

A Survey on Static Cache Analysis for Real-Time Systems

Mingsong Lv¹, Nan Guan², Jan Reineke³, Reinhard Wilhelm⁴, and Wang Yi⁵

- 1 Northeastern University, China
lumingsong@cse.neu.edu.cn
- 2 Northeastern University, China
guannan@ise.neu.edu.cn
- 3 Saarland University, Germany
<http://orcid.org/0000-0002-3459-2214>
reineke@cs.uni-saarland.de
- 4 Saarland University, Germany
<http://orcid.org/0000-0002-5599-7560>
wilhelm@cs.uni-saarland.de
- 5 Uppsala University, Sweden
yi@it.uu.se

Abstract

Real-time systems are reactive computer systems that must produce their reaction to a stimulus within given time bounds. A vital verification requirement is to estimate the Worst-Case Execution Time (WCET) of programs. These estimates are then used to predict the timing behavior of the overall system. The execution time of a program heavily depends on the underlying hardware, among which cache has the biggest influence. Analyzing cache behavior is very challenging due to the versatile

cache features and complex execution environment. This article provides a survey on static cache analysis for real-time systems. We first present the challenges and static analysis techniques for independent programs with respect to different cache features. Then, the discussion is extended to cache analysis in complex execution environment, followed by a survey of existing tools based on static techniques for cache analysis. An outlook for future research is provided at last.

2012 ACM Subject Classification Surveys and overviews

Keywords and phrases Hard real-time, cache analysis, worst-case execution time

Digital Object Identifier 10.4230/LITES-v003-i001-a005

Received 2015-04-23 Accepted 2016-05-11 Published 2016-06-29

1 Introduction

Real-time embedded systems not only exist in industry domains, such as automotive electronics, avionics, telecommunication, medical systems, etc., but are deeply immersed in our everyday life due to the rapid progress of mobile and embedded technology. A real-time system should not only provide logically correct functionality, but moreover, it must meet *timing requirements* as stated in the system specification [21]. In hard real-time systems, such as aerospace systems, a timing error may result in catastrophic consequences. A major task in real-time system verification is to analyze the timing behavior of the system before deployment in order to guarantee that no timing violation occurs at run time.

A real-time system is typically composed of many tasks that cooperate to provide the required functionality. To verify the satisfaction of the timing requirements of the system, one must first know how long each task (or program) may execute. However, this is not an easy problem, because




© Mingsong Lv, Nan Guan, Jan Reineke, Reinhard Wilhelm, and Wang Yi;
licensed under Creative Commons Attribution 3.0 Germany (CC BY 3.0 DE)

Leibniz Transactions on Embedded Systems, Vol. 3, Issue 1, Article No. 5, pp. 05:1–05:48



Leibniz Transactions on Embedded Systems

LITES Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** The distribution of execution times of a program.

the execution time of a program may vary widely as a result of many complex factors, such as data inputs, hardware features, execution contexts, etc. Among all the possible execution times (represented by the yellow range in Figure 1), the minimum and the maximum are called the Best-Case Execution Time (BCET) and the Worst-Case Execution Time (WCET), respectively. The main objective of program-level timing analysis is to estimate the WCET [130], which is then used in system-level timing analysis, such as a schedulability analysis [34, 115].

The common practice in industry has been, and partly still is, to *measure* the end-to-end execution latency of a task [127], by sampling its executions in different scenarios (depicted as the blue vertical lines in Figure 1). The maximal observed execution time increased by a safety margin is used as the WCET estimate of the measured program. This approach is called *dynamic timing analysis* in the real-time community. However, the worst case is not guaranteed to be covered by measurements. Thus, the observed WCET is in general an *underestimation* of the actual WCET. Analytical methods that cover all possible execution scenarios (without executing the analyzed program) and provide safe upper bounds on the WCET are desirable for hard real-time systems. They are usually called *static timing analysis*.

Unfortunately, such upper bounds cannot be easily estimated due to both the complexity of the program itself and the uncertainty from the execution environment. The program may execute different control flow paths depending on input, and these different paths may need different execution times. The execution platform may exhibit a dependence of the execution time of instructions on the execution state of the platform. This *execution state* consists of the occupancy of the platform resources. For example, an instruction may exhibit very different execution times depending on whether instruction or operand fetches hit or miss the cache. The execution environment, finally, may interfere with a program's execution by preempting its execution and thereby increasing the program's response time. Hence, these three factors all have an impact on the program's execution time.

Exhaustive exploration of the combined space of control flow paths and paths through the architectural state space is infeasible due to the size of this space. A conservative abstraction of the execution platform is typically used in static timing analysis to increase efficiency. This abstraction may adversely lead to an *overestimated* WCET. Efforts have to be exerted to reduce the overestimation as much as possible, to avoid the need to over-provision system resources.

All approaches to static timing analysis compute bounds on the execution times of a program starting with bounds on the execution times of individual instructions occurring at points in the program. Their execution times typically depend on the execution state of the platform. Depending on this state, an instruction's execution may suffer from *timing accidents*, which may increase the execution time by their associated *timing penalties*. For example, a memory

access may suffer from a cache miss, which increases its execution time by the cache miss penalty. The actual execution state is the result of the execution history. Different control flow paths through the program, in general, result in different execution states and may thus exhibit different execution times. A classification of an occurrence of a memory access as a cache hit or a cache miss must hold for all executions of this memory access.¹

Static timing analyses determine an invariant for each program point that describes all possible execution states when control reaches this program point. Such an invariant allows excluding many timing accidents, such as cache misses, pipeline stalls, etc., and safely allow subtracting the associated timing penalties from the worst-case upper bound on the execution times.

Among the hardware features to consider in timing analysis, *caches* have the biggest influence on execution time [59]. A precise analysis of the cache behavior does therefore have a great impact on the precision of the overall WCET estimation.

Cache is a small on-chip memory to bridge the speed gap between the processing unit and the much slower off-chip memory by storing a portion of the content from main memory. If a data request *hits* in the cache, it takes only very few processor cycles to deliver the data from the cache to the processing unit; otherwise, in the case of a *miss*, the CPU has to fetch the data from main memory, which nowadays consumes hundreds of processor cycles.

The role of cache analysis for WCET estimation is to predict the behavior of a program on the platform's caches. For example, cache analysis may provide a safe bound on the number of cache hits or misses when a program executes on some given platform; it may also categorize the accesses to memory blocks in programs as definite hits or misses.

In [130], WCET analysis techniques and tools are surveyed. Due to its importance in timing analysis and its complexity, cache analysis alone deserves an in-depth discussion.

The rest of the article is organized as follows. First, we give background knowledge on WCET estimation and caches in Sec. 2. Then, we present the problems and solutions for the intensively researched LRU caches in Sec. 3. A survey of the results on non-LRU caches is provided in Sec. 4. In Sec. 5, the discussions are extended to cache analysis in multi-tasking and multi-core environments where programs interfere with each other on shared caches. A summary of WCET analysis tools based on static cache analysis is given in Sec. 6. We present an outlook for future research at last.


2 Background Knowledge

We first present an established static WCET analysis framework to exhibit its main work flow and where cache analysis steps in. Then, the basic concepts on cache organization, behavior and analysis are introduced.

2.1 A Classical WCET Analysis Framework

The objective of static timing analysis is to compute safe lower and upper bounds on the execution times of programs. These are also called BCET and WCET estimates, respectively. The WCET is observed in a particular execution scenario with some execution context, such as data input and initial hardware state. Theoretically, the WCET is not computable; otherwise, one could solve the halting problem. In this article, we assume that all real-time programs terminate so that their WCET can be computed.

¹ Note that this is a slight simplification to ease understanding. Later on we will explain why it may make sense to partition sets of memory accesses by *contexts*. The distinction between the occurrence of a memory access or an instruction and one, several, or all of their executions is of utmost importance.



■ **Figure 2** The separated path and cache analyses framework for WCET estimation.

Most static analyses are performed on the binary code rather than the source code of the program, because the two need not have the same control flow due to compiler optimizations, and the source code does not determine the precise location of instructions and program variables in memory, which are needed for instruction- and data-cache analysis.


A naïve, straightforward analysis would enumerate all possible executions to find the largest execution time. However, this method does not scale. Consider a loop with a conditional branch inside. If we do not know whether or not the branch is taken in each iteration, the number of program paths to be explored is exponential with respect to the number of loop iterations. To tackle the complexity, the state-of-the-art analysis techniques adopt the framework in Figure 2.

The first step is to reconstruct the Control Flow Graph (CFG) of the program. A CFG is a directed graph, with each vertex representing an instruction and each edge representing the control flow. We say there is a *program point* right before each vertex in the following discussion. A CFG typically has a single entry and a single exit corresponding to the start and the end of the program. The analysis is then conducted on the CFG. In some work [93], WCET analysis is conducted in a modular way on program functions to reduce analysis overhead.

This step is followed by a *Value Analysis*, which computes enclosing intervals for all potential values of registers and local variables and also determines loop bounds. This is a more or less standard *Interval Analysis* as invented by P. and R. Cousot [30]. The next step is to compute an upper bound on the execution time of each instruction (C_i in Equation 1), which heavily depends on the underlying hardware features, such as pipelines [69], branch predictors [18, 28] and caches. *Cache analysis* is an important part of this step, which is often referred to as *micro-architectural analysis* or *low-level analysis*.

With the above results, the final task is to find the execution path that exhibits the longest execution time, typically referred to as WCET calculation. An established approach is the *Implicit Path Enumeration Technique* (IPET) [71], the main idea of which is to transform the problem of searching the worst-case execution path into searching the execution counts for each instruction such that the execution time is the largest. This can be formally modeled as an integer linear programming (ILP) problem, in which the execution time of a program is represented by the sum of execution latencies of all instructions. Thus, the WCET can be obtained by maximizing the execution time (the objective function of the ILP problem):

$$WCET = \max \sum_{i=1}^N C_i \times X_i \quad (1)$$



■ **Figure 3** A common memory architecture.

In Equation (1), N refers to the total number of instructions, which is a constant obtained from the CFG; C_i is the WCET bound for the i^{th} instruction, which has been computed in the third step; the variable X_i stands for the execution count of the i^{th} instruction. X_i is subject to constraints induced by the program structure. In the following, the variable d_{i_j} captures how often the edge from instruction i to instruction j is taken. Then, X_i must be equal to both the total execution counts of all its incoming edges and those of all outgoing edges², which can be expressed as follows.

$$\forall i, X_i = \sum_{\text{all incoming edges}} d_{*_i} = \sum_{\text{all outgoing edges}} d_{i_*} \quad (2)$$

Other program behavior can be constrained as well. For example, the loop iterations should be bounded in advance either manually or by automatic analysis [49]. They can be modeled as linear functions relating the execution counts of the loop body and the loop entry. All available constraints are expressed in one ILP problem, whose maximal solution bounds the WCET from above. To improve analysis efficiency, sequences of instructions (with no branch along the path) are combined into *basic blocks* and represented by a single vertex in the CFG.

The key feature of this framework is the separation of micro-architectural analysis from WCET calculation. In general, this approach is pessimistic. However, the sacrifice of precision is rewarded by a significant improvement in analysis efficiency.

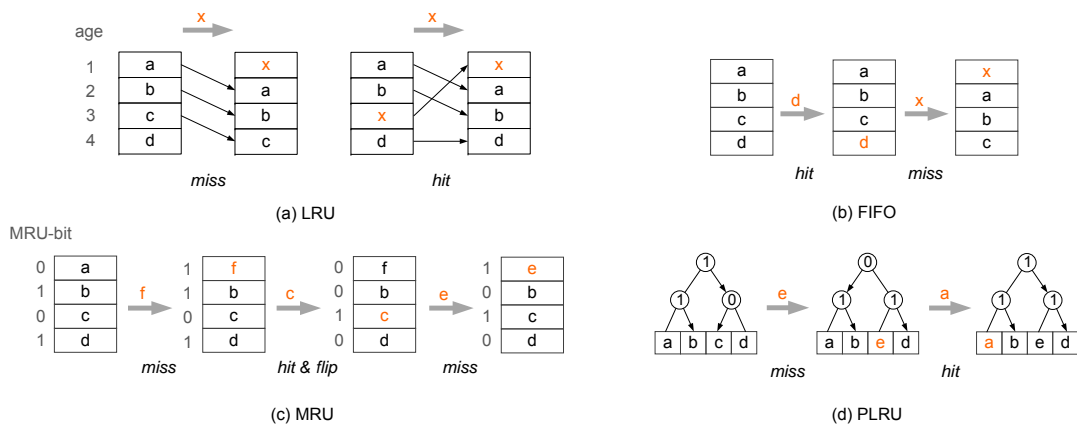
2.2 Cache Organization, Behavior and Analysis

Cache is a small, high-speed memory residing on the processor chip (shown in Figure 3) that stores a copy of a portion of the instructions and/or data in main memory. Each access to the cache results in either a *hit* or a *miss*. One can distinguish two types of cache misses. A *cold miss* occurs when a data element, absent from the cache, is loaded for the first time. If the cache is full and a cache miss occurs, a data element needs to be evicted. A *replacement miss* occurs when an evicted element is reused. Cache hits are the result of *memory reuse*.

2.2.1 Cache Organization and Behavior

In most processors, a cache line (the unit for cache access) contains multiple data elements. An access to one element causes the whole cache line to be loaded into the cache. As a result, a following access to another element of the same cache line also results in a cache hit. Besides, consecutive accesses to the same data element result in cache hits as well, an example of which is the execution of a loop. The above two types of reuses are commonly referred to as *spatial reuse* and *temporal reuse*, respectively, the pervasiveness of which is expressed by the well-known *locality principle*.

² For either the entry or the exit instruction, one can simply constraint the execution count to be exactly 1.



■ **Figure 4** Common cache replacement policies.

Some processors are equipped with two or more levels of caches, as a fine-grained trade-off between cost and speed. The lowest level³ (namely L1 cache) is usually divided into a private instruction cache and a private data cache, each of which is typically no larger than 32 KB and has an access latency of 1 – 2 cycles. If a memory access misses in the L1 cache, the L2 cache is queried. The capacity of L2 caches may range from hundreds of KB to several MB, with an access latency of around 10 cycles. In some high performance multi-core processors, an L3 cache may also be deployed to further expand cache capacity. Misses in the last level cache trigger accesses to the main memory via the *off-chip* memory bus, causing a delay in the order of hundreds of cycles.

Like other storage devices, addressing is an important feature of the cache design. Some processors adopt the *set-associative* organization, in which the address space is partitioned into independent *sets*. Every set has a fixed number of *ways*, each of which refers to a single cache line in every cache set. The total number of ways within a cache set is called *associativity*. To load a memory block, the processor first determines which cache set the block maps to. Then a lookup into the target set is performed for a free cache way. If all the cache ways are occupied, a *replacement policy* determines which old block to evict to make room for the new block. In this article, we consider four common policies illustrated in Figure 4, assuming a 4-way cache set.

The least-recently-used (LRU) policy replaces the block that has been used least recently. The illustration of LRU in Figure 4(a) is a first abstraction from the actual hardware implementation. Each cache way in an LRU cache set is associated with a fixed age, which is received by the block in the corresponding cache way. Figure 4(a) illustrates how the positions of the blocks are reordered upon a cache hit and a cache miss.

However, most commercial processors do not employ LRU, because it requires complex hardware implementation and further leads to higher power consumption. Non-LRU replacement policies, such as First-In-First-Out (FIFO), Most-Recently-Used (MRU) and Pseudo-LRU (PLRU), are adopted instead since they are simple to implement and still have similar average performance as LRU [58].

Figure 4(b) shows how the FIFO replacement policy works. A cache hit does not change the cache state. Upon a cache miss, all the memory blocks shift one position downwards, evicting the block in the bottom cache way; then the new block is installed in the top-most cache way. Again, this representation is an abstraction from the actual hardware implementation, which does not

³ A cache level is lower if it is closer to the processing unit; the highest level is typically called the last level.

shift memory blocks from one cache line to another, but rather maintains a modulo-4 counter to determine the next block to replace.

The MRU cache (shown in Figure 4(c)) maintains a bit for each cache way (called MRU-bit) to approximate the recency of access. Bit 1 means the block was visited recently. Upon a hit, the MRU-bit of the hit block is set to 1. Upon a miss, the top-most way with MRU-bit 0 is taken by the new block, and its MRU-bit is set to 1. Eventually there is only one cache way with MRU-bit 0. When this way is visited—in the example the access to c —the MRU-bit is turned to 1, and all the other MRU-bits are set to 0. This is called a *global flip*.

PLRU is a tree-based approximation of LRU (Figure 4(d)). It arranges the cache ways at the leaves of a binary tree with $k-1$ bits, where k is the cache associativity. Bit 0 and 1 on the branches indicate the left and the right subtrees, respectively. Following the bits downwards from the root, the cache line to be replaced or refilled can be found. After an access (either hit or miss) to a cache way, all tree bits along the path from it to the root are set to point away. It is possible that a cache set contains invalid cache lines. We assume the *tree-fill* policy, by which the line to be filled or replaced is always determined by the tree bits.

In most architectures, cache sets are completely independent of each other. This makes the independent analysis of programs' behavior on different cache sets possible. Throughout this article, we focus on the cache behavior in one set, and may use *cache* to refer to a cache set to simplify the presentation when appropriate.

2.2.2 Cache Analysis

The objective of cache analysis is to statically determine the cache behavior of a program. Its results can be used for performance analysis and optimization. The results may be of several types: one is the *classification of individual memory accesses in a program as hits or misses*. Such a classification of memory accesses can be used in a cooperating pipeline analysis to determine whether the pipeline may have to stall on an instruction or operand fetch. Another is a *bound on the number of cache loads in a segment of the program*. This allows, under certain conditions, just adding an accumulated penalty to the execution-time bound for memory accesses that could be neither classified as cache hits or misses. The former type of cache analysis could be called a *classifying cache analysis*, one instance of the latter, relevant for practice, is known as *persistence analysis*.

A typical use of the results of a cache analysis in real-time systems is estimating the BCET and WCET of programs. The bigger the percentage of hits that will happen during execution it can predict, the tighter the WCET estimation is. On the other hand, predicting a higher percentage of actual misses leads to tighter BCET estimation.

The designer of a cache analysis faces several questions: the first one is whether the analysis is to be a classifying or a persistence analysis. The second question is by which method cache behavior should be analyzed. Associated with the second question are the questions of the granularity of the analysis and the representation of the cache behavior properties.

Let us illustrate this rather abstract discussion with the example of classifying cache analyses. The most precise analysis would predict each *executed memory access* to be either a hit or a miss—here it is vital to make the difference between an *executed memory access* and the *occurrence of an instruction involving a memory access* in a program. We have assumed that all real-time programs terminate. Hence, any program would execute only finitely many memory accesses, so that such an analysis would in principle be possible. However, the corresponding analysis would, in general, not scale. On the other end are cache analyses that would classify *occurrences of memory accesses* as *always hit* or *always miss*, where *always* means for all executions of this occurrence of the memory access. However, experience has shown that the actual execution times of memory

accesses associated with one occurrence of a memory access may vary widely. This means that just taking their upper bounds may largely overestimate the memory-access costs. Precise and efficient analyses should attempt to classify subsets of executed memory accesses corresponding to one occurrence, such that the accesses in the subsets have some homogeneous timing behavior. The subsets would be characterized through control flow criteria, in the following called *contexts*. The most important examples for contexts are different iterations of loops.

To approach the second question raised above, one needs to identify the information about cache contents—in the following mostly called *concrete cache states* (CCS)—to be computed by a classifying cache analysis. This information would provide answers to the question, *are all memory accesses belonging to this occurrence (in this context) hits or are they all misses?* One solution would be to collect at each program point the set S of all concrete cache states that are possible when program control reaches this program point (in this context). Such an analysis would again not scale. Instead, one can represent sets of concrete cache states by *abstract cache states* (ACS). Each abstract cache state (compactly) represents a set of concrete cache states. As we will see later, two types of such abstract cache states are of interest. Consider the set S of all concrete cache states that are possible at a program point. An *abstract Must cache state* will represent the information: which memory blocks will be in each of the possible concrete cache states in S . This is obtained by some kind of intersection applied to the elements in S . Likewise, an *abstract May cache state* will represent the information: which memory blocks may be in one of the concrete cache states in S . This is obtained by some kind of union applied to the elements in S .

3 Analysis of LRU Caches

For decades, a majority of research on cache analysis has focused on caches with LRU replacement strategy. In this section, we survey the main analysis techniques with an emphasis on the approach based on Abstract Interpretation (AI) [38]. This technique is realized in the aiT tool of AbsInt [57], which is widely used in industry. Since programs spend most of their execution time in loops, a sub-section is dedicated to the analysis of the cache behavior in loops. The big picture on LRU caches is completed with further discussions on data cache and multi-level cache analyses.

3.1 Abstract-Interpretation-Based Approaches

The first cache analysis based on abstract interpretation (AI) was proposed by Ferdinand and Wilhelm in the 1990s [1, 38]. The overall approach works in two phases:

1. An AI-based cache analysis computes abstract cache states at all program points as part of a fixed-point solution;
2. These abstract cache states are queried in order to classify memory accesses.

3.1.1 A Short Introduction to Abstract Interpretation

Abstract interpretation [30] is a static program analysis method based on a semantics of the considered programming language. Instead of executing the program on the concrete domain of values, it executes an abstracted version of the program on an abstract domain of descriptions of values. In the case of cache analysis, the program abstraction only describes the memory-access behavior of the program, i.e., it performs all memory accesses that the program would execute. This abstracted program works on *abstract cache states*, which are descriptions of sets of concrete cache states. One abstract cache state is associated with each program point. Whenever the analysis encounters a memory access, it updates the abstract cache state in a way induced by the update that the processor would perform on the concrete cache states. Whenever the control flow



■ **Figure 5** An example to show the \sqsubseteq relation w.r.t. the Must domain.

of the program merges, e.g. at the end of a conditional or at the header of a loop, it combines the incoming abstract cache states in a sound way.

Gary Kildall [64] has recognized that the abstract domains of typical data-flow analyses are lattices, i.e., partially ordered sets where all subsets have least upper bounds. The partial order reflects the relative information content of two lattice elements. By convention, elements lower in the lattice represent more information than information higher in the lattice, i.e., an element a below or equal to an element b in the lattice, $a \sqsubseteq b$, contains no worse information than b . The domain of abstract cache states together with a partial order reflecting the amount of knowledge about cache contents, in this sense, also forms a lattice.

Consider two abstract Must cache states \hat{a} and \hat{b} (as shown in Figure 5). Abstract cache state \hat{a} represents just one concrete cache state, containing memory blocks $\{u, x, y, z\}$ with the ages 1, 2, 3, 4. Abstract cache state \hat{b} represents the set of concrete cache states with memory block u having age 1, x having age at most 3, z having age at most 4, and possibly one more (unknown) block at age 2, 3, or 4. $\hat{a} \sqsubseteq \hat{b}$ means that all the concrete cache states represented by \hat{a} are also represented by \hat{b} . In particular, this implies that all the memory blocks known to be contained in the concrete cache states described by \hat{b} , in the example above u, x, z , are also known to be contained in the concrete cache states described by \hat{a} . Furthermore, \hat{a} additionally tells us that, (1) block y is guaranteed to be in the cache while \hat{b} does not; (2) the age upper bound estimated for block x in \hat{a} is smaller than that in \hat{b} . Clearly, the abstract cache state \hat{a} contains better information than \hat{b} . As stated above, at control flow merge points cache analysis must combine the incoming information in a sound way. The operation applied to the incoming abstract cache states is the least upper bound, \sqcup , of the lattice. This is shown in Figure 6. As said above and made more precise later, it is some form of intersection. It determines (the best) safe information holding for all incoming paths.


The lowest element in the lattice of abstract cache states, \perp , called *bottom*, describes the empty set of concrete cache states. It is the initial analysis information at all program points but program entry. If we do not have any information about the cache contents at program entry, the highest lattice element, \top , called *top*, is used as initial analysis information. It describes the set of all concrete cache states, and thus the absence of information about cache contents.

The update functions are, in general, monotone, so that information, once computed, is not lost again. A fixed-point iteration over the control flow graph of the program is guaranteed to terminate and deliver the least fixed point as solution. Essential for termination is the finiteness of the lattice.

Let us summarize this short introduction to AI by listing the main ingredients of a particular abstract interpretation. The designer needs to choose or define an *abstract domain*, a lattice of abstract values, which are descriptions of (sets of) concrete values. The *partial order* defines the relative information content of two lattice elements. The *least upper bound* is the operation to join abstract values flowing to a program node through different control flow graph edges. *Abstract update functions* for an instruction reflect the instruction's effect on the incoming abstract values.

■ **Table 1** Cache Hit/Miss Classification.

Classification	Cache Access Behaviors Described	Analysis
Always Hit (AH)	Block is guaranteed to be in the cache upon each memory access	Must
Always Miss (AM)	Block is guaranteed not to be in the cache upon each memory access	May
Not Classified (NC)	Cannot be classified by any of the above classifications	/



■ **Figure 6** An example to demonstrate the Must join function.

3.1.2 Classification of Memory Accesses


Querying an abstract cache state, resulted from a cache analysis, for an accessed memory block may yield qualitative properties as listed in Table 1. To determine whether the memory access to m is *always hit* (AH), one simply checks the existence of m in the abstract Must cache reaching its program point. Similarly, to determine whether the memory access to m is *always miss* (AM), it suffices to know that m is not in the abstract may cache reaching its program point. If a memory access can be neither classified as AH nor as AM, it is classified as NC. An NC classification can have two reasons: (1) some of the executions of a memory access hit in the cache and others miss in the cache, or (2) the analysis method overapproximates the set of concrete cache states and thereby fails to deliver the correct classification. Research results show that these properties are able to cover most access behaviors for LRU caches [38]. The classifications can then be expressed as linear constraints on the execution cost of each instruction (basically each instruction generates a single memory access) and later integrated into a WCET estimation. In architectures without timing anomalies [22, 106], if the classification of a memory access is NC, it is safe to treat it as AM. The properties AH and AM in Table 1 are explored by independent analyses, which are now described in detail.

3.1.3 Must Analysis

The objective of a Must analysis is to compute a Must-ACS at each program point, *which represents the common cache contents in all possible executions leading to this program point*. An age is associated with each memory block in the Must-ACS, which is an upper bound of its ages in all CCS. We use a graphical representation to show a Must-ACS, in which blocks are grouped according to their ages, e.g., in Figure 6. The set of CCS represented by a given Must-ACS is formally defined by the *concretization function* below, where c and \hat{c} denote concrete and abstract cache states, respectively, and $age(c, m)$ refers to m 's age in cache state c (applies to both concrete and abstract states).

$$conc^{Must}(\hat{c}) = \{c \mid \forall m \in \hat{c} : m \in c \wedge age(c, m) \leq age(\hat{c}, m)\} \quad (3)$$

The Must-ACS at the program entry is initialized with \top representing all concrete cache states if the initial cache content is unknown; all other nodes are initialized with \perp representing the empty set of concrete cache states. A fixed-point computation is employed to compute the Must-ACS at each program point, during which two main operations over the Must-ACS are involved.



■ **Figure 7** An example to demonstrate the Must update function.

A *join function* combines several Must-ACS into a single Must-ACS when the control flow merges. The resulting Must-ACS takes the intersection of the sets of blocks in all the incoming states and assigns to each block its maximal age from the incoming states, as shown in Figure 6. An *update function* $\hat{U}(\hat{c}, x)$ defines how an abstract state \hat{c} is changed due to an access to memory block x , specifically, how the age of each block in the ACS is updated. Figure 7 shows an example. A correct Must update function guarantees that the age of each block in the computed ACS is a safe age upper bound for all possible represented CCS. For example in Figure 7, the age of d in \hat{c} implies that there could be a CCS represented by \hat{c} , in which d has an age of 4, such as c_2 . By loading x , we can no longer guarantee that d still stays in any resulting CCS. Thus, d has to be removed from the computed abstract state \hat{c}' .

3.1.4 May Analysis


May analysis computes a May-ACS at each program point, which represents all potentially cached contents in all possible executions leading to this program point. If block m does not exist in the May-ACS at the reaching program point, we can guarantee the access to m is AM. Unlike the Must-ACS, the age of each block in a May-ACS is the lower bound of its ages in all represented CCS, as expressed by the May concretization function below, with the same notions as in function (3).

$$\text{conc}^{\text{May}}(\hat{c}) = \{c \mid \forall m \in c : m \in \hat{c} \wedge \text{age}(\hat{c}, m) \leq \text{age}(c, m)\} \quad (4)$$

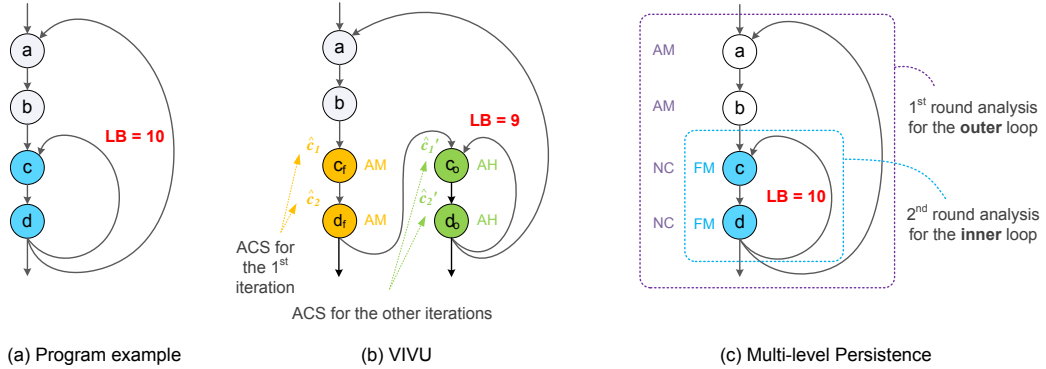
The May join function takes the union of the sets of blocks in all incoming May-ACS and assigns each block the minimal age in all incoming states. The May update function is exemplified in Figure 8, where \hat{c} is the May-ACS representing the concrete states c_1 and c_2 . Take memory block d for example. d 's age in the May-ACS is the minimum of those in c_1 and c_2 . After the access to x , d is evicted from c_2 . However, d still remains in the resulting May-ACS \hat{c}' , because \hat{c}' must soundly represent the other resulting CCS c_1' in which d remains. To determine whether the memory access to m is AM, it suffices to know that m is not in the May-ACS reaching its program point. A block m is not in the final May-ACS at a program point because either m has never been loaded or enough different blocks have been loaded to evict m from the cache. May analysis does not directly help with tighter WCET estimations, however, predicting more misses results in better estimations on BCET.

3.2 Improving Precision by Using Contexts

In practice, merely relying on classifying analyses, such as Must and May analyses, may still largely overestimate the memory-access cost and thus the WCET. Methods to improve the knowledge about the cache behavior are proposed by taking program structures into consideration, in particular loops. The cache behavior of programs in loops is somewhat special: the first iteration typically loads the contents into the cache; later iterations profit from the first iteration since



■ **Figure 8** An example to demonstrate the May update function.




■ **Figure 9** The ideas of VIVU and multi-level Persistence analysis.

accesses to the cached contents are hits. The concrete state on return to the loop header may thus be very different from that reaching the loop from the outside. This is also reflected in cache analysis where the abstract state upon return from the first iteration may be very different from that on the entrance to the loop. Naïvely applying the join function to these two abstract cache states would produce very bad information about the cache behavior of the loop. In such cases, a large percentage of the accesses to the memory blocks in the loop cannot be classified as either AH or AM by the previously introduced Must and May analysis. However, there are two alternative ways to tighten the WCET estimation. The first one exploits the above observation by virtually unrolling each loop followed by a Must analysis. An access to a memory block that in the first iteration would be classified as AM, and that in the other iterations would be classified as AH would then be classified as FM (first miss). The other alternative would be to bound the number of cache misses for all the accesses to a memory block within a certain program scope. These two analysis techniques will be now introduced. Note that the first analysis still is a classifying analysis for memory accesses, albeit with a new classification, FM, and the second is a bounding analysis for memory blocks in a scope.

3.2.1 Virtual Inlining & Virtual Unrolling (VIVU)

VIVU [85] can be used to improve cache analysis precision for loops. The idea is to analyze the first loop iteration separately from all other loop iterations. This is done by virtually unrolling the first iteration of the loop body⁴, so as to distinguish the behavior between the two contexts. Then, a Must analysis is applied to the program with the unrolled loop to find the AH memory

⁴ (1) The unrolling is called *virtual* since it is done by maintaining separate abstract cache states for the first iteration and the remaining iterations (e.g., \hat{c}_1 and \hat{c}'_1 in Figure 9(b)). For ease of understanding, we use a physically unrolled CFG to show the effects. (2) VIVU allows to unroll more than one iteration of the loop since iterations other than the first may have vastly different behavior and thus execution times. Here we assume only unrolling the first iteration to simplify presentation.



■ **Figure 10** Operations for Persistence analysis.

accesses in all but the first loop iterations. Figure 9(b) shows the results of unrolling the inner loop of the program in Figure 9(a). In the new CFG, c_f and c_o refer to the first and the other iterations of the inner loop, respectively (similar for d). Assume the cache is 2-way associative. Must analysis on the new CFG is able to classify c_o and d_o as AH, and thus c and d as FM.

3.2.2 Persistence Analysis

One can aim at the same analysis objective by a Persistence analysis. There are several possible notions of persistence of memory blocks one could aim at. These are:

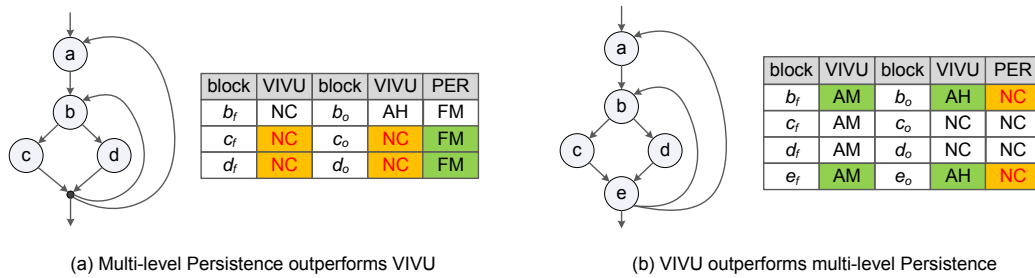
- persistence** execution causes at most one miss for the memory block,
- first miss** only the first access is a miss, all others are hits,
- no eviction** the block is never evicted after a possible miss.

For a memory block that is persistent in a program fragment, timing analysis can assume a bound of one cache load for all accesses within that program fragment.

The first Persistence analysis was proposed by Mueller et al. [88, 89] for computing First Miss classification of memory references. For set-associative caches, the basic idea of Mueller’s approach is to check whether all conflicting instructions in a loop fit into the cache. Later, Ferdinand and Wilhelm [38] proposed a Persistence analysis based on Abstract Interpretation. This analysis employs abstract cache states, Per-ACS, as do the Must and May analyses. The fact that a block has been visited and already evicted from the cache is modeled by assigning an age \top to the block, where \top is larger than the cache associativity. The Per-ACS at each program point represents the cache contents that are potentially visited and then guaranteed to remain in the cache. If a memory block exists in the Per-ACS at the end of the scope, then one can guarantee that at most one cache miss may occur for all the accesses to this block.

The concretization of a Per-ACS is a set of traces satisfying the persistence condition, i.e., at most one miss for each block with a non- \top age in the Per-ACS. More precisely, a Per-ACS captures upper bounds on the ages of memory blocks, assuming that they have already been accessed at least once in the execution of the program.

Figure 10(a) gives an example for Persistence update. All blocks in Per-ACS \hat{c} are potentially visited during program execution. By accessing x , d is no longer guaranteed to be in the cache since it has age 4 in \hat{c} , which is maintained by putting d in the \top -age line. Block f , already evicted from the cache before accessing x , remains unchanged in the \top -age line. The update function for Persistence analysis mainly needs to guarantee the maximal age for each block in the ACS is soundly maintained.



■ **Figure 11** Comparing VIVU with multi-level Persistence analysis.

At a control flow merge point, several Per-ACS are merged by the join function, which takes the *union* of the blocks in all the incoming Per-ACS and assigns each block the *maximal* age from the incoming states, as shown in Figure 10(b). Intuitively, the set union operation guarantees the resulting Per-ACS does not lose track of any potentially visited memory block; the maximal age ensures that we can safely predict whether a block is definitely persistent after its first access.

The Persistence analysis by [38] was recently found to be unsafe, due to an error in the update function, which may incorrectly underestimate the age of a block. The error was corrected by Cullmann [31] and Huynh [62]. Several different ways were proposed to restrict the set of memory blocks in a Per-ACS to the actual capacity of a cache set. The simplest way, yielding the least precise results, marks all memory blocks in a Per-ACS as non-persistent, if the Per-ACS contains more blocks than the associativity allows. Others check the number of conflicts between members of the Per-ACS; Huynh’s analysis employs fixed-point computation to collect for each block m a set of potentially conflicting blocks (blocks that are mapped to the same cache set with m and thus may age m). If the total number of conflicting blocks is no larger than $A - 1$, where A is the cache associativity, then one can safely draw the conclusion that m , once loaded, will persist in the cache. A similar idea was also applied in the Persistence analysis [90] by Mueller. Cullmann presents a number of similar persistence analyses [31]. The most precise of Cullmann’s analyses relies on a May analysis to make correct decisions on age update.

3.2.3 Analysis Scope

A bounding analysis, such as the Persistence analysis, is designed to investigate cache behavior within a *program scope*, in most cases a loop body. It is common for a program to have nested loops, where a block in an inner loop also belongs to the outer loop. A natural question would be: does the block have a different cache behavior for different loop levels? To distinguish a block’s behavior, the relevant loop nest(s) (the inner loop, the outer loop, or both) is/are unrolled in the VIVU approach. Multi-level approach were proposed by Mueller et al. [128] and Ballabriga and Cassé [10]. Ballabriga and Cassé [10] apply the Persistence analysis of [38] on the relevant scope, here specifically the relevant loop nest, to explore local cache behavior. Figure 9(c) illustrates the basic idea. Persistence properties regarding different loop nests for a memory block can be encoded as linear constraints (or other forms) and integrated into WCET computation for tighter estimations.

3.2.4 Comparing VIVU and Persistence Analysis

Since both VIVU and Persistence analysis are able to bound cache misses for a program scope, a straightforward question would be: which one is more precise? In fact, the two techniques are generally incomparable. For the program in Figure 11(a), c_o and d_o cannot be classified as AH

by Must analysis [38] of the VIVU approach, as neither c_o nor d_o is guaranteed to be accessed in any loop iteration (they belong to two respective conditional branches). On the other hand, multi-level Persistence analysis is successful in this example for c and d . In Figure 11(b), both b_o and e_o can be classified as AH by VIVU. However, none of b , c , d or e can be classified as FM by the multi-level Persistence analysis in the static cache simulation framework [90]. This is because the Persistence analysis [90] counts the conflicting blocks in a loop (mapped to the same cache set); if the number is larger than the cache associativity, none of the blocks in the loop can be classified as FM. If, otherwise, a different Persistence domain is adopted in the multi-level analysis, such as [31], b and e can be locally classified as persistent.

VIVU and Persistence analysis can also be compared in terms of analysis cost. On the one hand, VIVU may result in a more expensive micro-architectural analysis, having to distinguish multiple contexts. On the other hand, the results of Persistence analysis need to be encoded into constraints during implicit path enumeration. The influence of the two effects on analysis times has not yet been compared empirically.

3.3 Other Techniques


3.3.1 Static Cache Simulation (SCS)

There are essentially two approaches preceding AI-based approaches. Mueller et al. [88] developed *static cache simulation* to categorize memory references as *Always hit*, *First hit*, *First miss*, or *Always miss*. They were the first to propose to use *abstract cache states (ACS)*, starting for direct-mapped caches [88].

They later attempted to extend this approach to set-associative caches [89, 90, 128]. For set-associative caches, abstract cache states are defined to hold potentially cached memory blocks at all possible positions, i.e., ages. This results from using set union as join operation at control flow merge points. They also give an update scheme for abstract cache states, which, unfortunately, computes incorrect age lower bounds [89, 128]. The abstract cache states resulted after fixed-point iteration are then used to derive Always hit, First hit, First miss, or Always miss categorizations. It is far from trivial to derive Must information from information contained in their abstract cache states. Spatial locality can be easily exploited. Beyond that there is no obvious way to compute sound and precise Must information. In [90], Mueller employed dominator information in addition to abstract cache states to compute correct Must information.

3.3.2 Cache State Transition Graphs (CSTGs)

Also before AI-based approaches, Li et al. presented a technique that uses *Cache State Transition Graphs (CSTGs)* to model cache behavior [72]. A CSTG, built out of the CFG, models the cache-state transitions for a given cache set. A vertex in the CSTG stands for a possible concrete cache state, and each edge in the CSTG represents a possible transition from the source state to the destination state due to a memory access in the program. Instead of exploring qualitative properties, such as AH, AM and FM, the analysis tries to find a lower bound on cache hits for each memory block. The bounds can be modeled as linear constraints and combined into the ILP to obtain the WCET for the program. By explicitly enumerating the concrete cache states, the CSTG approach can provide good analysis precision. However, it does not scale with program size. Assume that there are M memory blocks mapped to each cache set with associativity K , the number of states in an CSTG can be calculated by $\sum_{i=0}^K \frac{M!}{(M-i)!}$ [72]. Note that the number of linear constraints is of the same scale as the number of CSTG states. In practice, the analysis efficiency is low due to the complexity of the resulting ILP problem.



■ **Figure 12** The analysis framework of McAiT.

3.3.3 Model-Checking-Based Methods

Model checking [26] is a powerful technique widely used in system-level timing analysis of real-time systems. Timed automata [6] have been used to model the cache behavior of programs, and a model checker has been employed to find the WCET. Existing work includes the McAiT tool [80], the METAMOC approach [32], and Gustavsson et al.’s analysis [50]. These works use the model checker (UPPAAL in their cases) as a black-box tool. They express the cache access behaviors of a program by the modeling language of the model checker, and verify whether the WCET of the program is bounded by a specified value as a reachability problem.

Figure 12 shows the architecture of the McAiT tool. McAiT first constructs the program automaton out of the CFG, which fully simulates the behavior of the program, such as the control flow and how the program accesses caches. For a given cache configuration, McAiT builds a timed automaton to model each cache. The execution of an instruction causes the program automaton to issue messages to the cache automaton via UPPAAL’s channel mechanism, and the cache automaton updates the cache state accordingly. The timed automata models for both the program and the cache are then explored using the UPPAAL model checker to find the WCET.

Essentially, the estimated WCET by a model checker is the actual WCET of the program, since all the possible executions are explored. Cache hits and misses for *each execution of memory accesses* are precisely reported. The major difference between this approach and the CSTG approach is that the possible cache states are not explicitly modeled in the automaton, but rather explored by the model checker. The main drawback of model-checking-based approaches is their lack of scalability, since an exponential state space has to be explored.

3.4 Data Caches

Modern processors are typically equipped with *data caches* to improve the performance of data accesses. Instructions are fetched from known addresses; so instruction fetching can be accurately analyzed. In contrast, data accesses are less predictable [76, 123].

3.4.1 Main Challenges


Before predicting hits/misses for data accesses, the set of data addresses accessed by each instruction needs to be determined, referred to as value analysis or address analysis [9, 129]. The threat to precision is that an imprecise value analysis may not be able to eliminate memory accesses that do not occur in a real execution. The problems are the following: Firstly, data manipulations using redirectable pointers make it hard to statically determine the data items actually accessed. Secondly, in the presence of dynamic data structures on the heap, the data addresses can only be determined at run-time (due to this problem, dynamic data structures are typically avoided in hard real-time systems). Lastly, value analysis may work with abstractions of memory addresses, such


```

1 for (i = 0; i < N; i++)
2   for (k = 0; k < N; k++)
3     for (j = 0; j < N; j++)
4       C[j][i] += A[k][i] * B[j][k];

```

■ **Figure 13** An example of matrix multiplication.



■ **Figure 14** Data cache analysis based on the pigeonhole principle.

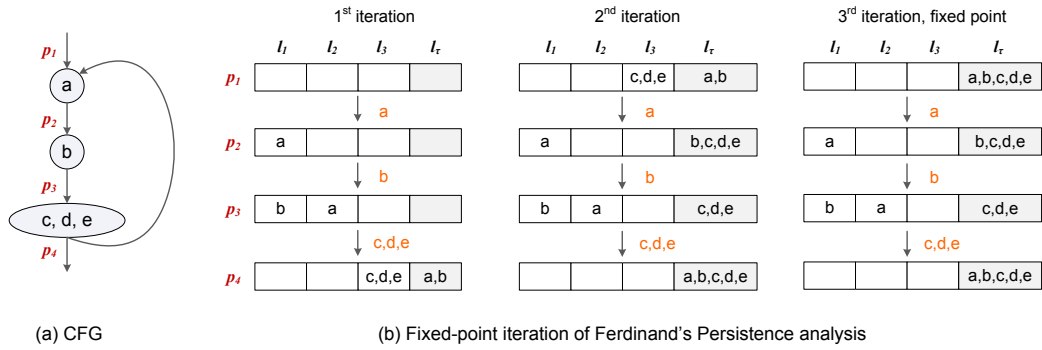
as intervals. As a result, the address range they compute may be overapproximated. Considering non-feasible data accesses in cache behavior analysis increases the probability of not being able to classify memory accesses as cache hits or misses. In addition, a memory access without a precisely determined address pollutes the information contained in an abstract data cache since the update function has to be applied to all potential concrete addresses.

Besides that, data cache analysis is challenged by another problem: executing an instruction may generate accesses to *multiple* data addresses. Figure 13 depicts a program with a matrix multiplication, in which line 4 generates accesses to different matrix elements (in different loop iterations). For this simple program, one can easily determine the data items accessed in each loop iteration. But, in general, data accesses could be very unpredictable due to input dependence of the array indexes. Consider accesses to array $A[x][y]$: if the values of x or y are not clear, one has to conservatively assume that any address in the whole array could be accessed. Furthermore, classifications of memory accesses as used for instruction-cache analysis (AH, AM, FM) may not be sufficient to describe data cache behavior.

3.4.2 Analysis without Input Dependence

Early work, such as the Cache Miss Equation (CME) framework [39], focused on analyzing programs with predictable data accesses. The underlying idea is to set up mathematical formulas (Linear Diophantine equations specifically) to precisely capture both spatial and temporal memory reuses by relating data addresses, loop induction variables and cache parameters. From the solution of the equations, one can check if a memory block is evicted from the cache before it can be reused. An upper bound on the number of misses can thus be obtained for WCET estimation.

However, only a small set of programs can be analyzed by the CME framework: (1) loops must be rectangular loops and perfectly nested; (2) array subscript expressions and the bounds of the loop index must be affine combinations of the enclosing loop indices; (3) no data/input-dependent conditions may exist. The CME framework has been later extended to allow function calls [124], conditionals only depending on the loop induction variables [124], and multiple loop nests [97]. Unfortunately, none of these methods can deal with input dependence. Clauss presented an approach of solving cache miss equations through the mapping to Ehrhart polynomials [27]. Still, the complexity of solving these polynomials is high. Another approach is the Presburger Arithmetic framework [23], which has similar restrictions and is computationally expensive.



■ **Figure 15** Ferdinand's Persistence analysis for data caches.

3.4.3 Analysis with Input Dependence

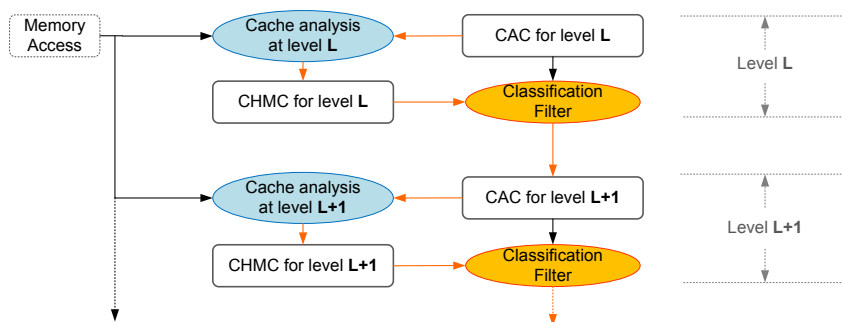
Earlier research to handle input dependence focused on direct-mapped caches. Kim et al. proposed an analysis method based on the pigeonhole principle [65]. Figure 14 shows 3 iterations in the execution of a loop in which a , b , c , and d can be accessed. If in total 9 memory accesses are generated, then at least $9 - 4 = 5$ among them must be cache hits. The 4 cache misses are cold misses. This work was later extended by Staschulat and Ernst to handle programs with unpredictable input dependency [112], in which cache misses are bounded according to data access types: (1) cache misses from predictable accesses are bounded by the pigeonhole principle; (2) cache misses from unpredictable accesses are tracked down by a miss counter and expressed with linear constraints. Unfortunately, these methods are still too restrictive. First, they only work for direct-mapped caches. Second, the loops must fit into the cache to utilize the pigeonhole principle. Essentially, these approaches correspond to simple Persistence analysis for the special case where programs fits into the cache.

AI-based analysis techniques are extended to analyze set-associative data caches with input dependency. Ferdinand extended the Persistence analysis [37] with a new update function to handle multiple memory accesses by one instruction. The basic idea and its drawback can be explained by the example in Figure 15.

Figure 15(a) gives the CFG of a loop in which p_1 to p_4 are program points. Note that each time the instruction after p_3 is executed, one of the blocks from $\{c, d, e\}$ could be accessed. Figure 15(b) shows the fixed-point iteration process, given a cache size of 3. The last column l_τ of each abstract state is used to collect the blocks that have been evicted from the cache (a common structure for most Persistence abstract domains). On the transitions from p_3 to p_4 , since it is not clear which of the three blocks (or their combination) is actually accessed, they pessimistically assume that all blocks in $\{c, d, e\}$ could be accessed and cause other blocks to age. Thus, both a and b are evicted from the cache (collected in l_τ). Moreover, since there is no knowledge on the access sequence of c , d and e , they receive an age of 3 when they are brought into the cache state in the 1st iteration. As a result, no cache hit can be predicted for this loop. As only one from c, d, e may be accessed in every iteration, slightly better results would be possible with a different transfer function.

Sen and Srikant developed a Must analysis for data caches [110]. The analysis can be combined with VIVU to discover the persistence property of data accesses. Despite some small differences, the age manipulation in Sen's Must analysis are similar to Ferdinand's Persistence analysis [37], and thus may lead to very pessimistic estimations.

Ferdinand's and Sen's analyses show that without modeling data access patterns, the abstract domain has to do very conservative age maintenance. Again, for the program in Figure 15(a), if by some means we know that the lifetime of c , d , and e do not overlap, then the analysis can be improved. For example, if the loop iterates for 30 times, c is only accessed in iterations 1 to 10, d



■ **Figure 16** The separate analysis architecture.

only in iterations 11 to 20, and e only in iterations 21 to 30, then c , d and e cannot evict each other in their lifetime. They are actually *persistent* (given a cache size of 3) once they are loaded into the cache, since any of them can only be aged by a and b . Based on this observation, Huynh et al. proposed scope-aware data cache analysis [62]. Each memory access is now associated with a *temporal scope* to model its lifetime, as an augmentation to the traditional AI-based analysis. In the update function, memory accesses that have no overlapping temporal scopes do not cause each other to age. In consequence, some non-existing access conflicts are excluded, and more persistent data accesses can be identified (such as c , d and e).

Hahn and Grund observe that cache analysis does not require knowledge of *absolute addresses* of memory accesses. Instead, it is sufficient to know about the *relation* between the addresses of different memory accesses: do they refer to the same cache block, a different cache block but the same cache set, or different cache blocks in different cache sets? Based on this insight, they developed *relational cache analysis* [51], which can classify accesses as cache hits even if the absolute address of the access is unknown.

To summarize, input dependence makes data-cache analysis a challenge. The main causes are imprecise address analysis and the inability to model and analyze data access patterns. Imprecision of the results may indicate that two memory accesses compete for the same cache set, while in reality they always go into different sets. The success of data cache analysis depends on whether *temporal and spatial locality* of data accesses can be precisely captured and analyzed.

3.5 Multi-Level Caches

Most modern processors adopt a multi-level cache design. Upon a memory access, the processor queries the memory hierarchy from the L1 cache down to main memory until the requested data or instruction is found. Regardless of the number of levels, the highest level cache is generally much faster than main memory, since the latter is accessed via the *off-chip* memory bus. To produce precise WCET estimations, cache analysis should be conducted all the way to the highest level, instead of merely on the L1 cache.


3.5.1 Separate vs. Integrated Approaches

Two major analysis frameworks for multi-level caches are the *separate analysis*, which analyzes caches level by level, and the *integrated analysis*, which deals with the cache hierarchy as a whole.

The separate analysis framework was first proposed by Mueller [87] and refined by Hardy and Puaud [55], who corrected a soundness problem. Figure 16 [55] shows the main work flow. In the analysis of all but the L1 cache, a key information is whether a data request actually leads to an access to this level. For example, if a memory access is predicted *always hit* at L1, then the

■ **Table 2** Computing CAC for level L and memory block r [55].

$CAC_{r,L-1} \backslash CHMC_{r,L-1}$	AM	AH	FM	NC
A	A	N	U	U
U	U	N	U	U
N	N	N	N	N



■ **Figure 17** The update function for U access [55].

L2 cache will not be visited. An interface across cache levels, called *Cache Access Classification* (CAC) is introduced to describe this information. The notion $CAC_{r,L}$ denotes the access property of block r to level L , which is evaluated to one of the following three cases:

- N (Never): the access to r is never performed at level L ;
- A (Always): the access to r is always performed at level L ;
- U (Uncertain): the access to r at level L can neither be excluded nor predicted.


The CAC values for level L are computed from both the Cache Hit/Miss Classification (CHMC) and the CAC for level $L-1$, which is shown in Table 2.

Trivially, if $CAC_{r,L-1} = A$ (or N), then the access to memory block r is always (or never) considered in the analysis of level L . However, handling the U classification requires special attention. In Mueller's multi-level analysis [87], memory accesses with U classification are "conservatively" treated as *always access* in the analysis of the current cache level. However, this treatment is demonstrated to be unsafe [55], since it may underestimate block ages, and thus incorrectly predict cache misses as hits. Hardy and Puaut corrected the problem by considering both possibilities for the U accesses in the update function (shown in Figure 17), which guarantees that the worst-case scenario is never missed.

However, Hardy and Puaut's analysis may suffer from precision problems. Note that in the above CAC computation, the FM classification is treated the same as NC, which means the information obtained by Persistence analysis at level $L-1$ is never leveraged in the analysis of level L . Actually, a block classified as FM at level $L-1$ causes at most one access to level L on its first access. Mueller in an earlier practice tried to solve this problem by unrolling the loop bodies [87], but this approach does not scale.

The component-wise separate analysis has several advantages. First, the analyzer has the flexibility to apply a different analysis method for each cache level, as long as the method produces hit/miss classifications as the interface across adjacent levels. This is desirable for architectures with different replacement policies for different cache levels, such as in the IBM Power 5 processor. Second, the overall analysis is scalable as long as the adopted single-level analysis is scalable.

However, the separate analysis may be pessimistic due to imprecise transfer of cache access information across cache levels. In contrast, *integrated analysis* [111] tries to build a holistic abstract domain for all cache levels, aiming to collect the information lost by the separate analysis.



■ **Figure 18** How a live cache helps to obtain a more precise join operation [111].

Consider a Must join operation of a separate analysis with a 2-way L1 cache and a 4-way L2 cache, shown in Figure 18. S_{iL_j} represents the abstract state S_i at Level j . Consider block x , which appears in $S1_{L2}$ on the left branch and in $S2_{L1}$ on the right branch. By a separate analysis, x does not appear in the joined state $S3$ of any level. If a subsequent access to x occurs, a cache hit can not be predicted. However, by evaluating both levels together, it can be seen that x does show up in *every incoming path*, so a subsequent access to x should be a cache hit, either in L1 or in L2 depending on the execution history. This is to say, x is guaranteed in the cache hierarchy at the joined program point. Unfortunately, this information is lost in the separate analysis.

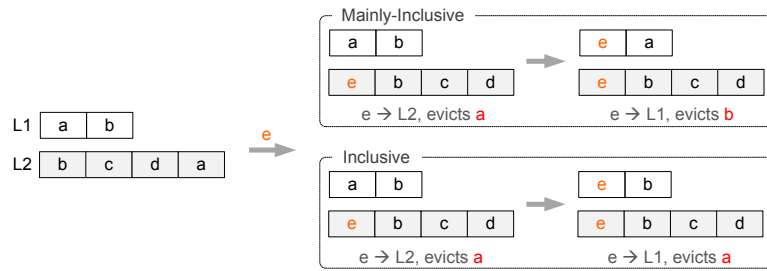
Based on this observation, Sondag and Rajan introduced a new component called *live cache* into the traditional abstract domains. At a join, a block is added to the live cache if it appears in some cache level in every input cache state, as depicted by $S3_{live}$ ⁵ in Figure 18. With live-cache information, one can now safely predict that x at least hits in the L2 cache. As reported in [111], the extra overhead by introducing live cache is acceptable for a 2-level cache hierarchy. However, analysis overhead increases with the number of cache levels, since an independent live cache is maintained for every pair of cache levels.

3.5.2 The Impact of Inclusiveness

The relationship between cache levels is a key design feature. In some processors, all data in level L must be contained in level $L+1$. Such caches are called (*strictly*) *inclusive* caches. For example in Figure 19, the access to e causes a to be evicted from L2, so a is forced to be removed from L1 to guarantee inclusion. Inclusive caches are favored in multi-cores: any data update in the shared L2 cache is automatically synchronized to the private L1 caches of all cores due to inclusion enforcement, which also achieves data coherency. Other processors adopt *exclusive* caches, in which data is guaranteed to be in at most one cache level. Exclusive caches are desirable for resource-limited systems since there is no data duplication in the cache hierarchy. The type adopted in previous discussions is called *mainly-inclusive*, which neither enforces inclusiveness nor exclusiveness. Inclusive caches are harder to analyze than exclusive caches, because the update to one cache level may cause further changes to both lower and higher cache levels. Such behavior may preclude the separate analysis flow.

Hardy et al. adapted the separate analysis [55], originally designed for mainly-inclusive caches, to both inclusive and exclusive caches [56]. The main idea is to first conduct an analysis assuming a mainly-inclusive cache, and to then modify the CHMC results to guarantee that the effects of cross-level cache updates are safely considered. For example in Figure 19, the analysis assuming

⁵ x has an older age in $S1_{L2}$ than in $S2_{L1}$. This is because Sondag's analysis assumes the write back policy. If x is evicted from the L1 cache, it is installed in the youngest position of the L2 cache. This means x can still suffer another $k - 1$ evictions in L2, where k is the cache associativity. The age of a block in the live cache describes how long it can stay in the *whole* cache hierarchy.



■ **Figure 19** Handling new behavior of inclusive and exclusive caches.

mainly-inclusive cache reports that a is AH at L1 but may be evicted from L2. To adapt this result to an inclusive cache, one must consider the possibility that a can be removed from the L1 cache due to an update in the L2 cache. As a result, a 's CHMC at L1 is modified from AH to NC. Similar problems exist in the May analysis as well. As a consequence to CAC computation, accesses to all cache levels except L1 are changed to Uncertain ($CAC_{r,L} = U$ for $L \geq 2$). All these modifications severely degrade the analysis precision.

Sondag and Rajan extended their integrated analysis to both inclusive and exclusive caches [111]. The resulting update and join functions for inclusive caches are very complex since once a block is accessed on some level L , the corresponding changes in other cache levels must be correctly considered. Analysis of exclusive caches has similar problems. To summarize, the inter-dependent updates among cache levels to enforce inclusion/exclusion brings new difficulties regardless of the analysis framework.

3.5.3 The Impact of Write Operations

A write to the cache occurs when a data variable receives a new value. Two levels of policies determine when and where to conduct the writing of data back to memory. The *write-through* policy requires that the new value is updated synchronously both in the cache and in main memory. In contrast, the *write-back* policy only marks the modified data as *dirty*, and performs the actual update of memory only when the data is evicted from the cache. A *write miss* occurs if the data to write are not in the cache. Under the *write allocate* policy, missed data are first loaded into the cache, and then updated with the new value, resulting in a cache miss followed by a cache hit. For the *non-write allocate* policy, the data are directly written to main memory, bypassing the caches.

The write-through policy is generally easy to handle in cache analysis, since data writes to a certain cache level incur no change to other cache levels. However, for the write-back policy, evicted dirty data are written to higher cache levels. Second, for the write allocate policy, a write operation always causes cache accesses regardless of hit or miss, which makes no difference from the read operation. However, for non-write allocate caches, a write miss never causes a cache access, so one cannot simply assume that each write operation changes the cache state, as is the case for reads. Like the inclusiveness enforcement, complex write operations cause cross-level cache updates, which is a challenge to the analysis.

Hardy's separate analysis framework has been extended to multi-level data caches by Lesage et al. [67] and to multi-level unified caches by Chattopadhyay and Roychoudhury [24]. However, both analyses assume a write-through policy. The abstract domains adopted in these methods are not able to handle write-back. Sondag and Rajan modeled write-back behavior in their integrated analysis [111]. It is shown that modeling write-back is easier by an integrated abstract domain, but one still has to carefully distinguish *possibly evicted blocks* from *definitely evicted blocks* at any level during the analysis to guarantee soundness. Definitely evicted blocks are identified from the

May information in Sondag’s method. Due to the access uncertainty of possibly evicted blocks, conservative update and join operations have to be employed, causing a loss in precision.

4 Analysis of Non-LRU Caches

In the past two decades, most research on cache analysis in the real-time domain was focused on the LRU replacement policy. The analysis of non-LRU replacement policies, i.e., those widely adopted in real-life processors, is still immature. In this section, we look into the challenges for non-LRU analysis and survey existing techniques.

4.1 Why Are Non-LRU Replacement Policies Hard to Analyze

To answer this question, we explore why it is hard to design *precise* and *efficient* abstract domains and the corresponding operations for non-LRU caches. We identify multiple challenges discussed in the following paragraphs.

4.1.1 Unsuitability of AH, AM, and FM Classifications

Under LRU replacement most memory accesses can be classified as AH, AM, or FM. Are these classifications equally suitable for non-LRU replacements? If not, what are alternatives that are better suited to characterize other policies’ behavior?

Unfortunately, these classifications are not as suitable for non-LRU replacements as they are for LRU. As shown in Guan et al.’s analysis [47], under FIFO replacement memory accesses may exhibit alternative hit and miss behavior so that none of the traditional classifications (AH, AM or FM) applies. Similarly, under MRU many cache accesses exhibit the *K-Miss* property [45]. An access classified as *K-Miss* suffers several misses (bounded by $K \leq \text{cache associativity}$) upon the first few accesses, and then persists in the cache. This kind of persistence property, however, is not captured by the FM classification. This demonstrates that one needs to better understand the specific cache behavior under different policies to come up with proper classifications.

4.1.2 Irregular and Non-Monotone Cache Update Behavior

The abstract domains and the corresponding transfer functions are designed to compute cache behavior invariants. An abstract domain is precise and efficient if (1) abstract states can compactly represent many concrete states, while preserving the information required for classification, and at the same time (2) transfer functions precisely capture the effect of a memory access on the concrete cache states. For example in the AI-based analyses for LRU, the block age bounds in the abstract states capture precisely the information required to classify blocks as cached or not. Further, this information can be precisely maintained by the transfer functions due to LRU’s regular cache update: a) whether or not a block’s age depends solely on its age relative to the accessed block’s age. Upper and lower bounds (in Must and May analyses) on the ages of blocks can be precisely updated due to the monotonicity of the operation, b) regardless of its previous age, and whether it was cached or not, the accessed block is always assigned the youngest age.

Unfortunately, most non-LRU replacement policies do not possess such monotone behavior. Take the FIFO replacement in Figure 4(b) for example. After a hit to d , d remains in the original position and is immediately evicted by the next access to x . The fact that d is *recently* accessed is not reflected by the update rules. An example of irregular behavior under MRU is shown in Figure 4(c). After f is installed into the cache, a subsequent hit to c followed by a miss to e evicts f out of the cache. However, block b , which is older than f , remains in the cache even after f is

evicted. The problem of PLRU is shown in Figure 4(d). In the state before a is accessed, the oldest block that will be evicted next is b . However, after a hit to a , the block to be evicted is changed to d .

To build efficient abstract domains for such replacement policies is very difficult. For example, in a Must analysis for FIFO [41], early determination of cache misses is helpful to better predict cache hits later. However, a very complex May analysis has to be designed to determine miss information as early as possible. For PLRU, a precise analysis must model the tree bits. The Must analysis for PLRU [40] by Grund employs a far more complex abstract domain, compared to LRU, to express information in the tree and predict cache hits.

4.1.3 The Influence of Initial States

Cache behavior heavily depends on execution history. In [104], it is shown that program performance under non-LRU replacement policies are very sensitive to the initial cache state, i.e., what remains in the cache before a program starts. This presents a challenge to obtain precise estimations. To illustrate the problem, assume currently m is accessed in a FIFO cache and we want to precisely estimate m 's lifetime. There are many possible situations: *Case 1*: m hits, but it has been in the cache for the longest time among the cached blocks, and thus it will be evicted upon the next cache miss. *Case 2*: the access to m is a miss and due to first-in, first-out behavior m will withstand another $k - 1$ (where k is cache associativity) cache misses without being evicted. Obviously, m 's remaining "life expectancy" in the two cases is rather different. If no knowledge on the initial cache states is available, a safe analysis (to predict hits) has to assume the worst case, i.e., *Case 1*. To be more precise, one may try to distinguish *Case 2* from *Case 1* by investigating whether the current access to m is a miss or not. Then, the analysis needs to know there are enough cache misses to evict any previously accessed m out of the cache, which again relies on the initial states. If a replacement policy can remove uncertainty from the initial states quickly, it will be easier to analyze.

To analytically model the effects of unknown initial states, Reineke et al. proposed a metric, *evict*, for a replacement policy [105], as listed in Table 3⁶. Intuitively, the value of *evict*(k) tells us after how long a sequence of pairwise different memory accesses, we can conclude that the cache only contains blocks from the access sequence, or, in other words, how long a sequence of pairwise different accesses is needed to evict all unknown cache contents from the cache.

The *evict*(k) results show that generally longer sequences have to be observed for non-LRU replacement policies. This property directly corresponds to the achievable precision by a May analysis to predict misses, and indirectly affects Must analysis, the precision of which partly depends on how much May information can be obtained during the Must analysis [41]. Similarly, the *fill* metric captures the number of pairwise different memory accesses required to reach a single cache state independently of the initial cache state. As can be seen in Table 3, the gap between LRU and other policies is even bigger for the *fill* metric.

The two metrics *evict* and *fill* discussed above relate to the precision of *classifying* analyses. In other work, Reineke and Grund [104] determined how strongly the *number* of cache misses may vary depending on the initial state, which is related to the precision of *bounding* analysis in the presence of uncertainty about the initial state. Their analysis demonstrates that the number of cache misses may vary strongly depending on the initial state under FIFO, PLRU, and MRU, while it may not vary much under LRU replacement. Further, it is shown that the empty cache

⁶ Table 3 extracts the HM case for *evict* and *fill* with $k > 2$ from the full results in [105], where k is cache associativity.

■ **Table 3** Predictability metrics [105].

Policy	$evict(k)$	$fill(k)$
LRU	k	k
FIFO	$2k - 1$	$3k - 1$
MRU	$2k - 2$	$3k - 4$
PLRU	$\frac{k}{2} \log_2 k + 1$	$\frac{k}{2} \log_2 k + k - 1$

■ **Table 4** Generalized predictability metrics [105].

Policy	$m_{ls}(k) = m_{ls}'(k)$	$evict'(k)$
LRU	k	k
FIFO	1	$2k - 1$
MRU	2	$2k - 2$
PLRU	$\log_2 k + 1$	∞

state is not necessarily the worst initial state for non-LRU policies. This presents severe problems for measurement-based WCET analysis approaches.

4.2 Predicting Cache Hits

The more hits can be predicted, the better the WCET bound. Must and Persistence analyses discussed earlier are used for this purpose. To predict a hit for a memory access to m , an analysis needs to ensure m has not been evicted since its last access. Intuitively, the more pairwise different blocks have been accessed since the last access to m , the higher the chance that m has been evicted from the cache. To capture this information, we introduce the following definition:

► **Definition 1** (Reuse Distance). Let p be a memory access sequence that ends with an access to memory block m . The *reuse distance*⁷ of m , denoted by $\mathbf{rd}_p(m)$, is the number of distinct blocks accessed along p since the previous access to m in p .


The notion of reuse distance coincides with the ages of memory blocks that are used in the analysis of LRU caches. For example, let $p_1 = \langle beabcd a \rangle$ and $p_2 = \langle abccdccba \rangle$, we have $\mathbf{rd}_{p_1}(a) = \mathbf{rd}_{p_2}(a) = 4$. Due to branches or input-dependent memory accesses, there can be multiple access sequences leading to the same access to block m in a program. So we define the *maximal reuse distance*, denoted by $\widehat{\mathbf{rd}}(m)$, as the maximal value of $\mathbf{rd}(m)$ over all possible memory access sequences leading to a particular access. We can evaluate analysis techniques by the maximal reuse distances for which they can predict cache hits.

Reineke et al. explored the *minimal life-span* [105] for different replacement policies, which is the minimal length of a sequence of pairwise different memory accesses necessary to evict a block that has just been accessed from the cache. The minimal life-span values are given in the $m_{ls}(k)$ column of Table 4. A slight variation of $m_{ls}(k)$ is $m_{ls}'(k)$, which considers the minimal number of pairwise different memory blocks required to evict a block that has just been accessed from the cache. Notice, the slight difference between the two notions: $m_{ls}(k)$ considers only sequences consisting of pairwise different accesses, whereas $m_{ls}'(k)$ allows multiple accesses to the

⁷ In the literature, this is also referred to as the (*LRU*) *stack distance* [86]. Also, note that the term reuse distance is used ambiguously in the literature, sometimes referring to the number of accesses since the previous access to m , and sometimes referring to the number of distinct blocks accessed since the previous access to m . We follow the latter notion.



■ **Figure 20** Levels to explore cache hits.



■ **Figure 21** Example programs.

same block. For all the considered policies, $mls(k)$ is equal to $mls'(k)$. The metric $evict'$, also listed in Table 4, will be discussed in Sec. 4.3. The $mls'(k)$ metric tells us how many of the most recently accessed blocks are guaranteed to be in the cache. By this result, a Must analysis can be constructed as follows: for any memory access to m , one can check if $\widehat{rd}(m) \leq mls'(k)$ holds for the given replacement policy. If yes, the memory access to m is guaranteed to be a hit. We say that such an analysis explores *Level I*, which is illustrated in Figure 20.

Notice that to predict cache hits, the Must analysis for LRU presented in Sec. 3.1 computes upper bounds on the maximal reuse distances of memory blocks. Similarly, the May analysis computes lower bounds on the minimal reuse distance of memory blocks to predict cache misses. As the notion of reuse distances is replacement policy-independent, these LRU Must analysis can thus safely be reused to predict hits for other policies, by relying on the policies value of $mls'(k)$.

However, the $mls'(k)$ values for non-LRU replacements are commonly small compared to cache associativity k , because they consider worst-case scenarios. In practice, it is unlikely that the worst case occurs at every program point. Thus, analyses tailored to a particular replacement policy can often go beyond *Level 1* in predicting hits.

For a program that can fit into the cache of size k , there is a strong intuition that each block of the program eventually persists in the cache, i.e., after some misses, the remaining accesses to each block are definitely cache hits. For such programs, the maximal reuse distance of any access is no larger than k (*Level II* in Figure 20). This property is attractive as it enables an efficient Persistence analysis that simply collects the set of different blocks accessed by a program. However, such a Persistence analysis is not correct for every replacement policy: it works for LRU, MRU and FIFO, but not for PLRU.

Figure 22 illustrates the cache state transitions when the loop in Figure 21(a) is executed, alternating between the three branches in the loop body on a 4-way PLRU cache. Each time a is accessed, the root bit points to the right subtree, so b , c , and d have to compete for the two cache lines on the right. Even though the loop can fit into the 4-way cache set, only block a is persistent. This example unveils a negative property of PLRU: it does not always make use of all its capacity [14].

It is crucial to investigate what process a block may go through before it finally persists in the cache. This is important to safely bound the number of misses that may occur.

For MRU, Guan et al.'s results [45] show that if for a block m , $\widehat{rd}(m) \leq k$ holds, m will eventually persist in the cache. However, m may suffer more than one miss before reaching the




Figure 22 An example that demonstrates that PLRU does not always use the cache’s entire capacity.

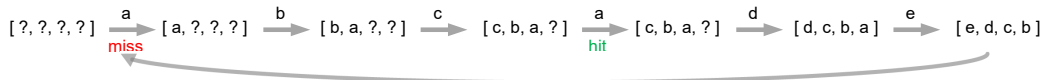



Figure 23 Alternative hit and miss behavior on FIFO.

stable state. The result is strong in that bounds on the number of misses are determined for all reuse distances in *Level II* for MRU. Consider the program in Figure 21(b): even if the program cannot fit into a 4-way MRU cache, block *a*’s number of misses can be bounded by a constant since $\widehat{\text{rd}}(a) \leq 4$.


For FIFO replacement, Grund and Reineke [42] show that if a loop entirely fits into the cache, each block suffers at most one miss and then persists in the cache⁸; otherwise, no guarantee is given. Guan et al. further explored this problem, and found that if a block *m* satisfies $\widehat{\text{rd}}(m) \leq k$, then even though *m* is not eventually persistent, it is still guaranteed to enjoy cache hits, which can be expressed by a bound on the number of misses that accesses to *m* may suffer [47]. For example in Figure 21(b), $\widehat{\text{rd}}(a) \leq 4$ holds for a 4-way FIFO cache. In the worst case, the loop alternatively takes the two branches, and *a* may suffer cache misses repeatedly. To evict *a* from the cache, both branches have to be taken, which causes *a* to enjoy a cache hit in the execution of one of the branches (shown in Figure 23). It can be shown that *a* suffers at most $\lfloor \frac{1}{2} \cdot x \rfloor + y$ misses, where *x* is the execution count of *a*, and *y* is the total number of times the loop is entered.

The only analysis to predict hits for PLRU for maximal reuse distances in *Level II* is a Must analysis proposed in [43]. The analysis presented is based on the following observation: to evict a block with the fewest possible accesses to distinct memory blocks, the three bits (assuming an associativity of 8) that are on the path from a cache line to the root of the tree need to be flipped in a particular order, namely from the bottom to the top. This is illustrated in Figure 24 for block *m*. Notice that flipping bits near the root before flipping all bits closer to the leaves does not contribute to evicting *m*, as these bits will eventually be flipped back before evicting *m*. The basic idea behind the analysis in [43] is to track two properties: a) the number of bits that already point towards a block (counting from the leaf of the tree), and b) the so-called “sub-tree distance” between pairs of blocks. The sub-tree distance between *a* and *b* captures which bits on the path from *a* to the root an access to *b* may flip. By analyzing these key properties, it is sometimes possible to predict that a block stays in the cache even if more than $\text{mls}'_{PLRU}(k) = \log_2 k + 1$ other blocks have been accessed.

⁸ Note that blocks do not necessarily encounter their misses in the first loop iteration. It may take several iterations for all the blocks to stabilize in the cache.



■ **Figure 24** A scenario in which block m is evicted with the minimal $mls'(k) = \log_2 k + 1$ accesses.



■ **Figure 25** Cache behavior under MRU in *Level III*.

To go beyond *Level II* means to explore whether blocks with maximal reuse distance larger than k still have cache hits. One needs to first show that the above fact does occur for some replacement policy, and second propose an analysis to discover the cache hits. Take MRU for example, Figure 25 shows that even if we have $\widehat{\mathbf{rd}}(a) = \widehat{\mathbf{rd}}(d) > 4$ for a 4-way cache, accessing a and d is always hit. But this phenomenon relies on the initial state at the entry of the loop. To explore such behavior, the abstract domain must be able to preserve very detailed information on cache states. This requirement makes it very hard to explore cache hits in *Level III* by abstract analysis methods. The abstractions introduced by Grund and Reineke for FIFO [41] are in principle able to predict *Level III* hits, however, they usually require a highly context-sensitive analysis to do so. *Level III* ends at $evict'(k)$, after which no more hits are possible.

4.3 Predicting Cache Misses

Predicting more misses tightens the estimated BCET, but it can also indirectly help with the analyses for some non-LRU replacement policies to predict more hits [41]. A concrete example is the case of FIFO explained in the discussion of the influence of initial states in Sec. 4.1. Furthermore, for multi-level cache analysis, predicting more misses for level L reduces the uncertainty of cache accesses on level $L + 1$, which leads to more precise overall estimations. Lastly, in micro-architectures with timing anomalies, if a memory access cannot be classified as a cache hit, both the cache hit and the cache miss cases need to be explored. Predicting cache misses may in such cases drastically reduce analysis times, as it allows to explore only the cache miss case.

To predict cache misses requires to show that memory blocks are not in the cache right before they are being accessed. Thus a May analysis, i.e., an analysis that overapproximates cache contents, is required to safely predict cache misses. May analyses can be constructed based on a variation of the $evict$ metric [105], which we denote by $evict'$. The difference between $evict$ and $evict'$ is the same as the difference between mls and mls' : any sequence s containing $evict'(k)$ distinct memory blocks is guaranteed to evict any prior cache contents not contained in the sequence s . In contrast, $evict$ refers only to sequences that never access the same memory block twice. Values of $evict'$ for common policies are listed in Table 4. For example, for FIFO, $evict'(k) = 2k - 1$. This means that after accesses to $2k - 1$ pairwise different blocks, the cache only contains elements from the accessed sequence. Then, the May analysis only needs to observe a sequence with $2k - 1$ pairwise different blocks other than m precluding the current access to m . On the other hand, for PLRU, $evict'(k) = \infty$. In other words, there are sequences of memory

accesses containing an arbitrary number of distinct memory blocks that do not evict all prior cache contents. We have seen an example of such a sequence in the previous section, which is illustrated in Figure 22.

May analyses based on *evict'* can be constructed by determining *lower bounds* on the reuse distances of memory blocks. The LRU May analysis presented in Sec. 3.1 does exactly that.

The *evict* metric suggests that less than $evict(k)$ accesses to pairwise different memory blocks do not allow to predict any misses, thereby constituting a limit on how much information a May analysis can obtain. However, this conclusion is built on the assumptions that the initial state is *completely unknown*, and that the access sequence consists of *pairwise different* memory accesses. There is thus hope that by tailoring an abstract domain to a specific replacement policy, more precise May analyses can be achieved. So far, such abstractions have only been built for the FIFO replacement policy [40, 41]. Due to limited space, we only explain the main intuitions and the key constituents of the two existing FIFO abstract domains.

Consider a 4-way FIFO cache and the access sequence $x \circ s \circ x$, where the first access to x is a miss and installs x into the “first-in” cache way, and then a sequence s that does not contain x is accessed followed by another access to x . To predict a miss for the second access to x , it suffices to check whether either of the following two properties holds:

- *Property 1*: The accesses in s result in at least 4 misses;
- *Property 2*: Before the second access to x , every cache way is occupied by a memory block from s .

The abstract domain proposed in [41], which we call $FIFO^\alpha$, checks *Property 1*. The key information to be maintained in the abstract domain is the number of *definite misses* after the first access to x , denoted by $dm(x)$. When $dm(x)$ reaches 4 during the analysis, one can predict misses for a future access to x . To better maintain definite misses, the number of *cache ways covered* by blocks accessed after the last access to x is maintained as auxiliary information.

A disadvantage of $FIFO^\alpha$ is that it only starts to predict misses after $2k - 1$ pairwise different memory blocks have been accessed, which is in line with the *evict* metric. Therefore, Grund and Reineke proposed another FIFO domain, which we denote by $FIFO^\beta$, so that cache misses can be predicted even if fewer than $2k - 1$ pairwise different blocks are accessed between two different accesses to the same block [40, 42]. The domain $FIFO^\beta$ checks *Property 2* to predict cache misses. For the above example, it explores if memory blocks accessed in sequence s eventually cover all the 4 cache ways. The exploration is based on a more powerful result (Lemma 4 in [42]):

If a sequence s contains l distinct blocks, then $l - k + 1$ cache ways must be occupied by the contents of s , regardless of the initial cache state.

Importantly, the effect of consecutive sequences adds up. For example, let $s = s_1 \circ s_2 = \langle a, b, c, d, e \rangle \circ \langle a, b, c, d, e \rangle$. The accesses to s_1 cover the $5 - 4 + 1 = 2$ most-recently-used ways in the cache set. Similarly, the accesses to s_2 contribute another $5 - 4 + 1 = 2$ to the covered positions. Then we can guarantee that access to s finally covered all the 4 cache ways⁹. This means the execution of s actually evicts x out of the cache and a miss on the second access to x can safely be predicted.

So far, no May analysis is known for PLRU. For MRU the best known May analysis is based on *evict'*. Precisely predicting misses for these two policies is still a challenge.

⁹ The effectiveness of this analysis depends on how a long sequence is partitioned. Grund has a systematic method to explore different partitionings for optimization [40].

4.4 The Relative Competitiveness Framework

Besides the above research, Reineke and Grund proposed the Relative Competitiveness framework [103] that allows to translate analysis results for one replacement policy to another policy. The promise is then to apply known LRU analyses to non-LRU caches.

A policy P is (k, c) -hit-competitive relative to policy Q if the number of cache hits $h_P(s)$ of P on sequence s is bounded from *below* by the number of cache hits $h_Q(s)$ of Q as follows: $h_P(s) \geq k \cdot h_Q(s) - c$. Similarly, a policy P is (k, c) -miss-competitive relative to policy Q if the number of cache misses $m_P(s)$ of P on sequence s is bounded from *above* by the number of cache misses $m_Q(s)$ of Q as follows: $m_P(s) \leq k \cdot m_Q(s) + c$.

By monotonicity of the two inequalities, they can also be applied to lower bounds on the number of hits and upper bounds on the number of misses: For example, given a lower bound on the number of hits of Q using hit-competitiveness a lower bound on the number of hits of P can be derived.

For $(k, c) = (1, 0)$ the notions of (k, c) -hit- and miss-competitiveness coincide. In this case, P “dominates” Q . In other words, P never incurs more misses than Q . In such a case we simply say that P is $(1, 0)$ -competitive relative to Q . Then, a Must analysis for Q is a valid Must analysis for P ; conversely, a May Analysis for Q is a valid May Analysis for P .

In [103] it is shown how to automatically compute the best values for (k, c) such that policy P is (k, c) -hit/miss-competitive relative to policy Q , for fixed associativities of the two policies. Depending on the similarity of P and Q this computation scales to associativities between 8 and 256.

The most interesting cases are those in which either P or Q is LRU, as precise analyses for LRU are known. Examples for hit-competitiveness results derived in this way are [101, 103]:

- An 8-way FIFO cache is $(\frac{1}{2}, \frac{7}{2})$ -hit-competitive relative to an 8-way LRU cache.
- An 8-way FIFO cache is $(\frac{2}{3}, 2)$ -hit-competitive relative to an 4-way LRU cache.
- An 8-way PLRU cache is $(\frac{1}{2}, \frac{3}{2})$ -hit-competitive relative to an 6-way LRU cache.
- An 8-way MRU cache is $(\frac{2}{3}, \frac{4}{3})$ -hit-competitive relative to a 4-way LRU cache.

To use this relation, assume 100 hits are predicted for a program on an 4-way LRU cache, then $\lceil \frac{2}{3} \times 100 - 2 \rceil = 65$ hits are guaranteed on an 8-way FIFO cache.


The metrics $mls'(k)$ and $evict'(k)$ are strongly related to $(1, 0)$ -competitiveness relative to LRU. In particular, let $mls'_P(k)$ and $evict'_P(k)$ denote the values of the two metrics under policy P . Then, P is $(1, 0)$ -competitive relative to $LRU(mls'_P(k))$ and $LRU(evict'_P(k))$ is $(1, 0)$ -competitive relative to P . For example:

- $LRU(2k - 1)$ is $(1, 0)$ -competitive relative to $FIFO(k)$.
- $LRU(2k - 2)$ is $(1, 0)$ -competitive relative to $MRU(k)$.
- $PLRU(k)$ is $(1, 0)$ -competitive relative to $LRU(\log_2 k + 1)$.

Cache analyses based on relative competitiveness can be pessimistic, because the relation holds for *any* possible workload. Moreover, the framework provides bounds on hits (or misses) for the whole program or alternatively program fragments rather than classifying independent memory access, except for the case of $(1, 0)$ -competitiveness. This makes it difficult to apply the approach in multi-level cache analysis, or in integrated analyses considering both caches and pipelines.

5 Execution Environments

Discussions so far have focussed on analyzing an independent program. Cache analysis is severely challenged in the presence of complex execution environment, such as multi-tasking systems or shared-cache multi-cores, where extra time delay due to interference on caches from other co-scheduled/running programs must be taken into account.



■ **Figure 26** The separate CRPD analysis framework.

5.1 Cache-Related Preemption Delay

An essential feature of real-time systems is preemption, which allows a higher priority task to preempt a lower priority task so that the higher priority one meets its deadline. However, preemptions may lead to extra cache misses: the execution of the preempting task may alter the cache state, so that once resumed, the preempted task needs to bring data back into the cache that was evicted as a consequence of the preemption. The extra delay due to cache reloading is commonly referred to as the *Cache-Related Preemption Delay* (CRPD). Empirical results [74] show that CRPD contributes significantly to the execution time, so it must be precisely estimated to obtain tight estimations of response times. Furthermore, it has also been shown that with CRPD, the synchronous release of all higher priority tasks does not represent the critical instance of single-core preemptive scheduling [134]. Clearly, preemptions introduce a new dimension of complexity into timing analysis.

The most intensively studied framework is *separate* CRPD analysis, in which the CRPD is treated as a separate overhead rather than as a part of the WCET of the preempted task. To bound the CRPD under LRU replacement, two approaches have been proposed, which are illustrated in Figure 26:

1. By analyzing the preempted task [2, 66, 91, 119, 113]: Additional misses can only occur for *useful cache blocks* (UCBs), i.e., blocks that may be cached and that may be reused later, resulting in cache hits. For LRU, the number of such UCBs is a bound on the number of additional misses due to preemptions. Static analyses have been proposed to safely approximate the set of UCBs.
2. By analyzing the preempting task [91, 113, 119, 121]: The preempting task may only cause additional cache misses in those cache sets that it modifies. Thus, analyses to compute bounds on the number of *evicting cache blocks* (ECBs) have been developed. A memory block is an ECB if it may be accessed during the preempting task’s execution. However, for set-associative caches, approaches based purely on ECBs have so far been either imprecise [17] or unsound [119], as shown in [17].

The CRPD is computed as the total time delay of all preemption-related cache misses. The final step is to take into account the computed CRPD bounds in a schedulability analysis framework, so that the Worst-Case Response Time (WCRT) of the preempted task can be obtained. All such approaches assume *timing compositionality* [52] so that the cost of additional cache misses can be accounted for separately.

5.1.1 Computing Useful and Evicting Cache Blocks

Since a preemption must happen before some instruction, here we first consider what happens at a particular program point. Most existing work adopts the UCB definition in [66]. A cache block m is useful at a given program point p , if:

1. m may be cached at p ;
2. m may be reused at some program point reachable from p without being evicted along the corresponding path.

To determine memory blocks that satisfy Condition (1), one needs to collect the set of blocks that may be cached by any possible program path from the starting point of the CFG to p , referred to as Reaching Cache Blocks (RCBs) and denoted by RCB_p . This corresponds to a May analysis as discussed in Sec. 3.1. To determine memory blocks that satisfy Condition (2), a set of Live Cache Blocks (LCBs), denoted by LCB_p , is computed similarly to RCB_p , however, by a *backward analysis*. Then, an overapproximation of the set of useful cache blocks at point p , UCB_p is obtained by the intersection of RCB_p and LCB_p . A bound on the CRPD is then obtained by taking the maximum size of UCB_p over all program points.

For direct-mapped caches, two major techniques exist: set-based analysis [66] and state-based analysis [91]. Both techniques rely on dataflow analyses to collect the RCBs and LCBs at each program point. State-based analysis maintains all possible concrete cache states at a program point. The analysis is precise, but does not scale to large programs. In contrast, set-based analysis maintains one abstract state at each program point, which collects the set of all possible cached blocks for each cache line. Staschulat and Ernst [113] proposed a scalable precision analysis that presents a trade-off between the above two analyses. The main idea is to pose a bound on the number of cache states maintained at each program point. Whenever the number of states goes beyond the limit, cache states are merged.

Regarding the analysis of the preempting task, note that (a) what matters is the size of the set of evicting cache blocks and not its actual contents, and (b) sizes that exceed the associativity of the cache do not have to be distinguished, as they will evict all prior cache contents anyway. For those reasons bounds on the number of ECBs can be obtained from bounds on the number of reaching cache blocks at the end of program execution, i.e., RCB_{end} .

5.1.2 CRPD Computation for Direct-Mapped Caches


For direct-mapped caches, the CRPD can be estimated by only considering the preempted task, which pessimistically assumes that each UCB of the preempted task could be evicted by the preempting task [66]. These techniques are classified as the UCB-Only approach by [4]. The CRPD can also be computed by only considering the preempting tasks [20, 121], which assumes any ECB of a preempting task may cause a preemption related cache miss (ECB-Only by [4]). Clearly, more precise CRPD can be computed by evaluating both the preempting and the preempted tasks. Specifically, the ECB-Only approaches have been improved by considering the preempted tasks, resulting in the UCB-Union class [118]; similarly, the UCB-Only approaches have been extended into the ECB-Union class [4].

Schedulability analysis needs to take the CRPD into account. Consider a widely adopted schedulability analysis [7] shown in Equation (5), where R_i is the response time, C_i is the WCET of a task, and T_j is the activation period. Equation (5) can be interpreted in the following way: the preemption cost of task τ_i preempted by τ_j , denoted by $\gamma_{i,j}$, is seen as an *extra* part of the execution time of the *preempting task* τ_j .

$$R_i = C_i + \sum_{\forall j \in hp(i)} \left\lceil \frac{R_i}{T_j} \right\rceil (C_j + \gamma_{i,j}). \quad (5)$$

In the presence of nested preemptions, as shown in Figure 27, the response time of τ_3 includes both the indirect cost of τ_1 preempting τ_2 (γ_a) and the direct cost of τ_2 preempting τ_3 (γ_b). The main problem is how to safely account for γ_a . Actually, γ_a can be considered in $\gamma_{3,1}$. Note that γ_a may be larger than γ_b , so a safe $\gamma_{3,1}$ needs to account for the maximal cost of τ_1 preempting any lower priority task, however not lower than τ_3 . Note that the ECB-Only approaches do not suffer from such nested preemption problems since they do not consider the preempted task.

A disadvantage of analyses by Equation (5) is: the worst-case delay $\gamma_{i,j}$ is always assumed for each preemption (τ_j preempting τ_i). As a result, some cache evictions can be included multiple



■ **Figure 27** An example of nested preemptions.



■ **Figure 28** An example of reordered misses.

times. To reduce this pessimism, other approaches [4, 114] adopted the schedulability test of Equation (6) instead, which evaluates the total cost of multiple preemptions of τ_j preempting τ_i as a whole. The computation of $\gamma_{i,j}^{sta}$ differentiates preemption scenarios, and thus can avoid unnecessary inclusion of cache evictions.

$$R_i = C_i + \sum_{\forall j \in hp(i)} \left(\left\lceil \frac{R_i}{T_j} \right\rceil C_j + \gamma_{i,j}^{sta} \right). \quad (6)$$

Altmeyer et al. provided a detailed classification of different approaches to bound the CRPD for direct-mapped caches and their relationship [4].


5.1.3 CRPD Computation for Set-Associative Caches

The computation of UCBs and ECBs can be solved by existing May analyses for set-associative caches. The main challenge is how to precisely and safely compute the “intersection” between the sets of UCBs and ECBs.

Let us first discuss a safety problem, given LRU replacement. Consider the case in Figure 28. Blocks a , b , c and d are all useful blocks. A preemption installs x into the cache set and thereby evicts a . The subsequent accesses of the preempted task to a , b , c and d are all cache misses even though the preempting task only evicted one cache block. This illustrates that there are two types of context-switch misses [17, 74]. The miss to a is a *replaced miss*, as a direct result of the preemption. In contrast, the misses to b , c and d are an indirect result of *reordering* of blocks by the LRU replacement policy. This example shows that even a single ECB can lead to a chain of misses to multiple UCBs, which cannot happen for direct-mapped caches. An example of an unsafe analysis is [118], which overlooked reordered misses.

One way to cope with this problem was proposed in [17]. As soon as there is a single ECB that maps to a particular cache set, all the UCBs that map to the same cache set are assumed to contribute to context-switch misses. This is obviously conservative, and it can be improved by obtaining more detailed information about the useful cache blocks. This information is captured by the notion of *resilience* introduced by [5].

The resilience $res(m)$ of a useful cache block m is the amount of “disturbance”, i.e. its ECBs, by a preempting task that the block may endure before becoming useless to the preempted task. Consider a useful cache block m for an 8-way LRU cache in Figure 29, where all blocks are mapped to the same cache set. The maximal age of m before its second access is 5. If the program is preempted at any point between the two accesses to m , for example at program point p , m will not be evicted from the cache as long as at most 3 ECBs from the preempting task map to the same set. So m ’s resilience is 3.



■ **Figure 29** The notion of resilience.

By computing lower bounds on the resilience of useful cache blocks, one can exclude many cache misses compared with the conservative assumption in [17]. However, nested preemptions must be very carefully handled. The ECBs from nested preempting tasks may accumulate to age a useful block. In this case, the ECBs of all possible preempting tasks must be considered, which may adversely introduce some pessimism.

The problem of reordered misses is rooted in LRU. A new policy called Selfish-LRU [102] has been proposed to eliminate reordered misses. The idea is to first evict cache blocks that do not belong to the currently active task.

For other replacement policies, such as FIFO and PLRU, the number of additional misses can even be greater than the number of UCBs, the number of cache ways, and the number of ECBs [17]. This makes it difficult to obtain precise CRPD bounds for these policies. An approach based on relative competitiveness [103] was sketched in [17] that allows bounding the total misses (intra- and inter-task misses) of a non-LRU policy from the results of LRU. Due to the generic nature of the relative competitiveness framework, the analysis results can be imprecise.

5.1.4 Other Analyses

It has been observed [2] that some pessimism is introduced by independently computing bounds on the CRPD and on the WCET. Consider the treatment of memory accesses to blocks that have been classified as useful cache blocks during the WCET analysis. If such accesses cannot be guaranteed to result in cache hits, a sound WCET analysis will also cover the cache miss case. However, in that case, while a preemption-related cache miss may occur in reality, it has already been accounted for in the computed WCET bound. Motivated by this observation, a notion of *definitely-cached* UCBs has been proposed in [2], which excludes such blocks from the CRPD computation. Excluding such blocks from the CRPD computation had previously been proposed by Schneider [107] in what he calls the “isolated method”. This approach relies on a coupling of WCET and CRPD analysis and may improve precision significantly.

Ramaprasad and Mueller [98, 100] presented an approach to response-time analysis for strictly-periodic task sets under fixed-priority scheduling, taking into account the CRPD in data caches. Due to their assumption of periodic tasks, they are able to simulate the scheduler during a hyperperiod of the system. Taking into account BCET and WCET estimates they can then accurately predict the number of preemptions, and to some extent even the preemption points within each job, which are taken into account in the CRPD analysis. In addition to the restriction to periodic task sets, this work shares the limitations of the data-cache analysis framework [97] that it builds upon, i.e., neither input-dependent memory accesses nor input-dependent control flows are supported. This approach was later extended to support a single non-preemptive region within each task [99].

CRPD analysis under dynamic priority scheduling has also been studied [63, 79]. The difference lies in the CRPD calculation for different schedulability tests of new scheduling policies. Lunniss et al. [78] compared the effectiveness of fixed priority scheduling and EDF in the presence of CRPD. Phavorin et al. [96] showed that common assumptions about optimality and sustainability of scheduling algorithms do not hold anymore, once CRPD is taken into account.

An alternative to bounding the cost of individual preemptions and taking this cost into account within schedulability analysis is to conservatively account for all possible preemptions within WCET analysis. This was first proposed by Schneider [107] in what he calls the “integrated method”. The advantage of such an approach is that it can take into account overlapping of memory latencies and computations in pipelined processors. This advantage, however, is usually outweighed by overestimating the number of additional misses that may occur, as an unbounded number of preemptions needs to be taken into account.

5.1.5 Limiting Preemptions


We have focused on approaches to bound the CRPD under fully-preemptive scheduling. Various mechanisms exist to reduce the number of preemptions, and thereby also the overheads introduced by preemptions. In the extreme case, tasks are scheduled non-preemptively, eliminating preemption cost entirely. However, under non-preemptive scheduling, long-running lower priority tasks often render task sets unschedulable. Prominent approaches that represent a compromise between the extremes of fully-preemptive and non-preemptive scheduling are *preemption thresholds* [125], *cooperative scheduling* [19], and *floating non-preemptive regions* [13].

Under fixed-priority scheduling with preemption thresholds, in addition to its priority, a task is associated with a preemption threshold. When a task is running, it can only be preempted by tasks whose priority is higher than the preemption threshold of the running task. This has been shown to sometimes improve schedulability and reduce the number of preemptions [125]. Schedulability analysis taking into account CRPD under preemption thresholds [16] relies on the same basic cache analysis concepts, i.e., UCBs and ECBs, as schedulability analyses under standard fixed-priority preemptive scheduling.

Under cooperative scheduling, each task allows the scheduler to preempt it at *fixed preemption points*, i.e., program locations at which it yields to the scheduler. Between preemption points, a task is executed non-preemptively. This introduces lower-priority blocking time, which may reduce schedulability, and requires analysis of the maximum blocking time [3]. The CRPD analyses described above, based on UCBs, ECBs, and resilience can be applied to fixed preemption points in a straightforward manner. A challenge is to select preemption points in a way that maximizes a task set’s schedulability. Fewer preemption points may result in lower CRPD but may increase the maximum blocking time. For programs consisting of straight-line code only, Bertogna et al. [15] provided an algorithm to optimally select preemption points. They assumed that the WCET is inflated accounting for preemptions at every preemption point by arbitrary preempting tasks, rather than accounting for the effect of preemptions within response-time analysis based on separate CRPD bounds.

In the floating non-preemptive region scheduling model [13], preemptions are limited as follows: When the highest priority task is executing, and a job of a higher priority task is released, the running job is not immediately preempted. Instead, a floating non-preemptive region of fixed length Q_i commences, where Q_i may depend on the task τ_i currently running. The currently running job is preempted after Q_i time units, unless it completes before. Again, CRPD analyses based on UCBs, ECBs, and resilience can be applied in this setting. However, as in the case of fixed preemption points, it is also possible to derive WCET bounds for all tasks that hold under the assumption that any two preemptions of a task are separated by at least Q_i time units. Given such WCET bounds, schedulability analyses do not need to further account for preemption costs. Marinho et al. [83, 84] provided analyses to bound a task’s WCET including CRPD taking into account the length of floating non-preemptive regions. Their algorithm consists of two phases:

1. First, a so-called preemption delay function f is computed. For any time t , $f(t)$ is a bound on the CRPD if the task is preempted after t time units. In order to compute f , they determine



■ **Figure 30** A common shared cache design in multi-cores.

a bound on the CRPD for each program point based on UCBs. Given lower and upper bounds on the execution time of each basic block, they determine lower and upper bounds on when a particular basic block may be running. Combining the two values yields f .

2. Then they incrementally determine based on f , one preemption at a time, the minimal progress over time in the execution of a program, when it can be preempted each Q_i time units.

If floating non-preemptive regions are short such an approach may be pessimistic as it does not take into account how often and by which tasks a task may actually be preempted. On the other hand, as in the “integrated method” of Schneider [107], it may take into account overlapping of memory latencies and computations in pipelined processors. To our knowledge, the two approaches have not been experimentally compared in the limited preemption context. This also applies to Bertogna et al. [15].

A comprehensive survey of the literature on cache-related preemption delays is given by Phavorin and Richard [95].

5.2 Shared Caches in Multi-Cores

Nowadays, multiple processing cores are deployed on a single die to fully exploit the real estate of the processor chip and to achieve high performance with low power consumption. A commonality among modern multi-core processors is the sharing of on-chip resources among multiple cores, such as the last-level cache (Figure 30), so that each core can potentially make use of the entire resource. However, tasks running in parallel on different cores compete for the shared resource, resulting in *inter-core conflicts*, also referred to as *inter-core interference*. Due to this interference, the execution time of a program now also depends on the resource-access behavior of the tasks running in parallel [133].

Inter-core interference on a shared cache is different from inter-task interference due to preemption. First, in a single-core preemptive system, a higher priority task does not suffer from interference by a lower priority task, while in multi-core systems, all tasks running in parallel on different cores interfere with each other independently of their priority level. Second, in a single-core preemptive system, a task can only suffer from interference by preempting tasks a small number of times, no more than the total number of releases of the higher priority tasks; in contrast, on multi-cores, interference on a shared cache may come between any two consecutive cache accesses of a task. Precisely analyzing all possible interleaving cache accesses on a shared cache is notoriously difficult due to the huge number of cases to consider.

One approach is to extend the AI-based analysis to take into account the interference on the shared cache [73]. The basic idea is similar to the resilience analysis in CRPD analysis. Assume that two tasks, A and B , run in parallel and share a k -way L2 cache. To estimate A 's WCET, multi-level analyses for A are first conducted without considering the interference from task B . Then, a second step analyzes task B to see whether its interference could cause the blocks of A

that are guaranteed to hit when A runs in isolation, to be evicted from the cache. Consider a block m of A : its maximal age, $age(m)$, in the cache can be extracted from a Must analysis. Then task B is analyzed to determine a bound on the number of interfering memory blocks that map to the same cache set as m , denoted by \mathcal{M} . If $\mathcal{M} \leq k - age(m)$ holds, m will remain in the cache even in the presence of B 's interference. Otherwise, m could be evicted from the cache due to B 's execution, and its classification needs to be changed accordingly.

A major drawback of the above approach is that the timing of cache conflicts is not considered, i.e., all potential cache conflicts computed from cache mapping are included. However, if by some means we know that the lifetimes of two conflicting tasks (or cache accesses) do not overlap, some cache conflicts can be safely excluded. This is a key property to tighten the estimations. Zhang and Yan proposed a technique to exclude infeasible conflicts by exploring conflicting pairs of cache accesses [135]. Liang et al. in [73] explore the overlapping of the lifetimes of co-running programs. The timing of cache conflicts can also be precisely captured by model checking. Gustavsson et al. used timed automata to model the behavior of programs on shared caches [50]. Infeasible conflicts can be precisely excluded when the UPPAAL model checker explores the system model. However, due to state space explosion, model checking based analysis can hardly scale beyond 2 cores. Another model checking based method was proposed in [132]. The SPIN model checker was adopted to exclude the infeasible cache conflicts, but the models did not explore the exact timing of the cache conflicts.

Another major analysis obstacle is that uncertainty, introduced by particular analysis technique or inherent to a hardware feature, may be amplified in the presence of shared caches. For example, in AI-based analyses, pessimistic age prediction makes the blocks of the interfered task less resilient; similarly, pessimistic age prediction for the interfering task leads to an overestimated number of conflicting blocks. Another source of uncertainty is the separation of cache behavior analysis and path analysis [120], which adversely introduces “architecturally-infeasible” paths. Pruning such infeasible paths can help to tighten WCET estimations. Banerjee et al. proposed a finer-grained abstract domain, which associates path information into the traditional Must and May abstract states to exclude non-existent cache states due to infeasible paths [12]. Chattopadhyay and Roychoudhury proposed another technique that improves the prediction for NC blocks by excluding infeasible paths using model checking [25]. Both techniques can be integrated into the analysis framework of [73] to more precisely estimate shared cache interference.

Even with the above techniques, real-time system design still faces a problem: if the shared cache is freely used, the worst-case performance of the tasks also degrades. Therefore, recent research tried to employ mechanisms that provide temporal isolation on shared caches, which both simplifies cache analysis and at the same time reduces the WCET. Cache partitioning [75, 116, 122] partitions the cache space among tasks by controlling page allocation to completely avoid cache conflicts among tasks on different cores. Cache locking [75, 116] locks the frequently used data in the cache so that hit/miss behavior is totally predictable. Another approach [54] tries to bypass the shared cache upon accesses to memory blocks with little reuse, which reduces cache interference. The idea of bypassing [54] was later extended to shared data caches [68], in which two heuristics are introduced to bypass indeterministic data references. On the system level, some further issues have to be solved. In multi-tasking systems, different tasks may try to lock the same cache segment, so scheduling of the lockings must be considered [126]. Regarding cache partitioning, the partitions assigned to the tasks may overlap in cache space. A task can only start execution if both the CPU and the cache partition are available. The schedulability tests must consider both the CPU and the cache constraints [46]. However, partitioning and locking have a side-effect of reducing the cache space available for each task. New techniques are expected for more intelligent resource allocation and arbitration, so that the WCET of the tasks can

■ **Table 5** WCET analysis tools supporting static cache analysis.

Tools	Instruction Cache	Data Cache	Multi-Level Cache	Non-LRU Cache	CRPD	Shared Cache
aiT	AI	AI		Pseudo-RR, PLRU, FIFO		
OTAWA	AI, ML-PER					
Chronos		Scope-Aware	Separate Framework			
Heptane	AI	AI	Separate Framework		Separate Analysis	
WCA	Model Checking			FIFO		
SWEET	AI					
METAMOC	Model Checking	Model Checking		Round Robin		
McAiT	AI		Separate Framework			Model Checking
Florida State, NC State, Furman University	SCS	SCS, CME	Separate Framework		Separate Analysis	
Chalmers University	Symbolic Execution	Symbolic Execution				

be further reduced (Unlike the general-purpose computing domain [136], cache management in real-time systems [82] optimizes the worst-case rather than the average-case performance.), and the schedulability of the overall system is improved.

6 Static Analysis Tools

In the past decades, a number of WCET analysis tools have been developed in both industry and academia. Table 5 lists the tools that support static cache analysis.

aiT [57] is the only static WCET analysis tool in routine use in industry. It has been *qualified*, i.e., admitted to certification of time-critical avionics subsystems of several Airbus planes by the European Aviation Safety Agency (EASA) and has been used in their certification. It is also used in other air and space companies in Europe, the United States, and China and in German automotive OEMs and their suppliers. It uses the AI-based analyses [31, 38] for both instruction and data caches. Besides LRU, the aiT tool can analyze three non-LRU replacement policies: Pseudo-Round-Robin [58] as well as, PLRU and FIFO based on the analyses described in [103] and [44].

The OTAWA tool [11] developed by the University of Toulouse, France, is an open framework for WCET analysis. OTAWA provides instruction cache analysis based on abstract interpretation [38] with the improvement of multi-level Persistence analysis [10].

Chronos [70] is a static WCET analysis tool from the National University of Singapore. It was originally designed with a highlight on pipeline analysis using the SimpleScalar simulator. The

latest version, Chronos 4.2, now supports the recent contributions of the group: scope-aware data cache analysis [62] and unified cache analysis [24].

Heptane [60] is a static WCET analysis tool developed by IRISA, France. The highlight of the tool is the separate analysis of multi-level caches [55]. It also support shared cache analysis by the technique extended from AI-based analysis [54].

WCA [109] from Vienna University of Technology and DTU is a WCET analysis tool for a Java processor, JOP [108], which uses *method cache* to store the instructions of a whole Java method. A method is fully loaded into the cache upon invocation and enjoys cache hits during its execution. On exit, the content of the caller function is reloaded into the method cache. The method cache is organized like a fully-associative FIFO cache with N blocks. The tool uses model checking to analyze the method cache, and it also provides a simple persistence analysis given that a code region can fit into the cache.

SWEET [117] is a WCET analysis tool currently maintained by Mälardalen University of Sweden. Although mainly focusing on flow analysis, it supports AI-based analysis for instruction caches [38].

The METAMOC tool [33] from Aalborg University of Denmark employs model checking for both instruction and data caches. It can analyze the round-robin replacement policy used by the ARM920T processor.

McAiT [80] is a WCET analysis tool jointly developed by Uppsala University of Sweden and Northeastern University of China. The tool supports L1 instruction cache analysis by the AI-based approaches [31, 38], and shared L2 cache analysis by model checking.

Other research prototypes include a tool from Florida State, North Carolina State, and Furman Universities, which adopts Static Cache Simulation [88] for both instruction and data cache analysis, and also supports data cache analysis using cache miss equations [97]. Another prototype from Chalmers University of Technology uses symbolic execution [77] for cache analysis.

The data provided in Table 5 might be imprecise, because the information can only be inferred from the publications instead of the tools in some cases. More comprehensive knowledge on existing WCET analysis tools can be found in [130] and the reports for the WCET Tool Challenge in 2011 [53], 2008 [61] and 2006 [48].

7 Future Research Directions

WCET estimation is a key task in timing analysis of real-time systems. Since caches may significantly affect execution time, the quality of cache analysis determines the precision of the estimated WCET. This article surveys the main challenges and analysis techniques for vast cache architectures. For decades, the LRU replacement policy has been well studied. The most valuable asset is that a comprehensive understanding of cache behavior and cache analysis were established by the ingenious researchers in related communities. Several future directions can be explored to bridge the gap.

Evaluation and Comparison of Different Approaches. The reader of this survey may be disappointed not to find evaluations and comparisons of the different methods for cache analysis. These are indeed hard to find in literature and are therefore subject of future research.

One of the obstacles to fair evaluation is the difficulty to obtain industrial software for experimentation; most industrial embedded software is not openly available. Another reason for the non-existence of good experimental evaluations is the dominance of the *Mälardalen Benchmark Suite*. The programs in this benchmark suite have a very special characteristic: They are small and contain tiny loops. They start with long straight-line code sequences for the initialization of the program variables. This alone makes them already problematic; the execution is independent

of the values of the input variables. Measuring this one execution would suffice if execution times were independent of the initial architectural state—which they often are not [104]. Analyzing the cache performance using the programs from this benchmark should stress the cache; access sequences without evictions do not provide any insights.

Analytical comparisons of the different approaches should, in principle, be possible. However, they have not been performed. For instance, the abstract cache state in Mueller’s static cache simulation [88, 90] is much like the abstract May cache in the abstract-interpretation-based approach [37, 38]. Intuition says that the precision of May information as obtained by abstract interpretation should therefore be the same as that obtained by static cache simulation. We would, however, assume that Must information as obtained by abstract interpretation is more precise than that obtained by static cache simulation since the computation of Must information from an abstract May cache employs rather strong conditions to eliminate contents from the May cache that cannot be guaranteed to be in all concrete caches.

Non-LRU Cache Analysis. Although LRU is highly predictable, it is practically more important to analyze non-LRU replacement policies since they are actually adopted in real-life processors. Must, May and Persistence analyses need to be established to fully characterize the cache behavior. Currently, the missing pieces are Persistence analysis for FIFO, Must and May analyses for MRU, Persistence and May analyses for PLRU. Furthermore, there are no techniques to analyze non-LRU data caches and multi-level caches, which are actually required to cover the whole cache hierarchy. For policies other than the above-mentioned ones, similar analysis targets should be fulfilled. However, we still lack a systematic way to construct abstract analyses for new replacement policies.

Application of Cache Analysis in Other Domains. So far the use of cache analysis has mostly been confined to WCET analysis. However, there is at least one more domain in which cache analysis can deliver valuable insights, namely security. *Side-channel attacks* recover secret inputs to programs from physical characteristics of the computation. Typical goals of such attacks are the recovery of cryptographic keys and private information about users. Characteristics that have been exploited for that purpose include execution time, cache behavior, memory and power consumption, and electromagnetic radiation. Doychev et al. have demonstrated that static cache analyses based on abstract interpretation can be used to derive guarantees on the amount of information leaked to an attacker [35].

Design and Analysis of Timing-Predictable Embedded Systems. Preemption delay analysis and multi-core shared cache analysis have to consider the interactions among tasks running in parallel. It is commonly acknowledged that inter-core interference not only harms cache analysis, but also degrades overall system performance. Academia has gradually come to a consensus [8, 29, 122]: the solution to this problem should be to regulate both the hardware [92, 131] and the software [36, 81, 94] so that the system behaves in a timely predictable manner. The grand challenge is to obtain predictability without sacrificing the performance provided by future powerful processors. Cache analysis will provide valuable insights to characterize tasks so that good design decisions can be made in resource allocation and arbitration, such as cache partitioning and cache-aware scheduling.

One important step in this direction would be to understand how *timing compositionality* [52] can be achieved. Due to complex interactions between caches and other microarchitectural components, such as branch predictors or out-of-order pipelines, provably sound WCET analyses can currently only be achieved by analyzing all of these components together in an *integrated* fashion. However, such an integrated approach is very unlikely to scale to multi-tasking systems

or even to the parallel execution of multiple tasks on a multi-core processor. In these scenarios, to limit analysis complexity, interference costs are better analyzed separately and then taken into account during schedulability analysis. Timing compositionality has, however, not been formally proven for models of *any* modern microarchitecture, leaving much of the recent work unapplicable to real systems.

Acknowledgement. This article was partially supported by National Natural Science Foundation of China (61370076 and 61532007), Collaborative Innovation Center of Major Machine Manufacturing in Liaoning, and State Key Laboratory of Synthetical Automation for Process Industries (PAL-N201503).

References

- 1 Martin Alt, Christian Ferdinand, Florian Martin, and Reinhard Wilhelm. Cache behavior prediction by abstract interpretation. In Radhia Cousot and David A. Schmidt, editors, *Static Analysis, Third International Symposium, SAS'96, Aachen, Germany, September 24-26, 1996, Proceedings*, volume 1145 of *Lecture Notes in Computer Science*, pages 52–66. Springer, 1996. doi:10.1007/3-540-61739-6_33.
- 2 Sebastian Altmeyer and Claire Burguière. A new notion of useful cache block to improve the bounds of cache-related preemption delay. In *21st Euromicro Conference on Real-Time Systems, ECRTS 2009, Dublin, Ireland, July 1-3, 2009*, pages 109–118. IEEE Computer Society, 2009. doi:10.1109/ECRTS.2009.21.
- 3 Sebastian Altmeyer, Claire Burguière, and Reinhard Wilhelm. Computing the maximum blocking time for scheduling with deferred preemption. In *Future Dependable Distributed Systems, 2009 Software Technologies for*, pages 200–204, March 2009. doi:10.1109/STFSSD.2009.12.
- 4 Sebastian Altmeyer, Robert I. Davis, and Claire Maiza. Improved cache related pre-emption delay aware response time analysis for fixed priority preemptive systems. *Real-Time Systems*, 48(5):499–526, 2012. doi:10.1007/s11241-012-9152-2.
- 5 Sebastian Altmeyer, Claire Maiza, and Jan Reineke. Resilience analysis: tightening the CRPD bound for set-associative caches. In Jaejin Lee and Bruce R. Childers, editors, *Proceedings of the ACM SIGPLAN/SIGBED 2010 conference on Languages, compilers, and tools for embedded systems, LCTES 2010, Stockholm, Sweden, April 13-15, 2010*, pages 153–162. ACM, 2010. doi:10.1145/1755888.1755911.
- 6 Rajeev Alur and David L. Dill. A theory of timed automata. *Theor. Comput. Sci.*, 126(2):183–235, 1994. doi:10.1016/0304-3975(94)90010-8.
- 7 Neil C. Audsley, Alan Burns, Mike M. Richardson, Ken Tindell, and Andy J. Wellings. Applying new scheduling theory to static priority pre-emptive scheduling. *Software Engineering Journal*, 8(5):284–292, 1993. URL: <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=238595>.
- 8 Philip Axer, Rolf Ernst, Heiko Falk, Alain Girault, Daniel Grund, Nan Guan, Bengt Jonsson, Peter Marwedel, Jan Reineke, Christine Rochange, Maurice Sebastian, Reinhard von Hanxleden, Reinhard Wilhelm, and Wang Yi. Building timing predictable embedded systems. *ACM Trans. Embedded Comput. Syst.*, 13(4):82:1–82:37, 2014. doi:10.1145/2560033.
- 9 Gogul Balakrishnan and Thomas W. Reps. Analyzing memory accesses in x86 executables. In Evelyn Duesterwald, editor, *Compiler Construction, 13th International Conference, CC 2004, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2004, Barcelona, Spain, March 29 – April 2, 2004, Proceedings*, volume 2985 of *Lecture Notes in Computer Science*, pages 5–23. Springer, 2004. doi:10.1007/978-3-540-24723-4_2.
- 10 Clément Ballabriga and Hugues Cassé. Improving the first-miss computation in set-associative instruction caches. In *20th Euromicro Conference on Real-Time Systems, ECRTS 2008, 2-4 July 2008, Prague, Czech Republic, Proceedings*, pages 341–350. IEEE Computer Society, 2008. doi:10.1109/ECRTS.2008.34.
- 11 Clément Ballabriga, Hugues Cassé, Christine Rochange, and Pascal Sainrat. OTAWA: an open toolbox for adaptive WCET analysis. In Sang Lyul Min, Robert G. Pettit IV, Peter P.uschner, and Theo Ungerer, editors, *Software Technologies for Embedded and Ubiquitous Systems – 8th IFIP WG 10.2 International Workshop, SEUS 2010, Waidhofen/Ybbs, Austria, October 13-15, 2010. Proceedings*, volume 6399 of *Lecture Notes in Computer Science*, pages 35–46. Springer, 2010. doi:10.1007/978-3-642-16256-5_6.
- 12 Abhijeet Banerjee, Sudipta Chattopadhyay, and Abhik Roychoudhury. Precise micro-architectural modeling for WCET analysis via AI+SAT. In *19th IEEE Real-Time and Embedded Technology and Applications Symposium, RTAS 2013, Philadelphia, PA, USA, April 9-11, 2013*, pages 87–96. IEEE Computer Society, 2013. doi:10.1109/RTAS.2013.6531082.
- 13 Sanjoy K. Baruah. The limited-preemption uniprocessor scheduling of sporadic task systems. In *17th Euromicro Conference on Real-Time Systems (ECRTS 2005), 6-8 July 2005, Palma de Mallorca, Spain, Proceedings*, pages 137–144.

- IEEE Computer Society, 2005. doi:10.1109/ECRTS.2005.32.
- 14 Christoph Berg. PLRU cache domino effects. In Frank Mueller and Frank Mueller, editors, *6th Intl. Workshop on Worst-Case Execution Time (WCET) Analysis, July 4, 2006, Dresden, Germany*, volume 4 of *OpenAccess Series in Informatics (OASICs)*, Dagstuhl, Germany, 2006. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/OASICs.WCET.2006.672.
 - 15 Marko Bertogna, Orgees Khani, Mauro Marinoni, Francesco Esposito, and Giorgio C. Buttazzo. Optimal selection of preemption points to minimize preemption overhead. In Karl-Erik Årzén, editor, *23rd Euromicro Conference on Real-Time Systems, ECRTS 2011, Porto, Portugal, 5-8 July, 2011*, pages 217–227. IEEE Computer Society, 2011. doi:10.1109/ECRTS.2011.28.
 - 16 Reinder J. Bril, Sebastian Altmeyer, Martijn M. H. P. van den Heuvel, Robert I. Davis, and Moris Behnam. Integrating cache-related pre-emption delays into analysis of fixed priority scheduling with pre-emption thresholds. In *Proceedings of the IEEE 35th IEEE Real-Time Systems Symposium, RTSS 2014, Rome, Italy, December 2-5, 2014*, pages 161–172. IEEE Computer Society, 2014. doi:10.1109/RTSS.2014.25.
 - 17 Claire Burguière, Jan Reineke, and Sebastian Altmeyer. Cache-related preemption delay computation for set-associative caches – pitfalls and solutions. In Niklas Holsti, editor, *9th Intl. Workshop on Worst-Case Execution Time Analysis, WCET 2009, Dublin, Ireland, July 1-3, 2009*, volume 10 of *OpenAccess Series in Informatics (OASICs)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany, 2009. doi:10.4230/OASICs.WCET.2009.2285.
 - 18 Claire Burguière and Christine Rochange. A contribution to branch prediction modeling in WCET analysis. In *2005 Design, Automation and Test in Europe Conference and Exposition (DATE 2005), 7-11 March 2005, Munich, Germany*, pages 612–617. IEEE Computer Society, 2005. doi:10.1109/DATE.2005.7.
 - 19 Alan Burns. Preemptive priority-based scheduling: An appropriate engineering approach. In Sang H. Son, editor, *Advances in Real-time Systems*, pages 225–248. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1995. URL: <http://dl.acm.org/citation.cfm?id=207721.207731>.
 - 20 José V. Busquets-Mataix, Juan José Serrano, Rafael Ors, Pedro J. Gil, and Andy J. Wellings. Adding instruction cache effect to schedulability analysis of preemptive real-time systems. In *2nd IEEE Real-Time Technology and Applications Symposium, RTAS'96, Boston, MA, USA, June 10-12, 1996*, pages 204–212. IEEE Computer Society, 1996. doi:10.1109/RTAS.1996.509537.
 - 21 Giorgio C. Buttazzo. *Hard Real-time Computing Systems: Predictable Scheduling Algorithms And Applications*. Springer-Verlag TELOS, Santa Clara, CA, USA, 2004.
 - 22 Franck Cassez, René Rydhof Hansen, and Mads Chr. Olesen. What is a timing anomaly? In Tullio Vardanega, editor, *12th International Workshop on Worst-Case Execution Time Analysis, WCET 2012, July 10, 2012, Pisa, Italy*, volume 23 of *OpenAccess Series in Informatics (OASICs)*, pages 1–12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2012. doi:10.4230/OASICs.WCET.2012.1.
 - 23 Siddhartha Chatterjee, Erin Parker, Philip J. Hanlon, and Alvin R. Lebeck. Exact analysis of the cache behavior of nested loops. In Michael Burke and Mary Lou Soffa, editors, *Proceedings of the 2001 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Snowbird, Utah, USA, June 20-22, 2001*, pages 286–297. ACM, 2001. doi:10.1145/378795.378859.
 - 24 Sudipta Chattopadhyay and Abhik Roychoudhury. Unified cache modeling for WCET analysis and layout optimizations. In Theodore P. Baker, editor, *Proceedings of the 30th IEEE Real-Time Systems Symposium, RTSS 2009, Washington, DC, USA, 1-4 December 2009*, pages 47–56. IEEE Computer Society, 2009. doi:10.1109/RTSS.2009.20.
 - 25 Sudipta Chattopadhyay and Abhik Roychoudhury. Scalable and precise refinement of cache timing analysis via model checking. In *Proceedings of the 32nd IEEE Real-Time Systems Symposium, RTSS 2011, Vienna, Austria, November 29 - December 2, 2011*, pages 193–203. IEEE Computer Society, 2011. doi:10.1109/RTSS.2011.15.
 - 26 Edmund M. Clarke, Orna Grumberg, and Doron A. Peled. *Model checking*. MIT Press, 2001.
 - 27 Philippe Clauss. Counting solutions to linear and nonlinear constraints through ehrhart polynomials: Applications to analyze and transform scientific programs. In Pen-Chung Yew, editor, *Proceedings of the 10th international conference on Supercomputing, ICS 1996, Philadelphia, PA, USA, May 25-28, 1996*, pages 278–285. ACM, 1996. doi:10.1145/237578.237617.
 - 28 Antoine Colin and Isabelle Puaut. Worst case execution time analysis for a processor with branch prediction. *Real-Time Systems*, 18(2/3):249–274, 2000. doi:10.1023/A:1008149332687.
 - 29 PREDATOR Consortium. The predator project page, 2011. URL: <http://www.predator-project.eu/consortium.htm>.
 - 30 Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In Robert M. Graham, Michael A. Harrison, and Ravi Sethi, editors, *Conference Record of the Fourth ACM Symposium on Principles of Programming Languages, Los Angeles, California, USA, January 1977*, pages 238–252. ACM, 1977. doi:10.1145/512950.512973.
 - 31 Christoph Cullmann. Cache persistence analysis: Theory and practice. *ACM Trans. Embedded Comput. Syst.*, 12(1s):40, 2013. doi:10.1145/2435227.2435236.
 - 32 Andreas Engelbrecht Dalsgaard, Mads Chr. Olesen, Martin Toft, René Rydhof Hansen, and Kim Guldstrand Larsen. METAMOC: modular execution time analysis using model checking. In Björn Lisper, editor, *10th International Workshop on Worst-Case Execution Time Analysis,*

- WCET 2010, July 6, 2010, Brussels, Belgium, volume 15 of *OpenAccess Series in Informatics (OASICs)*, pages 113–123. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany, 2010. doi:10.4230/OASICs.WCET.2010.113.
- 33 Andreas Engelbrecht Dalsgaard, Mads Chr. Olesen, Martin Toft, René Rydhof Hansen, and Kim Guldstrand Larsen. METAMOC: modular execution time analysis using model checking. In Björn Lisper, editor, *10th International Workshop on Worst-Case Execution Time Analysis, WCET 2010, July 6, 2010, Brussels, Belgium*, volume 15 of *OpenAccess Series in Informatics (OASICs)*, pages 113–123. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany, 2010. doi:10.4230/OASICs.WCET.2010.113.
 - 34 Robert I. Davis and Alan Burns. A survey of hard real-time scheduling for multiprocessor systems. *ACM Comput. Surv.*, 43(4):35, 2011. doi:10.1145/1978802.1978814.
 - 35 Goran Doychev, Dominik Feld, Boris Köpf, Laurent Mauborgne, and Jan Reineke. Cacheaudit: A tool for the static analysis of cache side channels. In Samuel T. King, editor, *Proceedings of the 22th USENIX Security Symposium, Washington, DC, USA, August 14-16, 2013*, pages 431–446. USENIX Association, 2013. URL: <https://www.usenix.org/conference/usenixsecurity13/technical-sessions/paper/doychev>.
 - 36 Heiko Falk and Helena Kotthaus. Wcet-driven cache-aware code positioning. In Rajesh K. Gupta and Vincent John Mooney, editors, *Proceedings of the 14th International Conference on Compilers, Architecture, and Synthesis for Embedded Systems, CASES 2011, part of the Seventh Embedded Systems Week, ESWeek 2011, Taipei, Taiwan, October 9-14, 2011*, pages 145–154. ACM, 2011. doi:10.1145/2038698.2038722.
 - 37 Christian Ferdinand. *Cache Behavior Prediction for Real-Time Systems*. PhD thesis, Saarland University, Saarbruecken, Germany, 1997. ISBN: 3-9307140-31-0.
 - 38 Christian Ferdinand and Reinhard Wilhelm. Efficient and precise cache behavior prediction for real-time systems. *Real-Time Systems*, 17(2-3):131–181, 1999. doi:10.1023/A:1008186323068.
 - 39 Somnath Ghosh, Margaret Martonosi, and Sharad Malik. Cache miss equations: a compiler framework for analyzing and tuning memory behavior. *ACM Trans. Program. Lang. Syst.*, 21(4):703–746, 1999. doi:10.1145/325478.325479.
 - 40 Daniel Grund. *Static Cache Analysis for Real-Time Systems – LRU, FIFO, PLRU*. PhD thesis, Saarland University, Saarbruecken, Germany, 2011. ISBN: 978-3-8442-1699-8.
 - 41 Daniel Grund and Jan Reineke. Abstract interpretation of FIFO replacement. In Jens Palsberg and Zhendong Su, editors, *Static Analysis, 16th International Symposium, SAS 2009, Los Angeles, CA, USA, August 9-11, 2009. Proceedings*, volume 5673 of *Lecture Notes in Computer Science*, pages 120–136. Springer, 2009. doi:10.1007/978-3-642-03237-0_10.
 - 42 Daniel Grund and Jan Reineke. 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*, pages 155–164. IEEE Computer Society, 2010. doi:10.1109/ECRTS.2010.8.
 - 43 Daniel Grund and Jan Reineke. Toward precise PLRU cache analysis. In Björn Lisper, editor, *10th International Workshop on Worst-Case Execution Time Analysis, WCET 2010, July 6, 2010, Brussels, Belgium*, volume 15 of *OpenAccess Series in Informatics (OASICs)*, pages 23–35. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany, 2010. doi:10.4230/OASICs.WCET.2010.23.
 - 44 Daniel Grund, Jan Reineke, and Gernot Gebhard. Branch target buffers: WCET analysis framework and timing predictability. *Journal of Systems Architecture – Embedded Systems Design*, 57(6):625–637, 2011. doi:10.1016/j.sysarc.2010.05.013.
 - 45 Nan Guan, Mingsong Lv, Wang Yi, and Ge Yu. WCET analysis with MRU caches: Challenging LRU for predictability. In Marco Di Natale, editor, *2012 IEEE 18th Real Time and Embedded Technology and Applications Symposium, Beijing, China, April 16-19, 2012*, pages 55–64. IEEE Computer Society, 2012. doi:10.1109/RTAS.2012.31.
 - 46 Nan Guan, Martin Stigge, Wang Yi, and Ge Yu. Cache-aware scheduling and analysis for multi-cores. In Samarjit Chakraborty and Nicolas Halbwachs, editors, *Proceedings of the 9th ACM & IEEE International conference on Embedded software, EMSOFT 2009, Grenoble, France, October 12-16, 2009*, pages 245–254. ACM, 2009. doi:10.1145/1629335.1629369.
 - 47 Nan Guan, Xinpeng Yang, Mingsong Lv, and Wang Yi. FIFO cache analysis for WCET estimation: a quantitative approach. In Enrico Macii, editor, *Design, Automation and Test in Europe, DATE 13, Grenoble, France, March 18-22, 2013*, pages 296–301. EDA Consortium San Jose, CA, USA / ACM DL, 2013. doi:10.7873/DATE.2013.073.
 - 48 Jan Gustafsson. WCET challenge 2006 – technical report. Technical Report ISSN 1404-3041 ISRN MDH-MRTC-206/2007-1-SE, Mälardalen University Sweden, January 2007. URL: <http://www.es.mdh.se/publications/1020->.
 - 49 Jan Gustafsson, Andreas Ermedahl, Christer Sandberg, and Björn Lisper. Automatic derivation of loop bounds and infeasible paths for WCET analysis using abstract execution. In *Proceedings of the 27th IEEE Real-Time Systems Symposium (RTSS 2006), 5-8 December 2006, Rio de Janeiro, Brazil*, pages 57–66. IEEE Computer Society, 2006. doi:10.1109/RTSS.2006.12.
 - 50 Andreas Gustavsson, Andreas Ermedahl, Björn Lisper, and Paul Pettersson. Towards WCET analysis of multicore architectures using UPPAAL. In Björn Lisper, editor, *10th International Workshop on Worst-Case Execution Time Analysis, WCET 2010, July 6, 2010, Brussels, Belgium*, volume 15 of *OpenAccess Series in Informatics (OASICs)*, pages 101–112. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany, 2010. doi:10.4230/OASICs.WCET.2010.101.

- 51 Sebastian Hahn and Daniel Grund. Relational cache analysis for static timing analysis. In Robert Davis, editor, *24th Euromicro Conference on Real-Time Systems, ECRTS 2012, Pisa, Italy, July 11-13, 2012*, pages 102–111. IEEE Computer Society, 2012. doi:10.1109/ECRTS.2012.14.
- 52 Sebastian Hahn, Jan Reineke, and Reinhard Wilhelm. Towards compositionality in execution time analysis: definition and challenges. *SIGBED Review*, 12(1):28–36, 2015. doi:10.1145/2752801.2752805.
- 53 Reinhard von Hanxleden, Niklas Holsti, Björn Lisper, Erhard Ploedereder, Armelle Bonenfant, Hugues Cassé, Sven Bünthe, Wolfgang Fellger, Sebastian Gepperth, Jan Gustafsson, Benedikt Huber, Nazrul Mohammad Islam, Daniel Kästner, Raimund Kirner, Laura Kovacs, Felix Krause, Marianne de Michiel, Mads Christian Olesen, Adrian Prantl, Wolfgang Puffitsch, Christine Rochange, Martin Schoeberl, Simon Wegener, Michael Zolda, and Jakob Zwirchmayr. WCET tool challenge 2011: Report. In *the 11th International Workshop on Worst-Case Execution Time (WCET) Analysis (WCET 2011)*, July 2011.
- 54 Damien Hardy, Thomas Piquet, and Isabelle Puaut. Using bypass to tighten WCET estimates for multi-core processors with shared instruction caches. In Theodore P. Baker, editor, *Proceedings of the 30th IEEE Real-Time Systems Symposium, RTSS 2009, Washington, DC, USA, 1-4 December 2009*, pages 68–77. IEEE Computer Society, 2009. doi:10.1109/RTSS.2009.34.
- 55 Damien Hardy and Isabelle Puaut. WCET analysis of multi-level non-inclusive set-associative instruction caches. In *Proceedings of the 29th IEEE Real-Time Systems Symposium, RTSS 2008, Barcelona, Spain, 30 November – 3 December 2008*, pages 456–466. IEEE Computer Society, 2008. doi:10.1109/RTSS.2008.10.
- 56 Damien Hardy and Isabelle Puaut. WCET analysis of instruction cache hierarchies. *Journal of Systems Architecture – Embedded Systems Design*, 57(7):677–694, 2011. doi:10.1016/j.sysarc.2010.08.007.
- 57 Reinhold Heckmann and Christian Ferdinand. Worst-case execution time prediction by static program analysis. *The whitepaper of aiT*, 2014. URL: https://www.absint.com/aiT_WCET.pdf.
- 58 Reinhold Heckmann, Marc Langenbach, Stephan Thesing, and Reinhard Wilhelm. The influence of processor architecture on the design and the results of WCET tools. *Proceedings of the IEEE*, 91(7):1038–1054, 2003. doi:10.1109/JPROC.2003.814618.
- 59 John L. Hennessy and David A. Patterson. *Computer Architecture: A Quantitative Approach*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 5th edition, 2011.
- 60 Heptane. The Heptane tool page, 2013. URL: http://www.irisa.fr/alf/index.php?option=com_content&view=article&id=29&Itemid=0&lang=en.
- 61 Niklas Holsti, Jan Gustafsson, Guillem Bernat, Clément Ballabriga, Armelle Bonenfant, Roman Bourgade, Hugues Cassé, Daniel Cordes, Albrecht Kadlec, Raimund Kirner, Jens Knoop, Paul Lokuciejewski, Nicholas Merriam, Marianne De Michiel, Adrian Prantl, Bernhard Rieder, Christine Rochange, Pascal Sainrat, and Markus Schordan. WCET 2008 – report from the tool challenge 2008 – 8th intl. workshop on worst-case execution time (WCET) analysis. In Raimund Kirner, editor, *8th Intl. Workshop on Worst-Case Execution Time (WCET) Analysis, Prague, Czech Republic, July 1, 2008*, volume 8 of *OpenAccess Series in Informatics (OASICs)*. Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), Schloss Dagstuhl, Germany, 2008. URL: <http://drops.dagstuhl.de/opus/volltexte/2008/1663>, doi:10.4230/OASICs.WCET.2008.1663.
- 62 Bach Khoa Huynh, Lei Ju, and Abhik 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*, pages 203–212. IEEE Computer Society, 2011. doi:10.1109/RTAS.2011.27.
- 63 Lei Ju, Samarjit Chakraborty, and Abhik Roychoudhury. Accounting for cache-related preemption delay in dynamic priority schedulability analysis. In *Design, Automation Test in Europe Conference Exhibition, 2007*, pages 1–6, April 2007. doi:10.1109/DATE.2007.364534.
- 64 Gary A. Kildall. A unified approach to global program optimization. In Patrick C. Fischer and Jeffrey D. Ullman, editors, *Conference Record of the ACM Symposium on Principles of Programming Languages, Boston, Massachusetts, USA, October 1973*, pages 194–206. ACM Press, 1973. doi:10.1145/512927.512945.
- 65 Sung-Kwan Kim, Sang Lyul Min, and Rhan 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*, pages 230–240. IEEE Computer Society, 1996. doi:10.1109/RTAS.1996.509540.
- 66 Chang-Gun Lee, Joosun Hahn, Yang-Min Seo, Sang Lyul Min, Rhan Ha, Seongsoo Hong, Chang Yun Park, Minsuk Lee, and Chong-Sang Kim. Analysis of cache-related preemption delay in fixed-priority preemptive scheduling. *IEEE Trans. Computers*, 47(6):700–713, 1998. doi:10.1109/12.689649.
- 67 Benjamin Lesage, Damien Hardy, and Isabelle Puaut. WCET analysis of multi-level set-associative data caches. In Niklas Holsti, editor, *9th Intl. Workshop on Worst-Case Execution Time Analysis, WCET 2009, Dublin, Ireland, July 1-3, 2009*, volume 10 of *OpenAccess Series in Informatics (OASICs)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany, 2009. doi:10.4230/OASICs.WCET.2009.2283.
- 68 Benjamin Lesage, Damien Hardy, and Isabelle Puaut. Shared data cache conflicts reduction for wcet computation in multi-core architectures. In *Proc. of the 18th Real-Time and Network Systems, Toulouse, France, 2010*.
- 69 Xianfeng Li, Abhik Roychoudhury, and Tulika Mitra. Modeling out-of-order processors for WCET

- analysis. *Real-Time Systems*, 34(3):195–227, 2006. doi:10.1007/s11241-006-9205-5.
- 70 Xianfeng Li, Liang Yun, Tulika Mitra, and Abhik Roychoudhury. Chronos: A timing analyzer for embedded software. *Sci. Comput. Program.*, 69(1-3):56–67, 2007. doi:10.1016/j.scico.2007.01.014.
- 71 Yau-Tsun Steven Li and Sharad Malik. Performance analysis of embedded software using implicit path enumeration. In *DAC*, pages 456–461, 1995. doi:10.1145/217474.217570.
- 72 Yau-Tsun Steven Li, Sharad Malik, and Andrew Wolfe. Cache modeling for real-time software: beyond direct mapped instruction caches. In *Proceedings of the 17th IEEE Real-Time Systems Symposium (RTSS'96)*, December 4-6, 1996, Washington, DC, USA, pages 254–263. IEEE Computer Society, 1996. doi:10.1109/REAL.1996.563722.
- 73 Yun Liang, Huping Ding, Tulika Mitra, Abhik Roychoudhury, Yan Li, and Vivvy Suhendra. Timing analysis of concurrent programs running on shared cache multi-cores. *Real-Time Systems*, 48(6):638–680, 2012. doi:10.1007/s11241-012-9160-2.
- 74 Fang Liu and Yan Solihin. Understanding the behavior and implications of context switch misses. *TACO*, 7(4):21, 2010. doi:10.1145/1880043.1880048.
- 75 Tiantian Liu, Yingchao Zhao, Minming Li, and Chun Jason Xue. Task assignment with cache partitioning and locking for WCET minimization on mp soc. In *39th International Conference on Parallel Processing, ICPP 2010, San Diego, California, USA, 13-16 September 2010*, pages 573–582. IEEE Computer Society, 2010. doi:10.1109/ICPP.2010.65.
- 76 Thomas Lundqvist and Per Stenström. A method to improve the estimated worst-case performance of data caching. In *6th International Workshop on Real-Time Computing and Applications Symposium (RTCSA'99)*, 13-16 December 1999, Hong Kong, China, pages 255–262. IEEE Computer Society, 1999. doi:10.1109/RTCSA.1999.811244.
- 77 Thomas Lundqvist and Per Stenström. A method to improve the estimated worst-case performance of data caching. In *6th International Workshop on Real-Time Computing and Applications Symposium (RTCSA'99)*, 13-16 December 1999, Hong Kong, China, pages 255–262. IEEE Computer Society, 1999. doi:10.1109/RTCSA.1999.811244.
- 78 Will Lunniss, Sebastian Altmeyer, and Robert I. Davis. A comparison between fixed priority and EDF scheduling accounting for cache related preemption delays. *LITES*, 1(1):01:1–01:24, 2014. doi:10.4230/LITES-v001-i001-a001.
- 79 Will Lunniss, Sebastian Altmeyer, Claire Maiza, and Robert I. Davis. Integrating cache related pre-emption delay analysis into EDF scheduling. In *19th IEEE Real-Time and Embedded Technology and Applications Symposium, RTAS 2013, Philadelphia, PA, USA, April 9-11, 2013*, pages 75–84. IEEE Computer Society, 2013. doi:10.1109/RTAS.2013.6531081.
- 80 Mingsong Lv, Nan Guan, Qingxu Deng, Ge Yu, and Wang Yi. Mcait – A timing analyzer for multicore real-time software. In Tevfik Bultan and Pao-Ann Hsiung, editors, *Automated Technology for Verification and Analysis, 9th International Symposium, ATVA 2011, Taipei, Taiwan, October 11-14, 2011. Proceedings*, volume 6996 of *Lecture Notes in Computer Science*, pages 414–417. Springer, 2011. doi:10.1007/978-3-642-24372-1_29.
- 81 Mohamed Abdel Maksoud and Jan Reineke. A compiler optimization to increase the efficiency of WCET analysis. In Mathieu Jan, Belgacem Ben Hedia, Joël Goossens, and Claire Maiza, editors, *22nd International Conference on Real-Time Networks and Systems, RTNS'14, Versailles, France, October 8-10, 2014*, page 87. ACM, 2014. doi:10.1145/2659787.2659825.
- 82 Renato Mancuso, Roman Dudko, Emiliano Betti, Marco Cesati, Marco Caccamo, and Rodolfo Pellizzoni. Real-time cache management framework for multi-core architectures. In *19th IEEE Real-Time and Embedded Technology and Applications Symposium, RTAS 2013, Philadelphia, PA, USA, April 9-11, 2013*, pages 45–54. IEEE Computer Society, 2013. doi:10.1109/RTAS.2013.6531078.
- 83 José Marinho, Vincent Nélis, Stefan M. Petters, and Isabelle Puaut. An improved preemption delay upper bound for floating non-preemptive region. In *7th IEEE International Symposium on Industrial Embedded Systems, SIES 2012, Karlsruhe, Germany, June 20-22, 2012*, pages 57–66. IEEE, 2012. doi:10.1109/SIES.2012.6356570.
- 84 José Marinho, Vincent Nélis, Stefan M. Petters, and Isabelle Puaut. Preemption delay analysis for floating non-preemptive region scheduling. In Wolfgang Rosenstiel and Lothar Thiele, editors, *2012 Design, Automation & Test in Europe Conference & Exhibition, DATE 2012, Dresden, Germany, March 12-16, 2012*, pages 497–502. IEEE, 2012. doi:10.1109/DATE.2012.6176520.
- 85 Florian Martin, Martin Alt, Reinhard Wilhelm, and Christian Ferdinand. Analysis of loops. In Kai Koskimies, editor, *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*, volume 1383 of *Lecture Notes in Computer Science*, pages 80–94. Springer, 1998. doi:10.1007/BFb0026424.
- 86 Richard L. Mattson, Jan Gecsei, Donald R. Slutz, and Irving L. Traiger. Evaluation techniques for storage hierarchies. *IBM Systems Journal*, 9(2):78–117, 1970. doi:10.1147/sj.92.0078.
- 87 Frank Mueller. Timing predictions for multi-level caches. In *In ACM SIGPLAN Workshop on Language, Compiler, and Tool Support for Real-Time Systems*, pages 29–36, 1997.
- 88 Frank Mueller and David B. Whalley. Fast instruction cache analysis via static cache simulation. In *Proceedings 28st Annual Simulation Symposium (SS'95)*, April 25-28, 1995, Santa Barbara, California, USA, pages 105–114. IEEE Computer Society, 1995. doi:10.1109/SIMSYM.1995.393589.
- 89 Frank Müller. Generalizing timing predictions to set-associative caches. In *Proceedings of the*

- Ninth Euromicro Workshop on Real-Time Systems, RTS 1997, 11-13 June, 1997, Toledo, Spain*, pages 64–71. IEEE Computer Society, 1997. doi:10.1109/EMWRTS.1997.613765.
- 90 Frank Müller. Timing analysis for instruction caches. *Real-Time Systems*, 18(2/3):217–247, 2000. doi:10.1023/A:1008145215849.
- 91 Hemendra Singh Negi, Tulika Mitra, and Abhik Roychoudhury. Accurate estimation of cache-related preemption delay. In Rajesh Gupta, Yukihiro Nakamura, Alex Orailoglu, and Pai H. Chou, editors, *Proceedings of the 1st IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis, CODES+ISSS 2003, Newport Beach, CA, USA, October 1-3, 2003*, pages 201–206. ACM, 2003. doi:10.1145/944645.944698.
- 92 Marco Paolieri, Eduardo Quiñones, Francisco J. Cazorla, Guillem Bernat, and Mateo Valero. Hardware support for WCET analysis of hard real-time multicore systems. In Stephen W. Keckler and Luiz André Barroso, editors, *36th International Symposium on Computer Architecture (ISCA 2009), June 20-24, 2009, Austin, TX, USA*, pages 57–68. ACM, 2009. doi:10.1145/1555754.1555764.
- 93 Kaustubh Patil, Kiran Seth, and Frank Mueller. Compositional static instruction cache simulation. In David B. Whalley and Ron Cytron, editors, *Proceedings of the 2004 ACM SIGPLAN/SIGBED Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES'04), Washington, DC, USA, June 11-13, 2004*, pages 136–145. ACM, 2004. doi:10.1145/997163.997183.
- 94 Rodolfo Pellizzoni, Emiliano Betti, Stanley Bak, Gang Yao, John Criswell, Marco Caccamo, and Russell Kegley. A predictable execution model for cots-based embedded systems. In *17th IEEE Real-Time and Embedded Technology and Applications Symposium, RTAS 2011, Chicago, Illinois, USA, 11-14 April 2011*, pages 269–279. IEEE Computer Society, 2011. doi:10.1109/RTAS.2011.33.
- 95 Guillaume Phavorin and Pascal Richard. Cache-related preemption delays and real-time scheduling: A survey for uniprocessor systems. Technical report, Laboratoire d'Informatique et d'Automatique pour les Systèmes, 2015. URL: <http://www.lias-lab.fr/publications/19296/survey.pdf>.
- 96 Guillaume Phavorin, Pascal Richard, Joël Goossens, Thomas Chapeaux, and Claire Maiza. Scheduling with preemption delays: anomalies and issues. In Julien Forget, editor, *Proceedings of the 23rd International Conference on Real Time Networks and Systems, RTNS 2015, Lille, France, November 4-6, 2015*, pages 109–118. ACM, 2015. doi:10.1145/2834848.2834853.
- 97 Harini Ramaprasad and Frank Mueller. Bounding worst-case data cache behavior by analytically deriving cache reference patterns. In *11th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS 2005), 7-10 March 2005, San Francisco, CA, USA*, pages 148–157. IEEE Computer Society, 2005. doi:10.1109/RTAS.2005.12.
- 98 Harini Ramaprasad and Frank Mueller. Bounding preemption delay within data cache reference patterns for real-time tasks. In *12th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS 2006), 4-7 April 2006, San Jose, California, USA*, pages 71–80. IEEE Computer Society, 2006. doi:10.1109/RTAS.2006.14.
- 99 Harini Ramaprasad and Frank Mueller. Bounding worst-case response time for tasks with non-preemptive regions. In *Proceedings of the 14th IEEE Real-Time and Embedded Technology and Applications Symposium, RTAS 2008, April 22-24, 2008, St. Louis, Missouri, USA*, pages 58–67. IEEE Computer Society, 2008. doi:10.1109/RTAS.2008.18.
- 100 Harini Ramaprasad and Frank Mueller. Tightening the bounds on feasible preemptions. *ACM Trans. Embedded Comput. Syst.*, 10(2):27, 2010. doi:10.1145/1880050.1880063.
- 101 Jan Reineke. *Caches in WCET Analysis: Predictability – Competitiveness – Sensitivity*. PhD thesis, Saarland University, 2009. URL: <http://www.epubli.de/shop/buch/Caches-in-WCET-Analysis-Jan-Reineke-9783941071698/12835>.
- 102 Jan Reineke, Sebastian Altmeyer, Daniel Grund, Sebastian Hahn, and Claire Maiza. Selfish-lru: Preemption-aware caching for predictability and performance. In *20th IEEE Real-Time and Embedded Technology and Applications Symposium, RTAS 2014, Berlin, Germany, April 15-17, 2014*, pages 135–144. IEEE Computer Society, 2014. doi:10.1109/RTAS.2014.6925997.
- 103 Jan Reineke and Daniel Grund. Relative competitive analysis of cache replacement policies. In Krisztián Flautner and John Regehr, editors, *Proceedings of the 2008 ACM SIGPLAN/SIGBED Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES'08), Tucson, AZ, USA, June 12-13, 2008*, pages 51–60. ACM, 2008. doi:10.1145/1375657.1375665.
- 104 Jan Reineke and Daniel Grund. Sensitivity of cache replacement policies. *ACM Trans. Embedded Comput. Syst.*, 12(1s):42, 2013. doi:10.1145/2435227.2435238.
- 105 Jan Reineke, Daniel Grund, Christoph Berg, and Reinhard Wilhelm. Timing predictability of cache replacement policies. *Real-Time Systems*, 37(2):99–122, 2007. doi:10.1007/s11241-007-9032-3.
- 106 Jan Reineke, Björn Wachter, Stephan Thesing, Reinhard Wilhelm, Ilia Polian, Jochen Eisinger, and Bernd Becker. A definition and classification of timing anomalies. In Frank Mueller, editor, *6th Intl. Workshop on Worst-Case Execution Time (WCET) Analysis, July 4, 2006, Dresden, Germany*, volume 4 of *OpenAccess Series in Informatics (OASICS)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2006. doi:10.4230/OASICS.WCET.2006.671.
- 107 Jörn Schneider. Cache and pipeline sensitive fixed priority scheduling for preemptive real-time systems. In *Proceedings of the 21st IEEE Real-Time Systems Symposium (RTSS 2000), Orlando, Florida, USA, 27-30 November 2000*, pages 195–204. IEEE Computer Society, 2000. doi:10.1109/REAL.2000.896009.

- 108 Martin Schoeberl. A java processor architecture for embedded real-time systems. *Journal of Systems Architecture – Embedded Systems Design*, 54(1-2):265–286, 2008. doi:10.1016/j.sysarc.2007.06.001.
- 109 Martin Schoeberl, Wolfgang Puffitsch, Rasmus Ulslev Pedersen, and Benedikt Huber. Worst-case execution time analysis for a java processor. *Softw., Pract. Exper.*, 40(6):507–542, 2010. doi:10.1002/spe.968.
- 110 Rathijit Sen and Y. N. Srikant. WCET estimation for executables in the presence of data caches. In Christoph M. Kirsch and Reinhard Wilhelm, editors, *Proceedings of the 7th ACM & IEEE International conference on Embedded software, EMSOFT 2007, September 30 – October 3, 2007, Salzburg, Austria*, pages 203–212. ACM, 2007. doi:10.1145/1289927.1289960.
- 111 Tyler Sondag and Hridesh Rajan. A more precise abstract domain for multi-level caches for tighter WCET analysis. In *Proceedings of the 31st IEEE Real-Time Systems Symposium, RTSS 2010, San Diego, California, USA, November 30 – December 3, 2010*, pages 395–404. IEEE Computer Society, 2010. doi:10.1109/RTSS.2010.8.
- 112 Jan Staschulat and Rolf Ernst. Worst case timing analysis of input dependent data cache behavior. In *18th Euromicro Conference on Real-Time Systems, ECRTS’06, 5-7 July 2006, Dresden, Germany, Proceedings*, pages 227–236. IEEE Computer Society, 2006. doi:10.1109/ECRTS.2006.33.
- 113 Jan Staschulat and Rolf Ernst. Scalable precision cache analysis for real-time software. *ACM Trans. Embedded Comput. Syst.*, 6(4), 2007. doi:10.1145/1274858.1274863.
- 114 Jan Staschulat, Simon Schliecker, and Rolf Ernst. Scheduling analysis of real-time systems with precise modeling of cache related preemption delay. In *17th Euromicro Conference on Real-Time Systems (ECRTS 2005), 6-8 July 2005, Palma de Mallorca, Spain, Proceedings*, pages 41–48. IEEE Computer Society, 2005. doi:10.1109/ECRTS.2005.26.
- 115 Martin Stigge and Wang Yi. Graph-based models for real-time workload: a survey. *Real-Time Systems*, 51(5):602–636, 2015. doi:10.1007/s11241-015-9234-z.
- 116 Vivy Suhendra and Tulika Mitra. Exploring locking & partitioning for predictable shared caches on multi-cores. In Limor Fix, editor, *Proceedings of the 45th Design Automation Conference, DAC 2008, Anaheim, CA, USA, June 8-13, 2008*, pages 300–303. ACM, 2008. doi:10.1145/1391469.1391545.
- 117 SWEET. The SWEET tool page, 2012. URL: <http://www.mrtc.mdh.se/projects/wcet/sweet/online/content/index.php>.
- 118 Yudong Tan and Vincent John Mooney III. Timing analysis for preemptive multitasking real-time systems with caches. *ACM Trans. Embedded Comput. Syst.*, 6(1), 2007. doi:10.1145/1210268.1210275.
- 119 Yudong Tan and Vincent John Mooney. Integrated intra- and inter-task cache analysis for preemptive multi-tasking real-time systems. In Henk Schepers, editor, *Software and Compilers for Embedded Systems, 8th International Workshop, SCOPES 2004, Amsterdam, The Netherlands, September 2-3, 2004, Proceedings*, volume 3199 of *Lecture Notes in Computer Science*, pages 182–199. Springer, 2004. doi:10.1007/978-3-540-30113-4_14.
- 120 Henrik Theiling, Christian Ferdinand, and Reinhard Wilhelm. Fast and precise WCET prediction by separated cache and path analyses. *Real-Time Systems*, 18(2/3):157–179, 2000. doi:10.1023/A:1008141130870.
- 121 Hiroyuki Tomiyama and Nikil D. Dutt. Program path analysis to bound cache-related preemption delay in preemptive real-time systems. In Frank Vahid and Jan Madsen, editors, *Proceedings of the Eighth International Workshop on Hardware/Software Codesign, CODES 2000, San Diego, California, USA, 2000*, pages 67–71. ACM, 2000. doi:10.1145/334012.334025.
- 122 Theo Ungerer, Francisco J. Cazorla, Pascal Sainrat, Guillem Bernat, Zlatko Petrov, Christine Rochange, Eduardo Quiñones, Mike Gerdes, Marco Paolieri, Julian Wolf, Hugues Cassé, Sascha Uhrig, Irakli Gulashvili, Michael Houston, Florian Kluge, Stefan Metzloff, and Jörg Mische. Merasa: Multicore execution of hard real-time applications supporting analyzability. *IEEE Micro*, 30(5):66–75, 2010. doi:10.1109/MM.2010.78.
- 123 Xavier Vera, Björn Lisper, and Jingling Xue. Data caches in multitasking hard real-time systems. In *Proceedings of the 24th IEEE Real-Time Systems Symposium (RTSS 2003), 3-5 December 2003, Cancun, Mexico*, pages 154–165. IEEE Computer Society, 2003. doi:10.1109/REAL.2003.1253263.
- 124 Xavier Vera and Jingling Xue. Let’s study whole-program cache behaviour analytically. In *Proceedings of the Eighth International Symposium on High-Performance Computer Architecture (HPCA’02), Boston, Massachusetts, USA, February 2-6, 2002*, pages 175–186. IEEE Computer Society, 2002. doi:10.1109/HPCA.2002.995708.
- 125 Yun Wang and Manas Saksena. Scheduling fixed-priority tasks with preemption threshold. In *6th International Workshop on Real-Time Computing and Applications Symposium (RTCSA’99), 13-16 December 1999, Hong Kong, China*, page 328. IEEE Computer Society, 1999. doi:10.1109/RTCSA.1999.811269.
- 126 Bryan C. Ward, Jonathan L. Herman, Christopher J. Kenna, and James H. Anderson. Outstanding paper award: Making shared caches more predictable on multicore platforms. In *25th Euromicro Conference on Real-Time Systems, ECRTS 2013, Paris, France, July 9-12, 2013*, pages 157–167. IEEE Computer Society, 2013. doi:10.1109/ECRTS.2013.26.
- 127 Ingomar Wenzel, Raimund Kirner, Bernhard Rieder, and Peter P. Puschner. Measurement-based timing analysis. In Tiziana Margaria and Bernhard Steffen, editors, *Leveraging Applications of Formal Methods, Verification and Validation, Third International Symposium, ISoLA 2008, Porto Sani, Greece, October 13-15, 2008. Proceedings*, volume 17 of *Communications in*

- Computer and Information Science*, pages 430–444. Springer, 2008. doi:10.1007/978-3-540-88479-8_30.
- 128 Randall T. White, Christopher A. Healy, David B. Whalley, Frank Mueller, and Marion 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*, pages 192–202. IEEE Computer Society, 1997. doi:10.1109/RTAS.1997.601358.
- 129 Randall T. White, Frank Mueller, Christopher A. Healy, David B. Whalley, and Marion G. Harmon. Timing analysis for data and wrap-around fill caches. *Real-Time Systems*, 17(2-3):209–233, 1999. doi:10.1023/A:1008190423977.
- 130 Reinhard Wilhelm, Jakob Engblom, Andreas Ermedahl, Niklas Holsti, Stephan Thesing, David B. Whalley, Guillem Bernat, Christian Ferdinand, Reinhold Heckmann, Tulika Mitra, Frank Mueller, Isabelle Puaut, Peter P.uschner, Jan Staschulat, and Per Stenström. The worst-case execution-time problem – overview of methods and survey of tools. *ACM Trans. Embedded Comput. Syst.*, 7(3), 2008. doi:10.1145/1347375.1347389.
- 131 Reinhard Wilhelm, Daniel Grund, Jan Reineke, Marc Schlickling, Markus Pister, and Christian Ferdinand. Memory hierarchies, pipelines, and buses for future architectures in time-critical embedded systems. *IEEE Trans. on CAD of Integrated Circuits and Systems*, 28(7):966–978, 2009. doi:10.1109/TCAD.2009.2013287.
- 132 Lan Wu and Wei Zhang. A model checking based approach to bounding worst-case execution time for multicore processors. *ACM Trans. Embedded Comput. Syst.*, 11(S2):56, 2012. doi:10.1145/2331147.2331166.
- 133 Jun Yan and Wei Zhang. WCET analysis for multi-core processors with shared L2 instruction caches. In *Proceedings of the 14th IEEE Real-Time and Embedded Technology and Applications Symposium, RTAS 2008, April 22-24, 2008, St. Louis, Missouri, USA*, pages 80–89. IEEE Computer Society, 2008. doi:10.1109/RTAS.2008.6.
- 134 Patrick Meumeu Yomsi and Yves Sorel. Extending rate monotonic analysis with exact cost of preemptions for hard real-time systems. In *19th Euromicro Conference on Real-Time Systems, ECRTS'07, 4-6 July 2007, Pisa, Italy, Proceedings*, pages 280–290. IEEE Computer Society, 2007. doi:10.1109/ECRTS.2007.15.
- 135 Wei Zhang and Jun Yan. Accurately estimating worst-case execution time for multi-core processors with shared direct-mapped instruction caches. In *15th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications, RTCSA 2009, Beijing, China, 24-26 August 2009*, pages 455–463. IEEE Computer Society, 2009. doi:10.1109/RTCSA.2009.55.
- 136 Xiao Zhang, Sandhya Dwarkadas, and Kai Shen. Towards practical page coloring-based multicore cache management. In Wolfgang Schröder-Preikschat, John Wilkes, and Rebecca Isaacs, editors, *Proceedings of the 2009 EuroSys Conference, Nuremberg, Germany, April 1-3, 2009*, pages 89–102. ACM, 2009. doi:10.1145/1519065.1519076.