigma}(e') & \text{else if } alias(e, e') \sqsubseteq sb{+}ds \\ \widehat{\sigma}(e') & \text{else if } \widehat{\sigma}(e) \le \widehat{\sigma}(e') \\ \widehat{\sigma}(e') + 1 & \text{else if } \widehat{\sigma}(e') + 1 < k \\ \infty & \text{otherwise} \end{cases}$$

Unsurprisingly, the transformer closely resembles the definition of its concrete counterpart $update_{LRU}$. (i) As in the concrete case, the accessed symbolic block is rejuvenated to age 0, as are all symbolic blocks that represent the same block. (ii) A symbolic block that is in the $sb{+}ds$ relation to the accessed block retains its age, which is safe, as seen by the following case distinction: Either the block is actually the accessed block and it should get age 0, or it maps to a different set and its age should be unchanged (first two cases of $update_{LRU}$). (iii) If the accessed symbolic block e is younger than symbolic block e', then e has already been accounted for in the age of e', and thus the age of e' should not increase. (iv) The age of a block cannot increase by more than one upon a single access, so the fourth case is always safe. (v) We do not distinguish ages beyond k, as it is not helpful to classify accesses as hits or misses. Instead we summarize these with the safe upper bound ∞.

As for the join operator and for unknown accesses, we prove that the access transformer is correct:

Lemma 3 (MCR Access Transformer Correctness). *For all $\widehat{\sigma} \in Sym\widehat{C}ache$ and $e \in \mathcal{A}$, we have:*

$$update(\gamma(\widehat{\sigma}), e) \subseteq \gamma(\widehat{update}_{\mathcal{A}\setminus\{\mathbf{X}\}}(\widehat{\sigma}, e))$$

The $\widehat{update}_{\mathcal{A}\setminus\{\mathbf{X}\}}$ transformer described above captures the effect of memory accesses. As the symbolic cache states are tied to the program state via the concretization function given in (3), changes to the loop variables need to be accounted for by appropriately adapting our symbolic cache states. We thus provide a second transformer, $\widehat{update}_\mathcal{S}$, which captures the effect of program statements on symbolic cache states.

We define $\widehat{update}_\mathcal{S}$ separately for each type of statement. The case of a back edge is arguably the most interesting one. Each symbolic block e needs to be replaced by its shifted version when i is incremented, so that the expression preserves its original value, which is achieved as follows:

$$\widehat{update}_\mathcal{S}(\widehat{\sigma}, backedge_i) := \{(Sh(e, i), b) \mid (e, b) \in \widehat{\sigma}\} \quad (4)$$

For example, $Sh(\{A, +, 4\}_i, i) = \{A - 4, +, 4\}_i$, which corresponds to replacing $A[i]$ by $A[i-1]$ upon incrementing i. One might wonder whether the set defined in Equation (4) actually defines a function. This is indeed the case for MCRs in normal form [9], [20] for which $Sh(\cdot, i)$ is bijective.

Entering a loop entails resetting the corresponding loop variable to i. However, unless the prior value of i is known, there is no way of rewriting expressions involving the variable i accordingly. Thus, in such cases the information for the corresponding MCRs is discarded:

$$\widehat{update}_\mathcal{S}(\widehat{\sigma}, entry_i) := \{(e, b) \mid (e, b) \in \widehat{\sigma} \wedge i \notin e\}$$

Finally, assume statements allow the analysis to substitute the corresponding loop variable by the assumed expression. This allows to retain information across multiple loops or in nested loops, e.g. in our running example where data cached in the first loop is reused in the second loop.

$$\widehat{update}_\mathcal{S}(\widehat{\sigma}, assume_{i,expr}) := \\ red(\{(e', b) \mid (e, b) \in \widehat{\sigma} \wedge e' = Sub(e, i, expr) \ne fail\}),$$

where $red(S) := \{(e, b) \mid (e, b) \in S \wedge \forall (e, b') \in S : b' \ge b\}$.

The substitution may result in multiple expressions becoming equal, e.g., $Sub(\{0, +, 2\}_i, i, 10) = Sub(\{10, +, 1\}_i, i, 10)$. Then $red(S)$ keeps the best bound and thereby ensures that the resulting relation is still a function.

This abstract transformer for statements is also correct:

Lemma 4 (Statement Transformer Correctness). *For all $\widehat{\sigma} \in Sym\widehat{C}ache$ and $s \in \mathcal{S}$, we have:*

$$update(\gamma(\widehat{\sigma}), s) \subseteq \gamma(\widehat{update}_\mathcal{S}(\widehat{\sigma}, s))$$

C. Analysis Correctness and Termination

We can now merge the statement and access transformers into a single one that deals with the three kinds of decorations:

$$\widehat{update}(\widehat{\sigma}, d) := \begin{cases} \widehat{update}_\mathcal{S}(\widehat{\sigma}, d) & \text{if } d \in \mathcal{S} \\ \widehat{update}_{\mathcal{A}\setminus\{\mathbf{X}\}}(\widehat{\sigma}, d) & \text{if } d \in \mathcal{A} \setminus \{\mathbf{X}\} \\ \widehat{update}_\mathbf{X}(\widehat{\sigma}, d) & \text{if } d = \mathbf{X} \end{cases}$$

Similarly to the collecting semantics we define the abstract semantics as the least solution of the following equations:

$$\widehat{R}(v_0) = \emptyset \quad (5)$$

$$\forall v \in V \setminus \{v_0\} : \widehat{R}(v) = \bigsqcup_{(u,d,v) \in E} \widehat{update}(\widehat{R}(u), d) \quad (6)$$

Equations (5) and (6) are the abstract counterpart of Equations (1) and (2). We can now state the main correctness theorem about our analyzer, which follows by standard Abstract Interpretation arguments from Lemmas 1, 2, 3, and 4:

Theorem 1 (Analysis Correctness). *For all $v \in V$, we have:*

$$R^C(v) \subseteq \gamma(\widehat{R}(v))$$

$$\begin{array}{ccccccc} & peel_0 & peel_1 & peel_2 & \cdots & peel_{15} & unroll_0 & unroll_1 \\ \text{Contexts:} & \bullet \rightarrow & \bullet \rightarrow & \bullet \cdots\cdots\cdots\cdots & \bullet \rightarrow & \bullet \rightarrow & \bullet \end{array}$$

Contexts: $peel_0 \rightarrow peel_1 \rightarrow peel_2 \cdots\cdots\cdots peel_{15} \rightarrow unroll_0 \rightarrow unroll_1$

Loop iterations: $\{0\}$ $\{1\}$ $\{2\}$ $\cdots$ $\{15\}$ $\{16, 18, \dots\}$ $\{17, 19, \dots\}$

Fig. 4: Peeling and unrolling contexts and their corresponding loop iterations.

VI. LOOP PEELING AND UNROLLING

A common problem that cache analyses by abstract interpretation suffer from is the loss of precision due to joins at the entry of loops. Indeed, the memory blocks loaded before a loop and within a loop usually differ. As a consequence the abstract cache states entering the loop and upon back edges from within the loop often have few, if any, memory blocks in common. A sound analysis can thus not conclude any blocks to be cached at the beginning of the loop body. One can avoid this issue by *loop peeling*, where the analysis distinguishes the first few iterations of the loop from the rest of the loop and maintains separate analysis information for each of these iterations. This allows the analysis to capture the "warm-up effect" commonly observed in loops iterating across arrays. The example in Figure 4 shows a loop for which the first 16 loop iterations are peeled, which is the optimal amount of peeling for our example from Section III.

Another problem that the basic analysis described in Section V suffers from is the lack of alignment information when establishing the alias relations between MCRs. For example, one cannot tell whether $A[i]$ and $A[i+1]$ map to the same block if no information about the alignment of $A[i]$ is available. Indeed, it can happen that $A[i]$ and $A[i+1]$ are separated by a block boundary when $A[i] \bmod BS = BS - 1$. The necessary alignment information can be obtained by *unrolling loops*, i.e. distinguishing consecutive loop iterations from each other. In the example in Figure 4 the loop is unrolled twice, distinguishing even from odd loop iterations. In our example from Section III we assumed a block size of 8 bytes and array cells of size 4 bytes. Provided knowledge about the base address of the array A, with loop unrolling, the alignment of accesses to $A[i]$ is fully determined.

A. Context-Sensitive Analysis

Given peeling and unrolling depths $MaxPeel \geq 0$ and $MaxUnroll > 0$, we define the following set of tags:

$$Tags := \{peel_x \mid 0 \leq x < MaxPeel\} \cup$$
$$\{unroll_x \mid 0 \leq x < MaxUnroll\}$$

These correspond to the nodes in the graph in Figure 4. We then define contexts as functions that associate a tag with each loop variable, i.e., $Ctxts = LoopVar \rightarrow Tags$. Then, $peel_x$ means that the loop variable has value x, and $unroll_x$ means that value of the loop variable is in $\{MaxPeel + MaxUnroll \cdot n + x \mid n \in \mathbb{N}\}$.

To avoid the precision loss at joins we lift our abstract domain to a context-sensitive domain $\widehat{SymCaches}$ that associates a symbolic cache state with each context:

$$\widehat{SymCaches} = Ctxts \hookrightarrow \widehat{SymCache}$$

These abstract states are updated as follows upon statements:

$$\widehat{update}_\mathcal{S}(\widehat{\sigma}, entry_i) := \lambda ctx \in Ctxts.$$
$$\begin{cases} \bigsqcup_{t \in Tags} \widehat{update}_\mathcal{S}(\widehat{\sigma}(ctx[i \mapsto t]), entry_i) & \text{if } ctx(i) = peel_0 \\ \bot & \text{otherwise} \end{cases}$$

Entering loop i corresponds to setting the loop variable i to zero. Thus, independently, of the previous tag for i, the new tag for i will be $peel_0$. The abstract value for this context is obtained by merging the values of all predecessor contexts, where i may be arbitrary (first case). Contexts in which the tag for i is not $peel_0$ are unreachable via entry edges (second case).

To define the update upon back edges we first capture the structure of the graph in Figure 4 via its set of edges $\mathcal{E}$:

$$\mathcal{E} := \{(peel_x, peel_{x+1}) \mid 0 \leq x < MaxPeel - 1\}$$
$$\cup \{(peel_{MaxPeel-1}, unroll_0)\}$$
$$\cup \{(unroll_x, unroll_{x+1}) \mid 0 \leq x < MaxUnroll - 1\}$$
$$\cup \{(unroll_{MaxUnroll-1}, unroll_0)\}$$

The set $\mathcal{E}$ captures how contexts evolve when taking back edges. Based on $\mathcal{E}$ we define $\widehat{update}_\mathcal{S}(\widehat{\sigma}, backedge_i)$:

$$\widehat{update}_\mathcal{S}(\widehat{\sigma}, backedge_i) := \lambda ctx \in Ctxts.$$
$$\bigsqcup_{\substack{ctx(i)=t' \\ (t,t') \in \mathcal{E}}} \widehat{update}_\mathcal{S}(\widehat{\sigma}(ctx[i \mapsto t]), backedge_i)$$

Assume statements and memory accesses do not modify loop variables. Thus, the update is simply applied pointwise to each context.

B. Refining Alias Relations using Context Information

Contexts provide information about the values of loop variables, which can be used to deduce the alignment of MCRs. To do so, we rely on an auxiliary function $eval_{\mathrm{mod}}(e, ctx)$ that partially evaluates an MCR e in context ctx obtaining one of the following results:

- $Exact(n)$, if the MCR is known to be exactly equal to n in context ctx.
- $Mod(n, p)$, if the MCR is known to be equal to n modulo p in context ctx.
- $Unknown$ if no such statement can be deduced.

We omit $eval_{mod}$ here for brevity; its definition is provided in the appendix.

Using $eval_{mod}$, we can refine the $alias$ function and use the context to deduce alignment relations. Given two MCRs e_1 and e_2, and a context ctx, we refine $alias$ as follows:

$$alias(e_1, e_2, ctx) :=$$

$$
\begin{cases}
sb & \text{if } n = e_1 - e_2 \in \mathbb{Z} \wedge a_1 \sqsubseteq Mod(n_1, BS) \\
& \quad \wedge \; a_2 \sqsubseteq Mod(n_2, BS) \wedge n - n_1 + n_2 = 0 \\
ss & \text{if } n = e_1 - e_2 \in \mathbb{Z} \wedge a_1 \sqsubseteq Mod(n_1, BS) \\
& \quad \wedge \; a_2 \sqsubseteq Mod(n_2, BS) \\
& \quad \wedge \; n - n_1 + n_2 \bmod NS \cdot BS = 0 \\
ds & \text{if } n = e_1 - e_2 \in \mathbb{Z} \wedge a_1 \sqsubseteq Mod(n_1, BS) \\
& \quad \wedge \; a_2 \sqsubseteq Mod(n_2, BS) \\
& \quad \wedge \; n - n_1 + n_2 \bmod NS \cdot BS \neq 0 \\
alias(e_1, e_2) & \text{otherwise}
\end{cases}
$$

where $a_1 = eval_{mod}(e_1, ctx)$ and $a_2 = eval_{mod}(e_2, ctx)$, $Exact(k) \sqsubseteq Mod(n, m)$ if $k = n \bmod m$, and $Mod(n', m') \sqsubseteq Mod(n, m)$ if $m | m'$ and $n = n' \bmod m$.

This refined alias function first looks at the difference $e_1 - e_2$ just like the non-refined version, except that the conditions to derive some relations are relaxed if the alignments (a_1 and a_2) of e_1 and e_2 are known.

In the first case, n_1 and n_2 are the offsets of e_1 and e_2 in their respective blocks. Thus, one can deduce the address of the block that e_1 maps to ($e_1 - a_1$), and compare it to the address of the block that e_2 maps to ($e_2 - a_2$). The equality of block addresses can be rewritten $n - n_1 + n_2 = 0$. If the equality holds, then e_1 and e_2 map to the same block.

The second case is similar, but we check an equality on cache sets instead of blocks. We thus consider alignments relative to sets, by evaluating e_1 and e_2 modulo $NS \cdot BS$. The equality is also checked modulo the same value because addresses that are $NS \cdot BS$ apart map to the same set.

The third case is analogous, except we check for expressions mapping to different sets instead of the same one. Finally, in cases were $eval_{mod}$ fails to evaluate e_1 and e_2 precisely, we rely on the version of $alias$ from Section V-B as a fallback.

VII. Implementation

We implemented the symbolic analysis in LLVMTA [26], [27], [28], a WCET analysis tool based on the LLVM compiler infrastructure. In particular, LLVMTA relies on LLVM to compile the program, which itself uses ScalarEvolution [6], [7] to perform optimizations. It was thus convenient to reuse this framework and convert ScalarEvolution expressions to our own MCR representation upon which we added support for the shifting and substitution operations. The main difficulty arising when converting ScalarEvolution expressions to MCRs is that ScalarEvolution (SCEV) expressions do not only contain integer constants but also LLVM values that belong to the LLVM intermediate representation (IR). Consider an array A that is allocated on the stack in a function f and then passed down to another function g accessing $A[i]$. A SCEV expression for such

an access would typically look like $\{\%A, +, 4\}_i$, where $\%A$ is a parameter of f. We rely on debug information to determine the register containing the value of $\%A$, and then query a dedicated constant value analysis to get the register value. This allows us to translate information available at the IR level down to the machine-code level at which our analysis is performed.

Several tricks are implemented to make the analysis more efficient. First, we rely on hash consing (https://en.wikipedia.org/wiki/Hash_consing) of MCRs to reduce the memory footprint of the analysis: when building an MCR, we check if it was already build before, and return a pointer on the old MCR when possible. In addition to saving memory, this allows us to cache and reuse the results of all operations involving MCRs.

Another trick to speed up the analysis is to avoid representing a symbolic cache state $\widehat{\sigma} \in \widehat{SymCache}$ as a single map of MCRs to ages. Instead, a cache state is split into several maps, which we called "virtual sets". We use one virtual set per physical cache set to store expressions that are known to map to this cache set. An additional virtual set is used for expressions whose corresponding cache set is unknown. When looking for "same block" MCRs (e.g. in $\widehat{update}_{\mathcal{A}\backslash\{\mathbf{X}\}}$), MCRs that map to a different virtual set than the accessed MCR can be excluded from the check, saving time. Virtual sets can also be shared between abstract states. Upon a memory access, if the set to which the accessed MCR is known, only the corresponding virtual set is modified. The remaining virtual sets can thus be shared between the old and the new abstract state, saving memory and avoiding copies.

Regarding the values of $MaxPeel$ and $MaxUnroll$, it is not possible to choose fixed values that would work well for every benchmark due to the presence of nested loops. For example, it is possible to peel the first 256 iterations of a single loop, but doing so for each loop of a loop nest of depth 3 would lead to the creation of 256^3 different contexts, blowing up the analysis complexity. We thus introduce the notion of a *peeling budget* in the analysis, which indicates the number of peeling contexts to create per loop nest. This budget is first spent on the innermost loop, then on the second innermost loop if it is possible to fully peel the innermost one, and so on. For example, consider a loop nest of depth 2, with loop bounds of 20 and 50 for the outer and inner loops, respectively. A peeling budget of 200 would lead to fully peeling the inner loop, because the loop bound of the inner loop is less than the current budget. Then the budget remaining for the outer loop would be 200/50, leading to a $MaxPeel$ value of 4 for the outer loop. We could introduce a similar notion for computing the $MaxUnroll$ value associated to each loop. Because this seemed unnecessary in many benchmarks, we chose to only unroll the innermost loop.

VIII. Experimental Evaluation

The aim of our experiments is to evaluate the following three aspects of our contributions:

1) The gain in accuracy obtained by performing cache analysis over a symbolic CFG.
2) Scalability when increasing the dataset sizes.

3) Scalability in terms of the cache geometry.

First, we demonstrate the properties of our analysis on the illustrative example from Section III. Then, we present experiments designed to assess the accuracy gain due to the symbolic approach and its scalability. All experiments are performed assuming a set-associative cache consisting of 8 cache sets, 8 cache ways, and cache lines of 64 bytes. We qualitatively contrast our work with other related work in Section IX.

In this evaluation we use the PolyBench [29] benchmarks. PolyBench has the advantage of providing a parametric dataset size, i.e. one can adapt the sizes of the data structures the algorithms iterate over. PolyBench provides 5 datasets size: *mini*, *small*, *medium*, *large*, and *extra large*, which is convenient to assess the scalability of our approach.

A. Behavior of the Symbolic Analysis on Illustrative Example

To verify that the symbolic analysis is behaving as expected, we analyze multiple variants of the program in Figure 1a from Section III. In all experiments, we use an array of $12 \cdot 1024 = 12288$ integers, but we vary the number of loop iterations in both loops between 4 and 12288, iterating back and forth across prefixes of the array. We then compare the following analyses:

- The symbolic analysis in optimal settings: we peel the exact number of iterations (1024) required to fill the cache, and we unroll enough iterations (128) to obtain perfect cache alignment information.
- Ferdinand's must analysis [10], [11] under the same settings, i.e. using the same $MaxPeel$ and $MaxUnroll$ values.
- Ferdinand's analysis where both loops are fully peeled.

In each of these analyses we configure LLVMTA to compute a bound on the number of cache misses.

We use Ferdinand's analysis as a baseline, as the symbolic analysis can be seen as a lifted version of Ferdinand's analysis to symbolic CFGs, and thus the observed differences can be directly attributed to operating symbolically.

Figure 5a shows the number of predicted misses when increasing the loop bounds. As expected, for low values of the loop bounds all analyses fully peel the loops, and achieve the same perfect results: The first loop incurs one miss in every 16 iterations, as 16 consecutive integers of 4 bytes fit in a 64-byte cache line. The second loop does not lead to any additional misses because the accessed data fits entirely in the cache. Once the loop bounds are big enough to fill the cache, to the right of the dashed vertical line, the predicted number of misses increases by 2 for every 16 loop iterations for both the symbolic analysis and Ferdinand's analysis if the loops are fully peeled. This is due to additional misses at the end of the second loop, which accesses blocks that were evicted at the end of the first loop. Indeed, the results of the symbolic analysis and of Ferdinand's analysis under full peeling are exact.

However, when the loop bounds exceed the number of peeled iterations, Ferdinand's analysis is unable to classify any access as a hit anymore. As a consequence, the bound on the number of potential misses increases with every access: spatial locality is not exploited because the analysis does not know the offset of the accesses inside a cache line.

Figure 5b shows the analysis runtime of the three analyses in terms of the loop bounds. Once the loop bounds exceed the $MaxPeel$ value, the analysis cost remains constant. Conversely, when increasing the value of $MaxPeel$ to match the loop bound, Ferdinand's analysis gets more and more expensive, quickly exceeding the cost of the symbolic analysis.

B. Accuracy of the Symbolic Analysis

In order to evaluate the benefits of the symbolic analysis in more realistic cases, we analyze the PolyBench benchmarks (with the default dataset size *large*), and compare its accuracy with Ferdinand's analysis. The cache configuration is fixed, but we vary the values of $MaxPeel$ and $MaxUnroll$. Indeed, both analyses perform very differently in terms of running time and accuracy when varying the peeling and unrolling settings, and comparing the two for a fixed setting would thus be difficult. So we set a runtime limit of one hour per benchmark and retain for each analysis the best achievable result within this time for each benchmark. Figure 6 shows that in these conditions, the symbolic analysis always outperforms Ferdinand's analysis. The geometric mean of the ratios of the bounds computed by the symbolic and non-symbolic analysis across all benchmarks is 0.335, significantly improving analysis accuracy.

C. Scalability Evaluation

We claim that the symbolic analysis runtime is largely independent of the number of loop iterations, as long as the number of loop iterations exceeds the number of peeled iterations. To support this claim, we ran the analysis using the same cache configuration and peeling/unrolling settings ($MaxPeel = 1024$, $MaxUnroll = 128$) for all the dataset sizes available in PolyBench. Figure 7 shows the analysis runtime for each benchmark and dataset size. Notice that the dataset size has a smaller impact on the analysis runtime than the benchmark itself, which suggests that the complexity of a benchmark's access patterns is more important than the number of accesses generated by the benchmark. As expected, analysis times for the *large* and *extra large* datasets are usually very close to each other even though the number of memory accesses in the XL case is 6.25 times higher on the average. For the smaller dataset sizes the loop bounds often do not reach the peeling settings, and thus the analysis cost still increases moving from XS to S, and sometimes also from S to M and L.

D. Impact of the Cache Geometry

To evaluate the impact of the cache geometry on the analysis runtime we designed two experiments.

In the first experiment, we investigate the impact of the associativity on the analysis runtime. We fix the cache line size to 64 bytes and the number of cache sets to 8, as in the previous experiments, and analyze associativities 8, 16, 32, and 64, corresponding to cache sizes of 4, 8, 16, and 32 KB, respectively. We run the symbolic analysis on all benchmarks of PolyBench for the *large* dataset. To enable the analysis

(a) Accuracy comparison when increasing the dataset's size.

(b) Analysis time comparison when increasing the dataset's size.

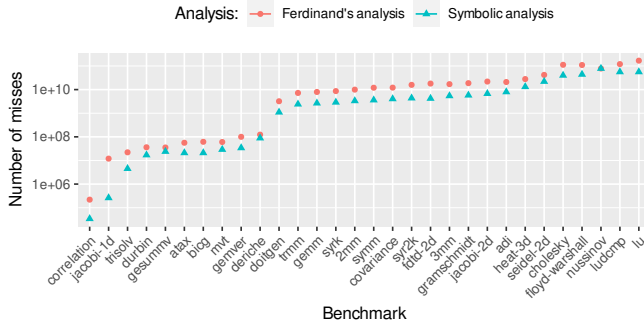Fig. 5: Accuracy and analysis time comparison on the running example.

Fig. 6: Accuracy comparison under a time constraint of 1 hour.

Fig. 8: Geometric mean of slowdowns relative to an analysis with associativity 8 across PolyBench for the *large* dataset.
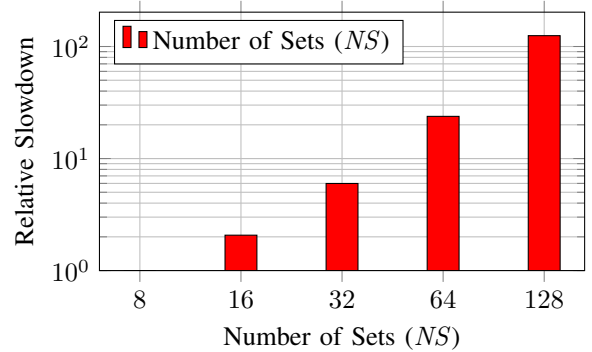

Fig. 7: Analysis runtimes for increasing dataset sizes.

Fig. 9: Geometric mean of slowdowns relative to an analysis with 8 cache sets across PolyBench for the *large* dataset.

to exploit the increased cache size, we double the peeling budget each time we double the associativity. Figure 8 shows the geometric mean of the slowdowns relative to an analysis with associativity 8. We observe a slowdown of 2.56, 10.7, and 70 at associativity 16, 32, and 64, respectively.

In the second experiment, we investigate the impact of the number of cache sets on the analysis runtime. Thus, we fix the cache line size to 64 bytes and the associativity to 8, and perform analyses for 8, 16, 32, 64, and 128 cache sets, corresponding to cache sizes of 4, 8, 16, 32 and 64 KB, respectively. Again, we double the peeling budget each time we double the number of cache lines. Figure 9 shows the geometric mean of the slowdowns relative to an analysis with 8 cache sets. We observe a slowdown of 2.07, 5.99, 23.8, and 125 at 16, 32, 64, and 128 cache sets, respectively.

In both experiments, we observe that the analysis runtime increases superlinearly with the cache size. Indeed, there are two effects at play here that are each individually expected to induce a linear slowdown: (i) the peeling budget is proportional to the cache size and thus the number of contexts increases linearly, and (ii) the abstract cache states grow linearly in the cache size. The effect of (ii) on the analysis runtimes is less pronounced when increasing the number of cache sets than when increasing the associativity due to the use of virtual sets, and we observe smaller slowdowns there.

IX. Related Work

Static cache analysis has received considerable attention in the context of WCET analysis. In the following, we focus on work targeted at data cache analysis. For a broader review of the literature consider the survey paper by Lv et al. [30].

At a high level, work on static cache analysis can be partitioned into classifying and bounding analyses:

- *Classifying analyses* [10], [11], [31], [32], [33], [34], [35], [13], [14], [36], [37] classify individual accesses in the program as hits or misses. Ferdinand's may and must analysis and our symbolic analysis fall into this class.

- *Bounding analyses* [38], [39], [11], [40], [41], [42], [19], [43] compute bounds on the number of misses that occur in a program fragment or in a subset of the program's accesses.

Let us first discuss related classifying analyses. We have already extensively discussed Ferdinand's LRU must analysis [10], [11] throughout the paper. It relies on a plain CFG abstraction, and precise analysis results for data caches are only possible if loops are fully unrolled.

Sen and Srikant [31] build upon LRU must analysis and make two contributions: (i) They introduce a new domain to analyze the set of memory addresses associated with a static memory reference called *circular linear progressions*. (ii) They introduce a new approach to context-sensitive analysis in which a loop is partitioned into n same-length regions that are further split into two parts. The first part is analyzed in "expansion mode", meaning that it is fully virtually unrolled, distinguishing all individual iterations, while the second part is analyzed in "summary mode". To achieve accurate results, the approach requires an unrolling value that is proportional to the number of loop iterations, similarly to Ferdinand's analysis.

Hahn and Grund [25], [44] introduce *relational cache analysis*, which tracks relations between memory accesses in the program following the lattice in Figure 3 similarly to our analysis. Wegener [23] proposes to judiciously apply loop peeling and unrolling to relational cache analysis. Their work is able to detect the exploitation of spatial and temporal locality within a given loop iteration (or within a sequence of loop iterations in case of unrolling). The fundamental limitation of [25], [44], [23] that our approach overcomes, is that their analysis never tracks more than a single symbol for each static memory reference (per unrolled iteration of the loop) in the program, whereas our analysis may dynamically generate an unbounded number of symbols for the same static reference due to the shifting operation upon loop back edges. As a consequence, in our example program, the temporal locality in the second loop would be entirely missed by relational cache analysis. The other major difference lies in our use of LLVM's ScalarEvolution framework to determine access expressions and loop bounds.

Let us now turn to bounding analyses. Kim et al. [38] determine a bound on the number of memory blocks accessed in a program. If at most m distinct blocks are accessed, and these fully fit into the cache, then at most m misses may occur. Such a cache persistence [19] argument only works in cases where the amount of accessed data is smaller than the cache itself, which is often not the case, e.g. in our illustrative example and in the entire PolyBench suite for larger dataset sizes.

Huynh et al. [40] present a persistence analysis that takes a different perspective, separately considering each memory block accessed in the program. For each such block, the analysis determines whether it is persistent, i.e., whether accesses to that block can result in more than one miss. This persistence classification is furthermore performed at different spatial and temporal scopes, e.g. distinguishing different intervals of loop iterations. As a result the analysis may be highly accurate. However the analysis complexity is at least linear in *both* the number of distinct memory blocks accessed by the program and the dynamic number of accesses performed ($> 10^{11}$ for several PolyBench benchmarks for the XL dataset), whereas our analysis is independent of both of these.

The approach of Sotin et al. [43] consists in encoding the program semantics and the cache replacement policy in a formula whose integral solutions correspond to cache misses, and to discharge this counting problem to an external solver [45]. The approach is however limited to counting misses associated to a single static memory reference inside a loop. Ad hoc extensions handling non-linear accesses, several accesses in the same loop, and analyzing nested loops are suggested, but it is not clear whether these approaches can be combined together to handle larger classes of programs.

Finally, there is a long and rich history of analytical cache models [46], [47], [48], [49], [50], [51], [52], [53], [54], [55] that determine the exact number of misses generated by loop nests. A common limitation of this line of work is that it cannot handle programs with input-dependent branches or memory accesses.

X. Conclusions and Future Work

We have introduced *symbolic data cache analysis* a novel analysis that systematically exploits a richer program abstraction than prior work, namely *symbolic control-flow graphs*, which can be obtained from LLVM's ScalarEvolution analysis. The experimental evaluation demonstrates that this new analysis outperforms classical LRU must analysis both in terms of accuracy and analysis runtime.

As a proof of concept, we have lifted the classical LRU must analysis to the symbolic level. Other existing analyses operating on plain CFGs could similarly be made symbolic, e.g. persistence analyses or classifying analyses for various replacement policies. It would also be interesting to investigate whether exact cache analysis on symbolic CFGs is possible along the lines of recent exact cache analyses on plain CFGs.

Another direction for future work is to apply the idea of symbolic cache analysis to even richer program abstractions, e.g. modeling operations on heap data structures.

REFERENCES

[1] R. Wilhelm, J. Engblom, A. Ermedahl, N. Holsti, S. Thesing, D. B. Whalley, G. Bernat, C. Ferdinand, R. Heckmann, T. Mitra, F. Mueller, I. Puaut, P. P. Puschner, J. Staschulat, and P. Stenström, "The worst-case execution-time problem - overview of methods and survey of tools," *ACM Trans. Embed. Comput. Syst.*, vol. 7, no. 3, pp. 36:1–36:53, 2008. [Online]. Available: https://doi.org/10.1145/1347375.1347389

[2] D. J. Bernstein, "Cache-timing attacks on AES," 2005. [Online]. Available: https://cr.yp.to/antiforgery/cachetiming-20050414.pdf

[3] G. Doychev, B. Köpf, L. Mauborgne, and J. Reineke, "CacheAudit: A tool for the static analysis of cache side channels," *ACM Trans. Inf. Syst. Secur.*, vol. 18, no. 1, pp. 4:1–4:32, Jun. 2015. [Online]. Available: http://doi.acm.org/10.1145/2756550

[4] S. Wang, P. Wang, X. Liu, D. Zhang, and D. Wu, "Cached: Identifying cache-based timing channels in production software," in *26th USENIX Security Symposium, USENIX Security 2017, Vancouver, BC, Canada, August 16-18, 2017*, E. Kirda and T. Ristenpart, Eds. USENIX Association, 2017, pp. 235–252. [Online]. Available: https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/wang-shuai

[5] C. Sung, B. Paulsen, and C. Wang, "CANAL: a cache timing analysis framework via LLVM transformation," in *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering, ASE 2018, Montpellier, France, September 3-7, 2018*, M. Huchard, C. Kästner, and G. Fraser, Eds. ACM, 2018, pp. 904–907. [Online]. Available: https://doi.org/10.1145/3238147.3240485

[6] LLVM, "ScalarEvolution." [Online]. Available: https://github.com/llvm/llvm-project/blob/main/llvm/include/llvm/Analysis/ScalarEvolution.h

[7] J. Absar, "Scalar Evolution - Demystified," 2018, European LLVM Developers Meeting. [Online]. Available: https://llvm.org/devmtg/2018-04/slides/Absar-ScalarEvolution.pdf

[8] O. Bachmann, P. S. Wang, and E. V. Zima, "Chains of recurrences - a method to expedite the evaluation of closed-form functions," in *Proceedings of the International Symposium on Symbolic and Algebraic Computation, ISSAC '94, Oxford, UK, July 20-22, 1994*, M. A. H. MacCallum, Ed. ACM, 1994, pp. 242–249. [Online]. Available: https://doi.org/10.1145/190347.190423

[9] R. van Engelen, "Efficient symbolic analysis for optimizing compilers," in *Compiler Construction, 10th International Conference, CC 2001 Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2001 Genova, Italy, April 2-6, 2001, Proceedings*, ser. Lecture Notes in Computer Science, R. Wilhelm, Ed., vol. 2027. Springer, 2001, pp. 118–132. [Online]. Available: https://doi.org/10.1007/3-540-45306-7_9

[10] C. Ferdinand, "Cache behavior prediction for real-time systems," Ph.D. dissertation, Saarland University, Saarbrücken, Germany, 1997, iSBN: 3-9307140-31-0. [Online]. Available: https://d-nb.info/953983706

[11] C. Ferdinand and R. Wilhelm, "Efficient and precise cache behavior prediction for real-time systems," *Real-Time Systems*, vol. 17, no. 2-3, pp. 131–181, Nov. 1999. [Online]. Available: https://doi.org/10.1023/A:1008186323068

[12] J. Reineke, D. Grund, C. Berg, and R. Wilhelm, "Timing predictability of cache replacement policies," *Real-Time Systems*, vol. 37, no. 2, pp. 99–122, Nov. 2007. [Online]. Available: https://doi.org/10.1007/s11241-007-9032-3

[13] V. Touzeau, C. Maïza, D. Monniaux, and J. Reineke, "Ascertaining uncertainty for efficient exact cache analysis," in *Computer Aided Verification - 29th International Conference, CAV 2017, Heidelberg, Germany, July 24-28, 2017, Proceedings, Part II*, ser. Lecture Notes in Computer Science, R. Majumdar and V. Kuncak, Eds., vol. 10427. Springer, 2017, pp. 22–40. [Online]. Available: https://doi.org/10.1007/978-3-319-63390-9_2

[14] ——, "Fast and exact analysis for LRU caches," *Proc. ACM Program. Lang.*, vol. 3, no. POPL, pp. 54:1–54:29, 2019. [Online]. Available: https://doi.org/10.1145/3290367

[15] G. Stock, S. Hahn, and J. Reineke, "Cache persistence analysis: Finally exact," in *IEEE Real-Time Systems Symposium, RTSS 2019, Hong Kong, SAR, China, December 3-6, 2019*. IEEE, 2019, pp. 481–494. [Online]. Available: https://doi.org/10.1109/RTSS46320.2019.00049

[16] F. Mueller, "Timing analysis for instruction caches," *Real-Time Systems*, vol. 18, no. 2, pp. 217–247, May 2000. [Online]. Available: https://doi.org/10.1023/A:1008145215849

[17] C. Cullmann, "Cache persistence analysis: Theory and practice," *ACM Trans. Embedded Comput. Syst.*, vol. 12, no. 1s, pp. 40:1–40:25, 2013. [Online]. Available: https://doi.org/10.1145/2435227.2435236

[18] Z. Zhang and X. D. Koutsoukos, "Improving the precision of abstract interpretation based cache persistence analysis," in *Proceedings of the 16th ACM SIGPLAN/SIGBED Conference on Languages, Compilers and Tools for Embedded Systems, Portland, OR, USA, June 18-19, 2015*, ser. LCTES 2015, 2015, pp. 10:1–10:10. [Online]. Available: https://doi.org/10.1145/2670529.2754967

[19] J. Reineke, "The semantic foundations and a landscape of cache-persistence analyses," *Leibniz Trans. Embed. Syst.*, vol. 5, no. 1, pp. 03:1–03:52, 2018. [Online]. Available: https://doi.org/10.4230/LITES-v005-i001-a003

[20] R. van Engelen, "Symbolic evaluation of chains of recurrences for loop optimization," Computer Science Dept., Florida State University, Tech. Rep. TR-000102, 2000.

[21] F. Martin, M. H. Alt, R. Wilhelm, and C. Ferdinand, "Analysis of loops," in *Compiler Construction, 7th International Conference, CC'98, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS'98, Lisbon, Portugal, March 28 - April 4, 1998, Proceedings*, ser. Lecture Notes in Computer Science, K. Koskimies, Ed., vol. 1383. Springer, 1998, pp. 80–94. [Online]. Available: https://doi.org/10.1007/BFb0026424

[22] L. Mauborgne and X. Rival, "Trace partitioning in abstract interpretation based static analyzers," in *Programming Languages and Systems, 14th European Symposium on Programming,ESOP 2005, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2005, Edinburgh, UK, April 4-8, 2005, Proceedings*, 2005, pp. 5–20. [Online]. Available: https://doi.org/10.1007/978-3-540-31987-0_2

[23] S. Wegener, "Computing same block relations for relational cache analysis," in *12th International Workshop on Worst-Case Execution Time Analysis, WCET 2012, July 10, 2012, Pisa, Italy*, ser. OASIcs, T. Vardanega, Ed., vol. 23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2012, pp. 25–37. [Online]. Available: https://doi.org/10.4230/OASIcs.WCET.2012.25

[24] D. A. Schmidt, "Trace-based abstract interpretation of operational semantics," *LISP Symb. Comput.*, vol. 10, no. 3, pp. 237–271, 1998.

[25] S. Hahn, "Towards relational cache analysis," Bachelor's Thesis, Saarland University, 2011. [Online]. Available: http://embedded.cs.uni-sb.de/publications/RelCanaBSC2011.pdf

[26] ——, "On static execution-time analysis," PhD Thesis, Saarland University, 2018. [Online]. Available: https://publikationen.sulb.uni-saarland.de/bitstream/20.500.11880/27440/1/dissertation.pdf

[27] S. Hahn, M. Jacobs, N. Hölscher, K. Chen, J. Chen, and J. Reineke, "LLVMTA: an LLVM-based WCET analysis tool," in *20th International Workshop on Worst-Case Execution Time Analysis, WCET 2022, July 5, 2022, Modena, Italy*, ser. OASIcs, C. Ballabriga, Ed., vol. 103. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, pp. 2:1–2:17. [Online]. Available: https://doi.org/10.4230/OASIcs.WCET.2022.2

[28] ——, "LLVMTA: an LLVM-based WCET analysis tool," https://gitlab.cs.uni-saarland.de/reineke/llvmta, 2022.

[29] L.-N. Pouchet, U. Bondugula, and T. Yuki, "Polybench v4.2.1." [Online]. Available: https://sourceforge.net/projects/polybench/

[30] M. Lv, N. Guan, J. Reineke, R. Wilhelm, and W. Yi, "A survey on static cache analysis for real-time systems," *Leibniz Trans. Embed. Syst.*, vol. 3, no. 1, pp. 05:1–05:48, 2016. [Online]. Available: https://doi.org/10.4230/LITES-v003-i001-a005

[31] R. Sen and Y. N. Srikant, "WCET estimation for executables in the presence of data caches," in *Proceedings of the 7th ACM & IEEE International conference on Embedded software, EMSOFT 2007, September 30 - October 3, 2007, Salzburg, Austria*, C. M. Kirsch and R. Wilhelm, Eds. ACM, 2007, pp. 203–212. [Online]. Available: https://doi.org/10.1145/1289927.1289960

[32] D. Grund and J. Reineke, "Abstract interpretation of FIFO replacement," in *Static Analysis, 16th International Symposium, SAS 2009, Los Angeles, CA, USA, August 9-11, 2009. Proceedings*, ser. Lecture Notes in Computer Science, J. Palsberg and Z. Su, Eds., vol. 5673. Springer, 2009, pp. 120–136. [Online]. Available: https://doi.org/10.1007/978-3-642-03237-0_10

[33] ——, "Precise and efficient FIFO-replacement analysis based on static phase detection," in *22nd Euromicro Conference on Real-Time Systems, ECRTS 2010, Brussels, Belgium, July 6-9, 2010*. IEEE Computer Society, 2010, pp. 155–164. [Online]. Available: https://doi.org/10.1109/ECRTS.2010.8

[34] ——, "Toward precise PLRU cache analysis," in *10th International Workshop on Worst-Case Execution Time Analysis, WCET 2010, July 6, 2010, Brussels, Belgium*, ser. OASICS, B. Lisper, Ed., vol. 15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany, 2010, pp. 23–35. [Online]. Available: https://doi.org/10.4230/OASIcs.WCET.2010.23

[35] S. Chattopadhyay and A. Roychoudhury, "Scalable and precise refinement of cache timing analysis via path-sensitive verification," *Real Time Syst.*, vol. 49, no. 4, pp. 517–562, 2013. [Online]. Available: https://doi.org/10.1007/s11241-013-9178-0

[36] F. Brandner and C. Noûs, "Precise and efficient analysis of context-sensitive cache conflict sets," in *28th International Conference on Real Time Networks and Systems, RTNS 2020, Paris, France, June 10, 2020*, L. Cucu-Grosjean, R. Medina, S. Altmeyer, and J. Scharbarg, Eds. ACM, 2020, pp. 44–55. [Online]. Available: https://doi.org/10.1145/3394810.3394811

[37] ——, "Precise, efficient, and context-sensitive cache analysis," *Real Time Syst.*, vol. 58, no. 1, pp. 36–84, 2022. [Online]. Available: https://doi.org/10.1007/s11241-021-09372-5

[38] S. Kim, S. L. Min, and R. Ha, "Efficient worst case timing analysis of data caching," in *2nd IEEE Real-Time Technology and Applications Symposium, RTAS '96, Boston, MA, USA, June 10-12, 1996*. IEEE Computer Society, 1996, pp. 230–240. [Online]. Available: https://doi.org/10.1109/RTTAS.1996.509540

[39] R. T. White, C. A. Healy, D. B. Whalley, F. Mueller, and M. G. Harmon, "Timing analysis for data caches and set-associative caches," in *3rd IEEE Real-Time Technology and Applications Symposium, RTAS '97, Montreal, Canada, June 9-11, 1997*. IEEE Computer Society, 1997, pp. 192–202. [Online]. Available: https://doi.org/10.1109/RTTAS.1997.601358

[40] B. K. Huynh, L. Ju, and A. Roychoudhury, "Scope-aware data cache analysis for WCET estimation," in *17th IEEE Real-Time and Embedded Technology and Applications Symposium, RTAS 2011, Chicago, Illinois, USA, 11-14 April 2011*. IEEE Computer Society, 2011, pp. 203–212. [Online]. Available: https://doi.org/10.1109/RTAS.2011.27

[41] N. Guan, X. Yang, M. Lv, and W. Yi, "FIFO cache analysis for WCET estimation: a quantitative approach," in *Design, Automation and Test in Europe, DATE 13, Grenoble, France, March 18-22, 2013*, E. Macii, Ed. EDA Consortium San Jose, CA, USA / ACM DL, 2013, pp. 296–301. [Online]. Available: https://doi.org/10.7873/DATE.2013.073

[42] N. Guan, M. Lv, W. Yi, and G. Yu, "WCET analysis with MRU cache: Challenging LRU for predictability," *ACM Trans. Embed. Comput. Syst.*, vol. 13, no. 4s, pp. 123:1–123:26, Apr. 2014. [Online]. Available: http://doi.acm.org/10.1145/2584655

[43] P. Sotin, Q. Vermande, and H. Cassé, "Data cache analysis by counting integer points," in *RTNS'2021: 29th International Conference on Real-Time Networks and Systems, Nantes, France, April 7-9, 2021*, A. Queudet, I. Bate, and G. Lipari, Eds. ACM, 2021, pp. 112–122. [Online]. Available: https://doi.org/10.1145/3453417.3453424

[44] S. Hahn and D. Grund, "Relational cache analysis for static timing analysis," in *24th Euromicro Conference on Real-Time Systems, ECRTS 2012, Pisa, Italy, July 11-13, 2012*, R. Davis, Ed. IEEE Computer Society, 2012, pp. 102–111. [Online]. Available: https://doi.org/10.1109/ECRTS.2012.14

[45] A. I. Barvinok, "A polynomial time algorithm for counting integral points in polyhedra when the dimension is fixed," in *34th Annual Symposium on Foundations of Computer Science, Palo Alto, California, USA, 3-5 November 1993*. IEEE Computer Society, 1993, pp. 566–572. [Online]. Available: https://doi.org/10.1109/SFCS.1993.366830

[46] S. Ghosh, M. Martonosi, and S. Malik, "Cache miss equations: An analytical representation of cache misses," in *Proceedings of the 11th international conference on Supercomputing, ICS 1997, Vienna, Austria, July 7-11, 1997*, S. J. Wallach and H. P. Zima, Eds. ACM, 1997, pp. 317–324. [Online]. Available: https://doi.org/10.1145/263580.263657

[47] ——, "Cache miss equations: A compiler framework for analyzing and tuning memory behavior," *ACM Trans. Program. Lang. Syst.*, vol. 21, no. 4, p. 703–746, Jul. 1999. [Online]. Available: https://doi.org/10.1145/325478.325479

[48] S. Chatterjee, E. Parker, P. J. Hanlon, and A. R. Lebeck, "Exact analysis of the cache behavior of nested loops," in *Proceedings of the 2001 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Snowbird, Utah, USA, June 20-22, 2001*, M. Burke and M. L. Soffa, Eds. ACM, 2001, pp. 286–297. [Online]. Available: https://doi.org/10.1145/378795.378859

[49] X. Vera and J. Xue, "Let's study whole-program cache behaviour analytically," in *Proceedings of the Eighth International Symposium on High-Performance Computer Architecture (HPCA'02), Boston, Massachusettes, USA, February 2-6, 2002*. IEEE Computer Society, 2002, pp. 175–186. [Online]. Available: https://doi.org/10.1109/HPCA.2002.995708

[50] X. Vera, N. Bermudo, J. Llosa, and A. González, "A fast and accurate framework to analyze and optimize cache memory behavior," *ACM Trans. Program. Lang. Syst.*, vol. 26, no. 2, pp. 263–300, 2004. [Online]. Available: https://doi.org/10.1145/973097.973099

[51] C. Cascaval and D. A. Padua, "Estimating cache misses and locality using stack distances," in *Proceedings of the 17th Annual International Conference on Supercomputing*, ser. ICS '03. New York, NY, USA: Association for Computing Machinery, 2003, p. 150–159. [Online]. Available: https://doi.org/10.1145/782814.782836

[52] K. Beyls and E. H. D'Hollander, "Generating cache hints for improved program efficiency," *J. Syst. Archit.*, vol. 51, no. 4, pp. 223–250, 2005. [Online]. Available: https://doi.org/10.1016/j.sysarc.2004.09.004

[53] W. Bao, S. Krishnamoorthy, L. Pouchet, and P. Sadayappan, "Analytical modeling of cache behavior for affine programs," *Proc. ACM Program. Lang.*, vol. 2, no. POPL, pp. 32:1–32:26, 2018. [Online]. Available: https://doi.org/10.1145/3158120

[54] T. Gysi, T. Grosser, L. Brandner, and T. Hoefler, "A fast analytical model of fully associative caches," in *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2019, Phoenix, AZ, USA, June 22-26, 2019*, K. S. McKinley and K. Fisher, Eds. ACM, 2019, pp. 816–829. [Online]. Available: https://doi.org/10.1145/3314221.3314606

[55] C. Morelli and J. Reineke, "Warping cache simulation of polyhedral programs," in *PLDI '22: 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation, San Diego, CA, USA, June 13 - 17, 2022*, R. Jhala and I. Dillig, Eds. ACM, 2022, pp. 316–331. [Online]. Available: https://doi.org/10.1145/3519939.3523714

APPENDIX A
PROOFS

Lemma 1 (Join Correctness). *For all $\widehat{\sigma}_1, \widehat{\sigma}_2 \in \widehat{SymCache}$:*

$$\gamma(\widehat{\sigma}_1) \cup \gamma(\widehat{\sigma}_2) \subseteq \gamma(\widehat{\sigma}_1 \sqcup \widehat{\sigma}_2)$$

Proof. Let $\widehat{\sigma}_1, \widehat{\sigma}_2 \in \widehat{SymCache}$, and $(\sigma_p, \sigma_c) \in \gamma(\widehat{\sigma}_1) \cup \gamma(\widehat{\sigma}_2)$. We assume without loss of generality that $(\sigma_p, \sigma_c) \in \gamma(\widehat{\sigma}_1)$. By definition of the concretization function, we have:

$$\forall e \in dom(\widehat{\sigma}_1) : \sigma_c(block(\llbracket e \rrbracket_{\sigma_p})) \le \widehat{\sigma}_1(e)$$

Thus, for any $e \in dom(\widehat{\sigma}_1) \cap dom(\widehat{\sigma}_2)$, we have:

$$\sigma_c(block(\llbracket e \rrbracket_{\sigma_p})) \le \widehat{\sigma}_1(e)$$
$$\le \max(\widehat{\sigma}_1(e), \widehat{\sigma}_2(e))$$
$$\le (\widehat{\sigma}_1 \sqcup \widehat{\sigma}_2)(e)$$

Thus, $(\sigma_p, \sigma_c) \in \gamma(\widehat{\sigma}_1 \sqcup \widehat{\sigma}_2)$. $\square$

Lemma 2 (Unknown Access Transformer Correctness). *For all $\widehat{\sigma} \in \widehat{SymCache}$, we have:*

$$update(\gamma(\widehat{\sigma}), \mathbf{X}) \subseteq \gamma(\widehat{update}_{\mathbf{X}}(\widehat{\sigma}))$$

Proof. Let $\widehat{\sigma} \in \widehat{SymCache}$. Let $(\sigma'_p, \sigma'_c) \in update(\gamma(\widehat{\sigma}), \mathbf{X})$.

We will show that $(\sigma'_p, \sigma'_c) \in \gamma(\widehat{update}_{\mathbf{X}}(\widehat{\sigma}))$.

Let $(\sigma_p, \sigma_c) \in \gamma(\widehat{\sigma})$ and $b \in \mathcal{B}$ such that $(\sigma'_p, \sigma'_c) = update((\sigma_p, \sigma_c), b)$. We know that $\forall e \in dom(\widehat{\sigma}) : \sigma_c(block(\llbracket e \rrbracket_{\sigma_p})) \le \widehat{\sigma}(e)$, and we want to prove that $\forall e \in dom(\widehat{\sigma}').\sigma'_c(block(\llbracket e \rrbracket_{\sigma'_p})) \le \widehat{\sigma}'$, where $\widehat{\sigma}' = \widehat{update}_{\mathbf{X}}(\widehat{\sigma})$.

Let $e \in dom(\widehat{\sigma}') = dom(\widehat{\sigma})$. Expanding all definitions, we have $\sigma'_c(block(\llbracket e \rrbracket_{\sigma'_p})) \le \sigma_c(block(\llbracket e \rrbracket_{\sigma_p})) + 1 \le \widehat{\sigma}(e) + 1 \le \widehat{\sigma}'(e)$.

Thus $(\sigma'_p, \sigma'_c) \in \gamma(\widehat{\sigma}')$, finishing the proof. $\square$

Lemma 3 (MCR Access Transformer Correctness). *For all $\widehat{\sigma} \in \widehat{SymCache}$ and $e \in \mathcal{A}$, we have:*

$$update(\gamma(\widehat{\sigma}), e) \subseteq \gamma(\widehat{update}_{\mathcal{A}\backslash\{\mathbf{X}\}}(\widehat{\sigma}, e))$$

Proof. Let $\widehat{\sigma} \in \widehat{SymCache}$, $e \in M(LoopVar)$ and $(\sigma_p, \sigma_c) \in \gamma(\widehat{\sigma})$. We will show that $(\sigma_p, update_{LRU}(\sigma_c, block(\llbracket e \rrbracket_{\sigma_p}))) \in \gamma(\widehat{update}_{\mathcal{A}\backslash\{\mathbf{X}\}}(\widehat{\sigma}, e))$. To ease reading, we introduce the following notations:

- $b = block(\llbracket e \rrbracket_{\sigma_p})$, the block e maps to.
- $\widehat{\sigma}' = \widehat{update}_{\mathcal{A}\backslash\{\mathbf{X}\}}(\widehat{\sigma}, e)$, the successor of $\widehat{\sigma}$ after the update.
- $\sigma'_c = update_{LRU}(\sigma_c, b)$, the successor of σ_c.

We then want to prove that: $(\sigma_p, \sigma'_c) \in \gamma(\widehat{\sigma}')$, i.e. $\forall e' : \sigma'_c(block(\llbracket e' \rrbracket_{\sigma_p})) \le \widehat{\sigma}'(e')$.

Let $e' \in dom(\widehat{\sigma}) \cup \{e\}$. Noting $b' = block(\llbracket e' \rrbracket_{\sigma_p})$, we want to show that $\sigma'_c(b') \le \widehat{\sigma}'(e')$. This is done by reasoning by case distinction, looking at how $\widehat{\sigma}'(e')$ is obtained from $\widehat{\sigma}(e')$.

- Assume $alias(e, e') \sqsubseteq sb$: By correctness of $alias$, we have: $b = b'$, which by definition of $update_{LRU}$ leads to $\sigma'_c(b') = 0$. We thus have $\sigma'_c(b') \le \widehat{\sigma}'(e')$. This case covers in particular the condition $e = e'$. In the remaining cases, we will thus assume that $e \ne e'$, and thus that $e' \in dom(\widehat{\sigma})$.

- Suppose $alias(e, e') \sqsubseteq sb + ds$: Again, by correctness of $alias$, we have either $b = b'$ or $set(b) \ne set(b')$. If $b = b'$, we are back to the previous case and $\sigma'_c(b') = 0 \le \widehat{\sigma}'(e')$. Otherwise, $set(b) \ne set(b')$, and by definition of $update_{LRU}$ we obtain $\sigma'_c(b') = \sigma_c(b')$. However, we have $\sigma_c(b') \le \widehat{\sigma}(e')$ because $(\sigma_p, \sigma_c) \in \gamma(\widehat{\sigma})$ and $e' \in dom(\widehat{\sigma})$. In addition, by case hypothesis, and definition of $\widehat{update}_{\mathcal{A}\backslash\{\mathbf{X}\}}$, we have: $\widehat{\sigma}'(e') = \widehat{\sigma}(e')$. Combining these inequalities together we obtain:

$$\sigma'_c(b') = \sigma_c(b') \le \widehat{\sigma}(e') = \widehat{\sigma}'(e')$$

- Otherwise, assume $\widehat{\sigma}(e) \le \widehat{\sigma}(e')$: By $\widehat{update}_{\mathcal{A}\backslash\{\mathbf{X}\}}$ and γ, we get $\widehat{\sigma}'(e') = \widehat{\sigma}(e') \ge \sigma_c(b')$. We then do one additional case distinction:
  - If $\sigma_c(b) \le \sigma_c(b')$, by $update_{LRU}$ we obtain:

  $$\sigma'_c(b') = \sigma_c(b') \le \widehat{\sigma}(e') = \widehat{\sigma}'(e')$$

  - Else $\sigma_c(b') < \sigma_c(b)$. We thus have:

  $$\sigma'_c(b') \le \sigma_c(b') + 1 \le \sigma_c(b) \le \widehat{\sigma}(e) \le \widehat{\sigma}(e') = \widehat{\sigma}'(e')$$

  Both subcases thus lead to $\sigma'_c(b') \le \widehat{\sigma}'(e')$ as required.

- Otherwise, suppose $\widehat{\sigma}(e') + 1 < k$: By definition of $\widehat{update}_{\mathcal{A}\backslash\{\mathbf{X}\}}$, we have $\widehat{\sigma}'(e') = \widehat{\sigma}(e') + 1$. Considering all cases in $update_{LRU}$, observe that $\sigma'_c(b') \le \sigma_c(b') + 1$, and so

$$\sigma'_c(b') \le \sigma_c(b') + 1 \le \widehat{\sigma}(e') + 1 = \widehat{\sigma}'(e')$$

- Finally, assume none of the conditions above apply: By definition $\widehat{update}_{\mathcal{A}\backslash\{\mathbf{X}\}}$, we have $\widehat{\sigma}'(e') = \infty$ and so trivially $\sigma'_c(b') \le \widehat{\sigma}'(e')$.

This finishes the proof. In all cases, we have $\sigma'_c(b') \le \widehat{\sigma}'(e')$ proving that $(\sigma_p, \sigma'_c) \in \gamma(\widehat{\sigma}')$, and thus:

$$update(\gamma(\widehat{\sigma}), e) \subseteq \gamma(\widehat{update}_{\mathcal{A}\backslash\{\mathbf{X}\}}(\widehat{\sigma}, e))$$

$\square$

Lemma 4 (Statement Transformer Correctness). *For all $\widehat{\sigma} \in \widehat{SymCache}$ and $s \in \mathcal{S}$, we have:*

$$update(\gamma(\widehat{\sigma}), s) \subseteq \gamma(\widehat{update}_{\mathcal{S}}(\widehat{\sigma}, s))$$

Proof. Let $\widehat{\sigma} \in \widehat{SymCache}$, $s \in \mathcal{S}$ and $(\sigma'_p, \sigma'_c) \in update(\gamma(\widehat{\sigma}), s))$. We will show that $(\sigma'_p, \sigma'_c) \in \gamma(\widehat{update}_{\mathcal{S}}(\widehat{\sigma}, s))$. In the remaining, we will note $\widehat{\sigma}' = \widehat{update}(\widehat{\sigma}, s)$.

By definition of $update$, $\sigma'_p \ne \perp_p$ and there exists $(\sigma_p, \sigma_c) \in \gamma(\widehat{\sigma})$ such that $(\sigma'_p, \sigma'_c) = update((\sigma_p, \sigma_c), s)$. From this, we deduce $\sigma'_c = \sigma_c$. and $\sigma'_p = update_{\mathcal{S}}(\sigma_p, s)$.

We want to show that for any $e' \in dom(\widehat{\sigma}')$, $\sigma_c(block(\llbracket e' \rrbracket_{\sigma'_p})) \le \widehat{\sigma}'(e')$.

Let $e' \in dom(\widehat{\sigma}')$. We proceed by case distinction on the value of s:

- If $s = entry_i$ for some i: Because $e' \in dom(\widehat{\sigma}')$, we have $i \notin e'$ and thus $\llbracket e' \rrbracket_{\sigma'_p} = \llbracket e' \rrbracket_{\sigma_p[i \mapsto 0]} = \llbracket e' \rrbracket_{\sigma_p}$. On the other side, we have: $\widehat{\sigma}'(e) = \widehat{\sigma}(e)$. Inserting these

equalities in the definition of $(\sigma_p, \sigma_c) \in \gamma(\widehat{\sigma})$, we obtain the desired inequality: $\sigma_c(block(\llbracket e' \rrbracket_{\sigma'_p})) \leq \widehat{\sigma}'(e')$.

- If $s = backedge_i$ for some i: By definition of $\widehat{update}_{\mathcal{S}}$, there is an e such that $e' = Sh(e, i)$ and $\widehat{\sigma}'(e') = \widehat{\sigma}(e)$. We thus have:

$$\sigma_c(block(\llbracket e' \rrbracket_{\sigma'_p})) = \sigma_c(block(\llbracket Sh(e,i) \rrbracket_{\sigma_p[i \mapsto \sigma_p(i)+1]}))$$
$$= \sigma_c(block(\llbracket e \rrbracket_{\sigma_p}))$$
$$\leq \widehat{\sigma}(e)$$
$$\leq \widehat{\sigma}'(e')$$

- If $s = assume_{i,expr}$ for some i and $expr \in MCR(LoopVar \setminus \{i\})$: By defintion of $\widehat{update}_{\mathcal{S}}$, we know there is an e such that $e' = Sub(e, i, expr) \neq fail$ and $\widehat{\sigma}'(e') = \widehat{\sigma}(e)$. We already deduced that $\sigma'_p \neq \bot_p$, which now implies that $\sigma'_p = \sigma_p$ and $\sigma_p(i) = \llbracket expr \rrbracket_{\sigma_p}$. We thus have:

$$\sigma_c(block(\llbracket e' \rrbracket_{\sigma'_p})) = \sigma_c(block(\llbracket Sub(e, i, expr) \rrbracket_{\sigma_p}))$$
$$= \sigma_c(block(\llbracket e \rrbracket_{\sigma_p[i \mapsto \llbracket expr \rrbracket_{\sigma_p}]}))$$
$$= \sigma_c(block(\llbracket e \rrbracket_{\sigma_p[i \mapsto \sigma_p(i)]}))$$
$$= \sigma_c(block(\llbracket e \rrbracket_{\sigma_p}))$$
$$\leq \widehat{\sigma}(e)$$
$$\leq \widehat{\sigma}'(e')$$

In every case, it thus holds that $\sigma_c(block(\llbracket e' \rrbracket_{\sigma'_p})) \leq \widehat{\sigma}'(e')$, proving that $(\sigma'_p, \sigma'_c) \in \gamma(\widehat{\sigma}')$ as desired. $\qquad\square$

Lemma 5 (Concrete Domain Completeness). *The set $D = \mathcal{P}((LoopVar \to \mathbb{N}) \times (\mathcal{B} \to \mathbb{N}))$, where $\mathcal{B}$ is the set of memory blocks accessed by the program, is a complete lattice.*

Proof. D being a powerset, $\bigcup S$ and $\bigcap S$ are obvious least upper and greatest lower bound for any subset S of D. $\qquad\square$

Lemma 6 (Abstract Domain Completeness). *The set $\widehat{SymCache} = M(LoopVar) \hookrightarrow \{0, \ldots, k-1, \infty\}$ is a complete lattice.*

Proof. Let $S \subseteq \widehat{SymCache}$. Consider $\bigsqcup S = \lambda e \in \bigcap_{\widehat{\sigma} \in S} dom(\widehat{\sigma}). \max(\{\widehat{\sigma}(e), \widehat{\sigma} \in S\})$. $\bigsqcup S$ is well defined even for an infinite subset S because $\{0, \ldots, k-1, \infty\}$ being finite, $\{\widehat{\sigma}(e), \widehat{\sigma} \in S\}$ always admit a maximum. $\bigsqcup S$ belongs to $\widehat{SymCache}$ and the upper-bound property $\forall \widehat{\sigma} \in S : \widehat{\sigma} \sqsubseteq \bigsqcup S$ is obvious. One can show in a similar way that $\bigsqcap S = \lambda e. \min(\{\widehat{\sigma}(e), \widehat{\sigma} \in S\})$ is a lower bound in $\widehat{SymCache}$ for any subset S of $\widehat{SymCache}$, making it a complete lattice. $\qquad\square$

Theorem 1 (Analysis Correctness). *For all $v \in V$, we have:*

$$R^C(v) \subseteq \gamma(\widehat{R}(v))$$

Proof. The semantics R^C and $\widehat{R}$ can be rewritten as the least fixpoints of the following functions:

$$F(R) = \lambda v. R_0 \cup \bigcup_{(u,d,v) \in E} update(R(u), d)$$
$$\widehat{F}(\widehat{R}) = \lambda v. \widehat{R_0}(v) \sqcup \bigsqcup_{(u,d,v) \in E} \widehat{update}(\widehat{R}(u), d)$$

where

$$R_0(v) = \begin{cases} \{(\lambda i.0, \sigma_c) \mid \sigma_c \in \Sigma_c\} & \text{if } v = v_0 \\ \emptyset & \text{otherwise} \end{cases}$$
$$\widehat{R_0}(v_0) = \emptyset$$

This is well-defined because the Knaster-Tarski fixpoint theorem guarantees the existence of these fixpoints. Indeed, F and $\widehat{F}$ are monotone by construction, and Lemmas 5 and 6 ensures that our concrete and abstract domains are complete lattices. In addition, the Knaster-Tarski theorem ensures that $R^C = lfp(F) = \bigcap Red(F) = \bigcap\{R \mid R \supseteq F(R)\}$ (the same relation holds for $\widehat{R}$ and $\widehat{F}$ but we will not need it).

From Lemmas 2 and 3 and 4 about the correctness of statement and access transformers, one can trivially derive that for any $\widehat{\sigma}$:

$$update(\gamma(\widehat{\sigma}, d) \subseteq \gamma(\widehat{update}(\widehat{\sigma}, d))$$

We thus have:

$$F(\gamma(\widehat{R})) = \lambda v. R_0 \cup$$
$$\bigcup_{(u,d,v) \in E} update(\gamma(\widehat{R}(u)), d)$$
$$\subseteq \lambda v. R_0 \cup \bigcup_{(u,d,v) \in E} \gamma(\widehat{update}(\widehat{R}(u), d))$$
$$\subseteq \lambda v. R_0 \cup \gamma\left(\bigsqcup_{(u,d,v) \in E} \widehat{update}(\widehat{R}(u), d)\right)$$
$$\subseteq \lambda v. \gamma(\widehat{R_0}) \cup \gamma\left(\bigsqcup_{(u,d,v) \in E} \widehat{update}(\widehat{R}(u), d)\right)$$
$$\subseteq \lambda v. \gamma\left(\widehat{R_0} \sqcup \bigsqcup_{(u,d,v) \in E} \widehat{update}(\widehat{R}(u), d)\right)$$
$$\subseteq \gamma(\widehat{F}(\widehat{R}))$$

Finally, we get the following inequalities:

$$\widehat{R} = lfp(\widehat{F}) \Rightarrow \widehat{F}(\widehat{R}) = \widehat{R}$$
$$\Rightarrow \gamma(\widehat{F}(\widehat{R})) = \gamma(\widehat{R})$$
$$\Rightarrow F(\gamma(\widehat{R})) \subseteq \gamma(\widehat{R})$$
$$\Rightarrow \gamma(\widehat{R}) \in Red(F)$$
$$\Rightarrow \gamma(\widehat{R}) \supseteq lfp(F) = R^C$$

$\qquad\square$

This section provides a formal definition of the $eval_{\text{mod}}$ function used in the submission. The role of $eval_{\text{mod}}$ is to evaluate an expression in a context that provides partial information about the current loop iteration. Thus, the return value of $eval_{\text{mod}}$ needs to represent partially known values, which is done using the following enumeration:

- $Exact(n)$ when the evaluated expression is known to be exactly n in the given context.
- $Mod(n, k)$ when only the residue modulo k is known to be n.
- $Unknown$ when no information can be derived in the given context.

We overload arithmetic operations (with $bop \in \{+, -, \times\}$) to operate on this partial knowledge as follows:

$$a_1 \; bop \; a_2 = \begin{cases} Exact(n_1 \; bop \; n_2) & \text{if } a_1 = Exact(n_1) \wedge a_2 = Exact(n_2) \\ Mod((n_1 \; bop \; n_2) \bmod k_2, k_2) & \text{if } a_1 = Exact(n_1) \wedge a_2 = Mod(n_2, k_2) \\ Mod((n_1 \; bop \; n_2) \bmod k_1, k_1) & \text{if } a_1 = Mod(n_1, k_1) \wedge a_2 = Exact(n_2) \\ Mod((n_1 \; bop \; n_2) \bmod \gcd(k_1, k_2), \gcd(k_1, k_2)) & \text{if } a_1 = Mod(n_1, k_1) \wedge a_2 = Mod(n_2, k_2) \wedge \\ & \quad \gcd(k_1, k_2) > 1 \\ Unknown & \text{otherwise} \end{cases}$$

We then define an auxiliary function $evalAux_{\text{mod}}$ as follows:

$$evalAux_{\text{mod}}(n, ctx, S) := Exact(n)$$
$$evalAux_{\text{mod}}(e_1 \; bop \; e_2, ctx, S) := evalAux_{\text{mod}}(e_1, ctx, S) \; bop \; evalAux_{\text{mod}}(e_2, ctx, S)$$

$$evalAux_{\text{mod}}(\{e_1, +, e_2\}_i, ctx, S) := \begin{cases} Unknown & \text{if } i \in S \\ evalAux_{\text{mod}}(e_1, ctx, S)+ & \text{if } i \notin S \wedge ctx(i) = peel_n \\ \quad evalAux_{\text{mod}}(e_2, ctx, S \cup \{i\}) \times Exact(n) & \\ evalAux_{\text{mod}}(e_1, ctx, S)+ & \text{if } i \notin S \wedge ctx(i) = unroll_n \\ \quad evalAux_{\text{mod}}(e_2, ctx, S \cup \{i\}) \times & \\ \quad Mod(MaxPeel + n, MaxUnroll) & \end{cases}$$

The function $eval_{\text{mod}}$ is then defined as $eval_{\text{mod}}(e, ctx) = evalAux_{\text{mod}}(e, ctx, \{\})$. The auxiliary function $evalAux_{\text{mod}}$ is recursive, and evaluates an expression by first evaluating its subexpressions in a bottom-up manner. The set S, initially empty, is used to track down the set of induction variables already met along the evaluation path in the top-down direction. Its role is to detect the evaluation of expressions of degree greater than 2 for a given induction variable, and return $Unknown$ in this case. Indeed, the provided formula would be invalid for such expressions when the loop induction variable is not exactly known. Consider $\{1, +, \{2, +, 1\}_x\}_x$ as an example of such an expression, which represents $\frac{(x+1)(x+2)}{2}$, and assume x is only known to be equal to 1 modulo 2. $x = 1$ would lead to $\frac{(x+1)(x+2)}{2} = 3$, but $x = 3$ would lead to $\frac{(x+1)(x+2)}{2} = 10$, which has a different residue modulo 2. Note that the tracking of loop-induction variables to avoid this case is simple but incomplete: it is possible that $evalAux_{\text{mod}}$ returns $Unknown$ in cases in which it would be possible to extract better information. However, in practice, most expressions are linear and thus $evalAux_{\text{mod}}$ almost always succeeds.