Integrating Path and Timing Analysis using Instruction-Level Simulation Techniques

Thomas Lundqvist and Per Stenström

Technical Report No. 98–3

Department of Computer Engineering
Chalmers University of Technology
S-412 96 Göteborg, Sweden

Göteborg, February 1998

Abstract

Previously published methods for estimation of the worst-case execution time on contemporary processors with complex pipelines and multi-level memory hierarchies result in overestimations owing to insufficient path and/or timing analysis. This not only gives rise to poor utilization of processing resources but also reduces the schedulability in real-time systems.

This paper presents a new method that integrates path and timing analysis to address these limitations. First, it is based on instruction-level architecture simulation techniques and thus has a potential to perform arbitrarily detailed timing analysis of hardware platforms. Second, by extending the simulation technique with the capability of handling unknown input data values, it is possible to exclude infeasible (or false) program paths in many cases, and also calculate path information, such as bounds on number of loop iterations, without the need for annotating the programs. Finally, in order to keep the number of program paths to be analyzed at a manageable level, we have extended the simulator with a path-merging strategy. This paper presents the method and particularly focuses on its path and timing analysis capabilities. It does this in the context of pipelined processors employing instruction and data caches.

1 Introduction

Static estimation of the worst-case execution time (WCET) has been identified as an important problem in the design of systems for time-critical applications. One reason is that most task scheduling techniques assume that such estimations are known before run-time to schedule the application tasks so that imposed timing constraints, e.g. deadlines, are met at run-time.

The WCET of a program is defined by the program path that takes the longest time to execute regardless of input data and initial system state. Ideally, a WCET estimation method should take as input a program and estimate a tight upper-bound on the actual WCET for a given hardware platform. There are two important sources to overestimations of the WCET. First, the estimation technique may encounter program paths that can never be executed regardless of the input data, usually referred to as *infeasible* (or false) program paths. Second, the timing model of the hardware platform may introduce pessimism because of simplifying timing assumptions.

The first problem can be addressed by requiring that the programmer provides path annotations [12] which clearly is both time-consuming and error-prone. A more attractive method is to automatically detect infeasible paths through static *path analysis methods* [3, 1]. As for the second problem, several *timing analysis* approaches have been proposed that statically estimate the execution time of a given path by taking into account the effects of, e.g., pipeline stalls and cache misses [7, 9, 4, 8, 6, 13, 11]. Unfortunately, since these timing analysis approaches typically charge fixed penalties caused by cache misses and pipeline hazards, it is difficult to handle the timing effect of dynamic interactions such as resource contention in the hardware platforms caused by buffering of instructions and memory requests. More seriously, no study to this date has shown how timing analysis methods can be integrated with path analysis methods to relieve the programmer from the burden of identifying and excluding infeasible paths from the analysis.

In this paper we present a new approach to WCET estimation that integrates path analysis with accurate timing analysis. Our WCET estimation method achieves this goal by (1) using instruction-level simulation techniques [10, 14] that conceptually simulate *all* feasible paths on arbitrarily detailed timing models of the hardware platform; (2) by eliminating many infeasible paths by extending architectural simulation techniques to handle unknown input data; and (3) by calculating bounds on number of loop iterations with no need for making annotations for statically known bounds.

A practical limitation of the method in its basic form is that the number of paths to simulate can easily be prohibitive, especially in loop constructs. We have therefore extended the architecture simulation method with a path-merging approach that manages to bound the number of simulated paths in a loop to just a few paths. This merging strategy reduces the number of simulated paths drastically and makes the approach useful for realistically sized programs. We apply the method to pipelined processors with instruction and data

caches and present how the merging strategy is implemented to encounter the worst-case effects of these features.

We start in Section 2 to put our work in the context of others. The next three sections are then devoted to describing our method focusing on the basic simulation approach in Section 3, the WCET algorithm in Section 4 and how detailed timing models of architectural features are integrated into the method in Section 5. An evaluation of the method focusing on the path analysis for seven test programs is presented in Section 6. In Section 7, we discuss the potential of the method along with its weaknesses before we conclude in Section 8.

2 Related Work

Because the prerequisites for tight estimations of the WCET are accurate path and timing analyses, let us review such proposed methods from the open literature.

First, in the absence of known input data, sufficient information must be available statically so that all execution paths have a finite length. We restrict the discussion to such programs. However, even if the WCET is statically decidable, infeasible (or false) program paths must be excluded from the analysis in order for the estimated WCET to be close to the actual WCET. Two fundamentally different approaches have been taken to address this problem. The first one is to rely on programmer-inserted annotations that provide such path information [12]. Clearly, letting the programmer provide this information is both time-consuming and error-prone. A more attractive method is therefore to exclude infeasible paths through static *path analysis*. One such method was recently proposed by Ermedahl and Gustafsson [3]. Their method, which is applied at the source-program level, establishes upper bounds on the number of loop iterations and also identifies infeasible paths. A similar technique was proposed by Altenbernd [1] who uses it to estimate the WCET of programs with fixed execution times for basic code blocks. While such path analysis methods can be used to estimate the WCET if tight upper-bounds on the execution time of high-level statements exist, WCET will be significantly overestimated for contemporary hardware platforms employing pipelining and caching. Therefore, such path analysis must be complemented with a timing analysis method that accurately estimates the execution time of the remaining (feasible) program paths, taking the detailed timing models of the hardware platform into account.

Contemporary microprocessors with complex pipelines and multilevel cache hierarchies pose a challenge for static timing analysis methods in that they must estimate the impact of, e.g., cache misses and pipeline stalls on the WCET. This problem has gained a lot of attention lately [7, 9, 4, 11]. Research has resulted in methods that can derive fairly tight estimates of single-issue pipelined processors with instruction caches. While there are differences between the approaches, all of them take as input the program flow graph

and statically analyze the program flow. In this process, the worst-case number of pipeline stalls and instruction cache misses are determined so as to estimate the WCET of each program path. More recently, these methods have also been extended to handle data caches [8, 6, 13, 11].

Unfortunately, these methods have only been demonstrated for pipelined processors with associated memory hierarchies that are much simpler than today's high-performance systems. Cache misses and pipeline hazards are typically modeled using fixed penalties, making it difficult to encounter dynamic interactions between hardware features, such as resource contention. This is a problem in high-performance systems, where we must account for the timing effects of, e.g., a write buffer between the first and second level cache and internal instruction queues in the processor. More seriously, available performance data shows that these methods do unnecessary overestimations because infeasible paths [6, 13] have not been eliminated from the analysis. No study so far has demonstrated how they perform when they are integrated with a path analysis method.

Our method, which is presented in the next three sections, integrates path and timing analysis in an efficient and accurate way to obtain tight WCET estimations. It thus has a potential to overcome the limitations of previous approaches.

3 The Approach

Consider programs for which the WCET is statically decidable; i.e., all possible execution paths in the program are finite in length. For example, bounds on number of iterations in loops are known before run-time, although they might be difficult to determine due to limitations in compiler analysis methods. Second, for the time being, we will consider processors with fixed instruction execution times although our method is applicable to arbitrary complex processor architectures with associated memory hierarchies which we show in Section 5. Given these assumptions, WCET could be conceptually determined by identifying the feasible path through the program with the longest execution time.

Instruction-level architectural simulation techniques have now matured so that the execution time of a program run on complex processor architectures can be accurately determined with a reasonable simulation efficiency [10, 14]. The advantages of using architectural simulation techniques are twofold. First, it is possible to make arbitrarily accurate estimations of the execution time of a program, given a combination of input data. Second, and presumably more importantly, when a given path through the program has been simulated, all input data independent dependencies will be resolved.

We first describe how traditional instruction-level simulation techniques can be extended to estimate WCET in Section 3.1. In Section 3.2 we show how the path analysis can be automated using this approach. Finally, in Section 3.3 we discuss some performance issues with the basic approach and how they are addressed.

3.1 Instruction-level simulation method

Instruction-level simulation techniques assume that input data is known and therefore only analyze a single path through the program associated with this input data. To find the WCET of a program using this approach, however, the program would have to be run with all possible combinations of input data which is clearly not feasible. Our approach is instead to simulate all paths through the program and in this process exclude the paths that are not possible regardless of input data. To do this, we have extended traditional instruction-level simulation techniques with the capability to handle unknown data and an extended semantics for each data-manipulating instruction to correctly perform arithmetics with the unknown data values as follows.

Each data type is extended with an element denoted *unknown*. In general, when this element is a source operand the destination operand will be unknown as well. Examples of the extended semantics for some common instruction types can be seen in Table 1. Consider for example an add instruction, ADD T, A, B, that operates on 32-bit unsigned integers, $Z := \{0 \ldots 2^{32} - 1\}$ and is defined as $T \leftarrow A + B$. In the extended semantics, it is instead defined on $\overline{Z} := Z \cup \{unknown\}$ with the semantics seen in Table 1.

Instruction type	Example	Semantics
ALU	ADD T,A,B	$T \leftarrow \begin{cases} unknown & \text{if } A = unknown \vee B = unknown \\ A + B & \text{otherwise} \end{cases}$
	AND T,A,B	$T \leftarrow \begin{cases} unknown & \text{if } A = unknown \wedge B \neq 0 \vee A \neq 0 \wedge B = unknown \\ 0 & \text{if } A = 0 \vee B = 0 \\ A \text{ and } B & \text{otherwise} \end{cases}$
Compare	CMP A,B	$A - B$ and update the condition code (cc) register. May set bits in the cc-register to *unknown*.
Conditional branch	BEQ L1	Test bit in cc-register to see if we should branch. If bit is *unknown* simulate both paths.
Load	LD R,A	Copy data from memory at address A to register R (the data can be *unknown*). If address is *unknown*, set R to *unknown*.
Store	ST R,A	Copy data from register R to memory at address A (the data can be *unknown*). If address is *unknown*, all memory locations are assigned *unknown* (a more efficient solution is discussed in Section 3.3).

Table 1: Extended semantics of instructions.

The load and stores need special treatment, since the reference address used may be *unknown*. For loads, this results in an *unknown* being loaded into a register. For stores, however, an unknown address can affect an arbitrary memory location. Therefore, the correct action would be to assign the value *unknown* to all memory locations to capture the worst-case situation. This is of course a major limitation. For the time being, we will

assume that all addresses are statically known. However, we will discuss efficient solutions to overcome this limitation later in Section 3.3.

The semantics for a conditional branch is also special. When a conditional branch whose branch condition is *unknown* is encountered, both paths must be simulated. An important consequence of this is that when the branch condition is known, the extended simulation technique will exclude paths that can never be taken. Let's review this path analysis capability in some more detail in the next section.

3.2 Path analysis

To understand how the extended simulation technique can automate path analysis, consider the program below which sums values in the upper-right triangle of matrix b.

```
  int fun(int b[4][4])
  {
1   int i, j, sum = 0;
2
3   for (i = 0 ; i < 4 ; i++)
4     if (b[i][i] > 0)
5       for (j = i+1 ; j < 4 ; j++)
6         sum += b[i][j];
7   return sum;
  }
```

For simplicity, we reason about the program in a high-level language, even if the simulation of course is done at the instruction level. In the beginning of the simulation, the data values in matrix b is treated as unknown input and all elements are assigned the value *unknown*. When the simulation is started, execution will continue until line 4 is reached where the branch condition of the conditional branch corresponding to the if-statement depends on an unknown value. Consequently, the two possible execution paths originating from this conditional branch have to be simulated. One path continues through lines 3,4, the other one through lines 5,6...3,4, until they hit the same unknown conditional branch again during the second iteration of the outer loop. Continuing in this way, we will finally have $2^4 = 16$ paths that reach the end of the program. The total WCET is the maximum one found among these paths.

An important advantage of the instruction-level simulation approach is that all conditions that depend on data values that are known statically will be computed during the simulation. For example, input independent bounds on the number of loop iterations that are expressed with complex functions that are difficult for static compiler analysis methods to evaluate, will be computed in this method. This eliminates the need for program annotations to, e.g., bound the number of loop iterations. Moreover, many paths that cannot be executed, regardless of the input data, will be excluded from the analysis and will therefore not affect the WCET estimation.

3.3 Discussion

Ideally, for each combination of input data, the goal is to eliminate all infeasible program paths. The domain of values we use, however, makes it sometimes impossible to correctly analyze mutually exclusive paths. Consider for example the statements below where b is *unknown*.

```
if (b < 100)
  fun1()
if (b > 200)
  fun2()
```

Variable b will be *unknown* in both conditions, forcing the simulation of four paths even if `fun1()` and `fun2()` are mutually exclusive and the number of feasible paths are three. The effect of this will be a potential overestimation of the WCET which is caused by the fact that we only distinguish between known and unknown values. However, there is nothing in the approach that would hinder us from extending the domain of values to, e.g., ranges of values. By doing this, it would be possible to handle cases like the one above. The simulation technique we have implemented, and which we evaluate later in the paper, uses the simpler domain. However, we will discuss the implications of using ranges in Section 7.

As mentioned above, store instructions with an unknown reference address need special treatment. To efficiently handle these accesses, we identify all data structures that are accessed with unknown store instructions as *unpredictable data structures*. These are mapped by the linker into a memory area which can only respond with unknown values. This means that stores to this area can be ignored and loads will always return *unknown* as a result, thus assuring us of a correct estimation of the WCET with no added cost in simulation time.

We identify the unpredictable data structures by letting the simulator output a list of all unknown stores it encounters. With the help of a source-code debugger, it is possible to connect each store instruction to a data structure which is marked as unpredictable. Eventually, the linking phase is redone to properly map the marked data structures to the 'unknown' area of memory. After this step, a correct estimation may be done. The approach mentioned above can be used for statically allocated data structures only. For unpredictable dynamically allocated data structures we do not yet have any solution.

A key problem with the simulation approach in this section is the explosion in number of paths to be simulated. If a loop iteration has n feasible paths and the number of loop iterations is k, the number of paths to simulate is n^k. Fortunately, good heuristics exist to drastically reduce the number of paths we need to simulate. We have used a path merging strategy, which forms the basis of our WCET method presented in the next section.

4 The WCET Method

To reduce the number of paths that have to be explored in the path analysis, we apply a path merge strategy. This reduces the number of paths to explore from n^k down to n for a loop containing n feasible paths in each iteration and doing k iterations. In each loop iteration, all n paths are explored but in the beginning of the next iteration all these n paths are merged into one. Thus, the number of simulated paths are less than or equal to n. We first describe how the merging operation is performed in Section 4.1. This operation is used in the WCET algorithm which we present in Section 4.2. Finally, in Section 4.3 we discuss how to implement the method with a reasonable time complexity.

4.1 Merging algorithm

In order to understand how the merging of two paths is done, consider the example program in Figure 1 running on a simple architecture with fixed instruction execution times. When the loop is entered, the data in array `b[]` is considered to be *unknown*. During the first iteration of the loop, both paths are simulated. In the second iteration of the loop, when the simulator encounters the unknown conditional branch on line 2, two paths can be merged; p_A going through lines 1,2,3,1,2, and p_B going through lines 1,2,5,1,2. If the current WCET of p_A and p_B are 30 and 20, respectively, the longer path p_A is chosen and p_B can be discarded throughout the simulation. However, to make a valid estimation of the worst-case execution path throughout the program execution, the impact of p_B on the total WCET must be taken into account. For example, the variable `sum` can be assigned different values in the two paths. Therefore, all variables whose values are different, or are *unknown*, must be assigned the value *unknown* in the resulting path of the merge operation.

Formally, the algorithm views a path, p_A, as consisting of a WCET for the path, $p_A.wcet$, and the system state at the end of the path, $p_A.state$. The state of a path is the partial result of the execution including, e.g., the content of all memory locations, registers, and status bits that can affect the WCET in the future. In order to compute the state of the path resulting from the merging of several paths, the state of the merged paths is compared and *unknown* values are assigned to locations that contain either *unknown* or values that differ in the state of the merged paths. We denote this operation as the *union* operation of path states.

The merging algorithm is described in Algorithm 1. It creates a new path p_C from two paths p_A and p_B. The program counter, which must not differ in the two paths merged, is copied to the new path. The new WCET is the maximum of the WCET of the two original paths. Finally, states of the merged paths are unioned. The union operation between values is defined as:

$$a \cup b \Leftrightarrow \begin{cases} unknown & \text{if } a = unknown \vee b = unknown \\ unknown & \text{if } a \neq b \\ a & \text{otherwise} \end{cases}$$

```
1    for (i = 0 ; i < 100 ; i++)
2      if (b[i] > 5)
3        sum = sum + a * i
4      else
5        sum = sum + 1
```

Figure 1: Loop example.

Algorithm 1 Merging two paths p_A and p_B creating p_C.

{PCT = program counter, CC = condition code}
Require: $p_A.state.PCT = p_B.state.PCT$
$p_C.state.PCT \leftarrow p_A.state.PCT$

$p_C.wcet \leftarrow \max(p_A.wcet, p_B.wcet)$

$p_C.state.CC \leftarrow p_A.state.CC \cup p_B.state.CC$
for all registers i **do**
$p_C.state.R[i] \leftarrow p_A.state.R[i] \cup p_B.state.R[i]$
end for
for all memory positions a **do**
$p_C.state.MEM[a] \leftarrow$
$p_A.state.MEM[a] \cup p_B.state.MEM[a]$
end for

4.2 WCET algorithm

In order to implement the WCET simulation technique and the merging algorithm, one important issue is in which order paths should be simulated. Consider a loop with two possible paths in each iteration. In order to merge these two paths, they must have been simulated the same number of iterations. To accomplish this, the WCET algorithm uses loop information from the control flow graph of the program.

The algorithm (see Algorithm 2) starts the simulation from the beginning of the program. Whenever an unknown conditional branch is found, the simulation is stopped and the algorithm selects, as the next path to simulate, the path that has made the least progress, *a minimum progress path*, in terms of loop iterations. If this path is not unique, all paths that have made the same progress and are at the same position in the program, *equal progress paths*, are merged into one before simulation is continued.

By always selecting the path that has made least progress, the algorithm makes it certain that all paths in a loop iteration is simulated before a new iteration begins. In fact, merging will occur every time two paths that have made equal progress meet at the same position in the program. This makes an exponential growth of the number of paths impossible.

Algorithm 2 WCET algorithm handling merge.

```
{A is set of active paths, C completed paths}
{∅ = empty set, \ = set minus}

A ← ∅
C ← ∅

p ← starting null path
Simulate(p)
if p reached end of program then
  C ← C ∪ {p}
else {p reached an unknown conditional branch}
  A ← A ∪ {p}
end if

while A not empty do
  p ← minimal progress path in A
  A ← A \ {p}
  for all paths q with equal progress as p do
    A ← A \ {q}
    p ← merge(p, q)
  end for
  {Path p ends with a branch forcing a split}
  for each possible branch target i do
    p_i ← copy of p
    Simulate(p_i) along target i
    if p_i reached end of program then
      C ← C ∪ {p_i}
    else {p_i reached an unknown conditional branch}
      A ← A ∪ {p_i}
    end if
  end for
end while

wcet ← max_{p∈C} p.wcet
```

4.3 Performance considerations

The performance of this method depends on how many paths we need to simulate, how often merging is done and how fast the actual merge operation is. How often the merging should be done is a complex question. If we merge often, it is likely that we will have fewer paths to simulate. However, when doing a merge we can lose information in a way that makes us fail to eliminate infeasible paths. Thus, merging less often can lead to a fewer number of simulated paths.

The algorithm we have evaluated, merge as often as possible. This makes an exponential

growth of the number of paths impossible, e.g. in loops. Typically, each unknown branch found during the simulation would lead to the creation of yet another path, which later would result in an additional merge operation. The total number of paths simulated, as well as the number of merge operations, can in some cases grow in proportion to the number of loop iterations done in the program. One example is a loop with one unknown exit condition. The simulation of this loop would produce one new path (the exit path) to simulate each iteration. All these paths would, unless they reach the end of the program, be merged resulting in an equal number of merge operations.

For each merge operation, one must union the content of all registers and memory locations. This might be a quite slow process if the amount of memory is large. We can speed up this operation considerably by utilizing the fact that paths that are to be merged, often shared a long history of execution before they got split up. By only recording changes made to the system state since the time where the two paths were created, we can quickly identify the parts of memory where the two system states differ. As an example, suppose that the system we model contains 1 Mbyte of main memory. Then one can divide this memory into small fixed size pages (say 512 bytes each) and each path only keeps the accessed pages in its system state. In this way, only a few pages of memory need to be compared during merging.

5 Timing Analysis

The WCET algorithm in the previous section can estimate WCET for hardware platforms with fixed instruction execution times. In this section, we extend it to model the timing of pipelined processors with caches. To demonstrate the approach, we particularly focus on systems with separate, direct-mapped instruction and data caches and a single-issue pipelined execution unit. However, the approach extends to more sophisticated architectures.

In order to update the WCET properly during simulation, the simulator must of course be extended to model the timing of caches and pipelines. In the context of caches, the simulator must model the impact of cache misses on the execution time. And in the context of pipelines, the simulator must account for the impact of structural, data, and control hazards [5] on the execution time. With this capability, it is possible to make a safe estimation of the execution time of a given path through the program.

A critical issue is how to carry out a merge operation. To do this, one must be able to estimate the impact of the system state on the future execution time. Such state information is exemplified by the content of the tag memory in the caches (which affects future misses) and resources occupied by an instruction, such as data-path components and registers (which affects future structural and data hazards). The merging operation introduced in Section 4 must be extended to handle such state information, which we will

refer to as *timing state*.

To merge the timing state, we could use the same general principle as used when merging the content of memory locations; for all locations where the timing state differs we assert a pessimistic value, such as *unknown* in the merged timing state. For example, in the case of most caches *unknown* means that a cache block is invalid, and the unioning of two cache timing states makes all cache blocks whose tags differ invalid. We call this method the *pessimistic merge*, since it incurs a severe pessimism in the estimation. Fortunately, a method will be presented in the next section that avoids this pessimistic merge operation in many cases. We first present the general method in the next section and then apply it to caching and pipelining in Sections 5.2 and 5.3, respectively.

5.1 Optimistic merge approach

Our approach to reduce the pessimism when merging, is based on the idea that if it was known in advance which path belongs to the worst case path through the program, one could update the worst-case execution time with the execution time of that path and also choose the timing state of that path when merging and then discard the timing state of the other path. Then, no pessimism would be incurred on a merge operation. While it is not possible to know in advance which of the two paths belongs to the worst-case path, a good guess would be that the longer of the two belongs to the worst-case path. If we also estimate how big effect the timing state of the shorter path has on the future execution time, we can make sure whether it is correct to use the WCET of the longer path along with its timing state when merging two paths. This approach is formulated in the following algorithm where we assume that the worst-case execution times of the long and the short paths are $WCET_L$ and $WCET_S$, respectively.

1. Estimate the *worst-case penalty* ($WCET_P$) that the short path would incur on the future execution time that will not be incurred by the long execution path.
2. If $WCET_P < WCET_L - WCET_S$ then use the timing state and WCET of the long path in the merge operation.
3. Otherwise, the pessimistic merge approach must be used.

In order to make this algorithm useful, we must clearly define what we mean with timing state and worst-case penalty. This is done in the next two sections in the context of caches and pipelining.

5.2 Instruction and data cache analysis

In the following, we first describe how to estimate the future effect on the execution time when discarding the cache state of the short path. Also, in order to handle the case when

it is not possible to discard, we need to define how to do a pessimistic merge between two cache states. Both instruction and data caches can be described and treated using the same principles. Therefore, we make no distinction between them.

Consider the timing state of two caches, C_L and C_S, belonging to the long and short path, respectively. A cache state is represented by an array of tags showing how blocks are currently mapped in the cache. To estimate the possible effect on the future execution time of discarding C_S, we must consider all cases where C_S may lead to a greater number of future cache misses than C_L would. The worst case is found if we imagine that all cache blocks resident in C_L but not found in C_S will be needed in the future. We would then lose all these potential misses if we discard C_S. To find the total possible loss of execution time when discarding C_S, we go through all cache blocks and compare the tag in C_S with the tag in C_L and sum up the cache miss penalties caused by all differences found. If the accumulated penalty is smaller than the difference in WCET between the two paths, no merging is needed.

While pessimistic merging might be needed if it is not possible to discard C_S, it is not always necessary to invalidate all blocks found different. By invalidating a block in C_L we make it impossible for C_S to cause any additional misses in the future. Thus, for each block we invalidate in C_L we reduce the additional penalty that C_S may incur compared to C_L. Eventually, the worst-case penalty may became small enough to proceed and actually discard C_S.

5.3 Pipeline analysis

Because it is straight-forward to take into account timing effects of pipelines in an architectural simulator [14], we will here focus on an approach to define the timing state needed to make an optimistic (or pessimistic) merge operation.

A pipeline will incur penalties caused by resource contention (structural hazards) and register dependencies (data hazards). An approach to estimate the worst-case effect these hazards can have on the future execution time would be to associate with the resources (data-path components and registers) timing state that indicates when the resources will be released in the worst case. For example, when a register is used as destination operand, the timing state of this register is updated with time information when the operand will be available in the worst case. The worst-case penalty a path may incur on the future execution time is the time it takes for the last resource to be freed. Control hazards can be resolved in isolation for each instruction and will not influence any future instruction.

Consider now the timing state of two pipelines, P_L and P_S, belonging to the long and short path, respectively. We want to estimate the possible future effect on the execution time that P_S may lead to compared to P_L, and must in this case consider all future hazards that P_S can lead to, which P_L cannot lead to. For each data-path component and register we determine whether P_S can lead to a greater delay due to hazards than P_L would. The

maximum difference in delay found is the worst-case penalty that P_S, but not P_L, would incur on the future execution time. If this penalty is smaller than the difference in WCET between the two paths, no merging is needed and we can discard P_S.

If we are not allowed to discard P_S, we must do the pessimistic merge instead. This can be done using the same principle as for the caches; we change P_L in a way that reduces the possible future effect on the execution time that P_S may lead to compared to P_L. For each possible structural or data hazard resulting from P_S but not from P_L, we can change P_L in a way that makes it as bad as P_S. In this way we can gradually reduce the additional future execution time that P_S may lead to, until at last, we are allowed to discard P_S.

We are currently implementing the approach described in this section in a simulator with the purpose of evaluating the accuracy of the WCET estimations for a PowerPC system.

6 Experimental Results

We have estimated the WCET of seven benchmark programs, running on an idealized architecture with no caches, and where all instructions execute in one clock cycle. Thus, the evaluation focuses on the path analysis aspects of the method.

6.1 Methodology

A WCET simulator has been constructed by extending an existing instruction-level simulator, PSIM [2], which simulates the PowerPC instruction set. The original simulator has been extended with the capability of handling unknown values. Also, the WCET algorithm described in Section 4.2 has been added to control the path exploration and merging. No cache and pipeline simulation was enabled and the execution time is equal to the simulated instruction count.

The GNU compiler (gcc 2.7.2.2) and linker has been used to compile and link the benchmarks. No optimization was enabled. The simulated run-time environment contains no operating system; consequently, we disabled all calls to system functions like disk I/O and the `printf` function in the benchmarks.

6.1.1 Benchmarks and metrics

An overview of the seven benchmark programs can be seen in Table 2. There are four small programs: *matmult*, *bsort*, *isort*, and *fib*, and three larger programs: *DES*, *jfdctint*, and *compress*.

The two benchmarks *fib* and *compress* both contain a loop where the exit condition depends on unknown input data. In order to bound the number of iterations when simulating, we need to add manual annotations. This is not supported in our implementation.

Name	Description
matmult	Multiplies 2 10x10 matrices
bsort	Bubblesort of 100 integers
isort	Insertsort of 10 integers
fib	Calculates n:th element of the Fibonacci sequence for $n \leq 30$
DES	Encrypts 64-bit data
jfdctint	Does a discrete cosine transform of an 8x8 pixel image
compress	Compresses 50 bytes of data (downscaled version of compress from SPEC CPU95 benchmark suite)

Table 2: The benchmark programs used.

Instead, we have added an extra exit condition in the loops. In *fib* we have added the condition: $i \leq 30$ because we know that input data is always in this range. In *compress* we bound an inner loop whose iteration variable is j, using the current iteration count, i, of the outer loop: $j \leq i$. This is a safe but very pessimistic bound, but we found it difficult to prove that a tighter bound could be used.

The true WCETs of all programs have been measured by running the programs in the simulator with the worst case input data. This works fine for all programs except *compress*, where the worst case input data is hard to find. Instead, a random sequence of 50 bytes has been used as input.

Besides the estimated WCET from the extended simulator, we also did a manual estimate of the structural WCET, i.e., the execution time of the longest structural path in the control flow graph of the program and using fixed bounds on the number of iterations of all loops. This figure would represent a WCET estimation method that eliminates no infeasible paths and uses fixed iteration bounds for loops. The purpose of doing this is to analyze the capability of the method to eliminate infeasible paths.

6.2 Estimation results

Table 3 compares the measured true WCET with the estimated WCET from the simulator and the manually derived structural estimation for the benchmark programs (WCET is expressed in clock cycles, ratio is the WCET relative to the measured, true WCET).

For all benchmarks, except *compress*, we find that the method succeeds in finding the true WCET. In *compress*, the overestimation is caused by the inner loop. As mentioned previously, we bound this loop using the very pessimistic condition $j \leq i$, but when measuring the true WCET, we found that this inner loop is actually only doing one single iteration. For *compress*, we do not know if the measured WCET is actually the true WCET. The true WCET is probably higher than the measured one.

Two of the benchmarks, *matmult* and *jfdctint*, have no infeasible paths at all, and only one path was simulated. In *DES* however, there exist infeasible paths caused by data

Program	Measured true WCET	Estimated WCET	Ratio	Estimated structural WCET	Ratio
matmult	7063912	7063912	1	7063912	1
bsort	292026	292026	1	572920	1.96
isort	2594	2594	1	4430	1.71
fib	697	697	1	697	1
DES	118675	118675	1	119877	1.01
jfdctint	6010	6010	1	6010	1
compress	9380	49046	5.2	161161	17.2

Table 3: The estimated and true WCET of benchmark programs.

dependencies between different functions. These infeasible paths were eliminated and only one path was simulated. In *bsort* and *isort*, all infeasible paths were not eliminated. Still, this did not lead to any over-estimation, since all simulated infeasible paths were shorter than the worst-case path found.

If we take a look at the estimated structural WCET of the programs, we see that the WCET is grossly over-estimated for *bsort*, *isort*, and *compress*. In *bsort* and *isort* it depends entirely on using a fixed iteration count for an inner loop which is normally bounded by the outer loops current iteration count. This leads to an over-estimation of a factor of two for the loop and influences *bsort* more than *isort* because of the greater number of iterations done in *bsort*. In *compress* there is a similar inner loop which is forced to have a fixed iteration bound again causing an over-estimation of a factor of two. In addition, there exists a very long infeasible path that extends the structural estimate. This path is eliminated when estimating with our simulation method. The tiny over-estimation in *DES* results from infeasible paths that are not possible to eliminate.

One strength of doing the analysis on the instruction level can be seen when looking at *DES*. In the source code, one can find several conditional expressions which should indicate several possible feasible paths through the program. However, the compiler (gcc with no optimization enabled) automatically generates code without any branches for these conditional expressions and the resulting program has only a single feasible path. In summary, the method in this paper appears promising in eliminating infeasible paths.

7 Discussion

As the results indicate, to get a tight estimation of the WCET of a program, it is crucial that we eliminate infeasible paths, especially in the presence of loops with dynamic bounds. Of equal importance is an accurate timing analysis. This has not been evaluated in this study, but the potential can be seen when looking at *DES*, *matmult*, and *jfdctint*, where

the simulator only needs to simulate one path through the program. This path can be simulated with an arbitrary detailed timing model and will always give us the true WCET. Thus, by eliminating infeasible paths we can concentrate on the feasible ones, and make a more accurate timing analysis.

A big advantage of integrating the path and timing analysis can be seen when comparing with other solutions where the path and timing analyses are kept separated. If the path analysis is done first, we would need a way to represent the path information generated from the path analysis, and the timing analysis phase must be able to utilize this information. On the other hand, if the timing analysis is done first, we would be forced to work with fixed WCETs for blocks of statements when doing the path and WCET calculation. These problems are not present in our method, which does the path and timing analysis simultaneously.

The simulator uses, in our method, a very simple domain where values can be either known or unknown. This domain performs quite well compared to a more complex domain, e.g., based on intervals of values which is used by [3, 1]. However, as we saw in Section 3.3, overestimations may sometimes arise for mutually exclusive paths. There is no inherent problem in extending our method to a more powerful domain. The result would only be a slower simulator needing more memory. Our choice of domain results in an additional 1 bit of memory for each 32-bit word of memory to hold the known/unknown status. An interval representation would need 2 extra words for each word of memory. Also, a more complex semantics would be needed, which would result in a slower execution of each instruction. On the other hand, the more complex domain might be preferable for some applications, if it manages to cut more infeasible paths and thereby gain speed and accuracy compared to our simple domain.

A more serious problem with our simple domain is that if all exit conditions of a loop is input data dependent, we get a completely unknown upper bound on the number of iterations in the loop. In this case we must use a manual annotation or add a known exit condition by modifying the loop condition in the program. For example, the loop in the program below, where b is unknown input data, will never do more than 100 iterations regardless of b. This fact cannot be represented with our simple domain, and during the simulation the loop will get a completely unknown exit condition, forcing us to add annotations or modify the program.

```
if (b < 100)
  for (i = 0 ; i < b ; i++)
    sum = sum + i;
```

In the example above the simple domain was causing the problem. A similar problem can also arise when merging. The union operation used when merging may cause information needed to bound a loop to be lost. If this happens, we are also forced to annotate or change the program.

Our method uses the same basic idea regarding path analysis as the method used by Altenbernd in [1], but he did not integrate any timing analysis. His method uses precalculated execution times on each basic code block and uses this to prune paths during the path exploration. However, in some cases, the method suffers from complexity problems. We use a path merging strategy instead, which guarantees a manageable number of paths, and also makes it possible to integrate the timing analysis.

8 Conclusions and Future Work

In this paper we have presented a new method for estimating the WCET of a program. This method integrates path and timing analysis and thereby has the potential to do tight estimations by eliminating infeasible paths and concentrating the timing analysis on the feasible ones.

Our results concerning the path analysis aspects show that many infeasible paths were indeed eliminated in the studied benchmark programs, but it remains to be shown how well the method performs when considering the timing analysis as well. This is the topic of our current and future work. Currently, we are estimating the WCET of the benchmark programs when run on a system with a data cache and a pipeline. The evaluation of this system will be the topic of a forthcoming paper.

Acknowledgments

We are deeply indebted to Dr. Jan Jonsson of Chalmers for his constructive comments on previous versions of this manuscript. This research is supported by a grant from Swedish Research Council on Engineering Science under contract number 221-96-214.

References

[1] P. Altenbernd. On the false path problem in hard real-time programs. In *Proceedings of the 8th Euromicro Workshop on Real-Time Systems*, pages 102–107, June 1996.

[2] A. Cagney. PSIM, a PowerPC simulator. ftp://ftp.ci.com.au/pub/psim/index.html.

[3] A. Ermedahl and J. Gustafsson. Deriving annotations for tight calculation of execution time. In *Proceedings of EUROPAR'97*, pages 1298–1307, August 1997.

[4] C. A. Healy, D. B. Whalley, and M. G. Harmon. Integrating the timing analysis of pipelining and instruction caching. In *Proceedings of the 16th IEEE Real-Time Systems Symposium*, pages 288–297, December 1995.

[5] J. L. Hennessy and D. A. Patterson. *Computer Architecture: A Quantitative Approach, 2ed.* Morgan Kaufmann, 1996.

[6] S.-K. Kim, S. L. Min, and R. Ha. Efficient worst case timing analysis of data caching. In *Proceedings of the 2nd IEEE Real-Time Technology and Applications Symposium*, pages 230–240, June 1996.

[7] Y.-T. S. Li, S. Malik, and A. Wolfe. Efficient microarchitecture modeling and path analysis for real-time software. In *Proceedings of the 16th IEEE Real-Time Systems Symposium*, pages 298–307, December 1995.

[8] Y.-T. S. Li, S. Malik, and A. Wolfe. Cache modeling for real-time software: Beyond direct mapped instruction caches. In *Proceedings of the 17th IEEE Real-Time Systems Symposium*, pages 254–263, December 1996.

[9] S.-S. Lim, Y. H. Bae, G. T. Jang, B.-D. Rhee, S. L. Min, C. Y. Park, H. Shin, K. Park, and C. S. Kim. An accurate worst case timing analysis technique for RISC processors. In *Proceedings of the 15th IEEE Real-Time Systems Symposium*, pages 97–108, December 1994.

[10] P. Magnusson, F. Dahlgren, H. Grahn, M. Karlsson, F. Larsson, A. Moestedt, J. Nilsson, P. Stenström, and B. Werner. Simics/sun4m: A virtual workstation. In *Proceedings of USENIX98*, 1998. To appear.

[11] G. Ottosson and M. Sjödin. Worst-case execution time analysis for modern hardware architectures. In *Proceedings of ACM SIGPLAN Workshop on Language, Compiler, and Tool Support for Real-Time Systems*, June 1997.

[12] P. Puschner and C. Koza. Calculating the maximum execution time of real-time programs. *The Journal of Real-Time Systems*, pages 159–176, 1989.

[13] R. T. White, F. Mueller, C. A. Healy, D. B. Whalley, and M. G. Harmon. Timing analysis for data caches and set-associative caches. In *Proceedings of the 3nd IEEE Real-Time Technology and Applications Symposium*, pages 192–202, June 1997.

[14] E. Witchel and M. Rosenblum. Embra: Fast and flexible machine simulation. In *Proceedings of ACM SIGMETRICS '96*, pages 68–79, 1996.