INRIA

PROARTIS: Probabilistically Analysable Real-Time Systems

Francisco J. Cazorla[1,5] — Eduardo Quiñones[1] — Tullio Vardanega[2] — Liliana Cucu[3] — Benoit Triquet[4] — Guillem Bernat[6] — Emery Berger[1,7] — Jaume Abella[1] — Franck Wartel[4] — Michael Houston[6] — Luca Santinelli[3] — Leonidas Kosmidis[1] — Code Lo[3] — Dorin Maxim[3]

[1] Barcelona Supercomputing Center.

[2] University of Padua

[3] Institut national de recherche en informatique et en automatique

[4] Airbus France

[5] Spanish National Research Council (IIIA-CSIC)

[6] Rapita Systems

[7] University of Massachusetts Amherst

N° 7869

26 january 2012

Embedded and Real Time Systems

Rapport de recherche

PROARTIS: Probabilistically Analysable Real-Time Systems

Francisco J. Cazorla[1,5] , Eduardo Quiñones[1] , Tullio Vardanega[2] ,
Liliana Cucu[3] , Benoit Triquet[4] , Guillem Bernat[6] , Emery Berger[1,7] ,
Jaume Abella[1] , Franck Wartel[4] , Michael Houston[6] , Luca Santinelli[3] ,
Leonidas Kosmidis[1] , Code Lo[3] , Dorin Maxim[3]

[1] Barcelona Supercomputing Center.
[2] University of Padua
[3] Institut national de recherche en informatique et en automatique
[4] Airbus France
[5] Spanish National Research Council (IIIA-CSIC)
[6] Rapita Systems
[7] University of Massachusetts Amherst

Abstract: Static Timing Analysis is the state-of-the-art practice to ascertain the timing behaviour of current-generation real-time embedded systems. The adoption of more complex hardware to respond to the increasing demand for computing power in next-generation systems exacerbates some of the limitations of Static Timing Analysis. In particular, the effort of acquiring (1) detail information on the hardware to develop an accurate model of its execution latency as well as (2) knowledge of the timing behaviour of the program in the presence of varying hardware conditions, such as those dependent on the history of previously executed instructions. We call these problems the Timing Analysis Walls.

In this vision-statement paper we present Probabilistic Timing Analysis, a novel approach to the analysis of the timing behaviour of next-generation real-time embedded systems. We show how Probabilistic Timing Analysis attacks the Timing Analysis Walls; we then illustrate the mathematical foundations on which this method is based and the challenges we face in the effort of efficiently implementing it. We also present experimental evidence that shows how Probabilistic Timing Analysis reduces the extent of knowledge about the execution platform required to produce probabilistically-safe and tight WCET estimations.

Key-words: timing analysis

PROARTIS: Probabilistically Analysable Real-Time Systems

Résumé : Notre article présente une analyse probabiliste de durée d'exécution d'un programme. Nous présentons les bases mathématiques de notre approche et les défis à surmonter pour rendre l'approche applicable dans des cas réels.

Mots-clés : durée d'exécution

Contents

1 Introduction

1.1 Problem Statement

The market for Critical Real-Time Embedded Systems (CRTES), which among others includes the avionics and automotive sectors, is experiencing an unprecedented growth, and is expected to continue to steadily grow for the foreseeable future [8]. Let us for instance look at the automotive domain: a state-of-the-art high-end car, which currently embeds up to 70 Electronic Control Units (ECUs), is predicted to embed many more [7] to account for the inclusion of such new increasingly sophisticated functions as Advanced Driver Assistance Systems (ADAS). For CRTES of this kind it is imperative to ensure the timing correctness of system operation: some form of *Worst-Case Execution Time* (WCET) analysis is needed to that end.

The competition on functional value, measured in terms of application services delivered per unit of product faces CRTES industry with rising demands for greater performance, increased computing power, and stricter cost-containment. The latter factor puts pressure on the reduction in the number of processing units and ECUs used in the system, to which industry responds by looking at more powerful processors, with aggressive hardware acceleration features like caches and deep memory hierarchies.

In this evolving scenario, it must be acknowledged that the industrial application of current WCET analysis techniques [28], which accounts for a significant proportion of total verification and validation time and cost of system production, yields far from perfect results. IBM has for example found that 50% of the warranty costs in cars are related to electronics and their embedded software, and that 30% of those costs are related to timing flaws. These instances of incorrect operation cost industry billions of Euros annually [7].

Current state-of-the-art timing analysis techniques can be broadly classified in two complementary strands [28]: static timing analysis; and measurement-based analysis.

Measurement-based analysis techniques rely on extensive testing performed on the real system under analysis using stressful, high-coverage input data, recording the so-called high watermark execution time, i.e. the longest observed execution time; and adding to it an engineering margin to make safety allowances for the unknown. However, the safeness of the engineering margin is extremely difficult – if at all possible – to determine, especially when the system may exhibit discontinuous changes in timing due to pathological cache access patterns or other unanticipated timing behaviour.

Static timing analysis techniques rely instead on the construction of a cycle-accurate model of the system and a mathematical representation of the application code which makes it possible to determine the timing behaviour on that model. The mathematical representation is then processed with linear programming techniques to determine a safe upper-bound on the WCET, providing stronger guarantees than measurement-based approaches. Static approaches have one important limitation though: they are expensive to carry out owing to the need to acquire exhaustive knowledge of all factors, both hardware and software, that determine the execution history of the program under analysis. Some processor architectures may dramatically increase this cost. Others, possibly subject to intellectual property restrictions or incomplete documentation, may even make it altogether impossible, and construction of the timing model must resort to observations.

In order to appreciate the complexity of acquiring complete knowledge of execution history, consider a cache model with a Least Recently Used (LRU) replacement policy. The accuracy in predicting the hit/miss outcome of a memory access depends

on knowing the full sequence and addresses of the previous memory accesses made by the program up to the point of interest, in order to build a complete and correct representation of the cache state. Any reduction of the available knowledge, e.g. when the addresses of some memory accesses are unknown, leads to a rapid degradation of the tightness of the WCET estimation. In fact, partial knowledge can lead to results as inaccurate as those obtained with no information at all.

1.2 PROARTIS Vision and Approach

Our view is that the cost of acquiring the required knowledge to perform trustworthy analysis can be significantly reduced by adopting a hardware/software architecture whose execution timing behaviour eradicates dependence on execution history. One way to achieve this independence is via introducing randomisation in the timing behaviour of the hardware and possibly of the software (while the functional behaviour is left unchanged), coupled with new probabilistic analysis techniques. An example of such hardware is a cache memory in which, in the event of a miss, the victim is randomly selected, from any position in the cache. We call this unit of eviction/replacement, cache entry. Under this cache configuration, the probability of hit/miss for an access has a small dependence on execution history, in particular, the number of cache misses between the access under consideration and the previous access to the same address. Note that the hit/miss probability is different from the frequency of events. For instance, a memory instruction may have a 50% hit probability if every time it accesses cache we flip a coin and hit if and only if we get heads. Conversely, if the instruction hits and misses alternately, that instruction does not have a 50% hit probability but a 50% hit frequency. This is so because the outcome, and hence the latency, of each access is fully deterministic.
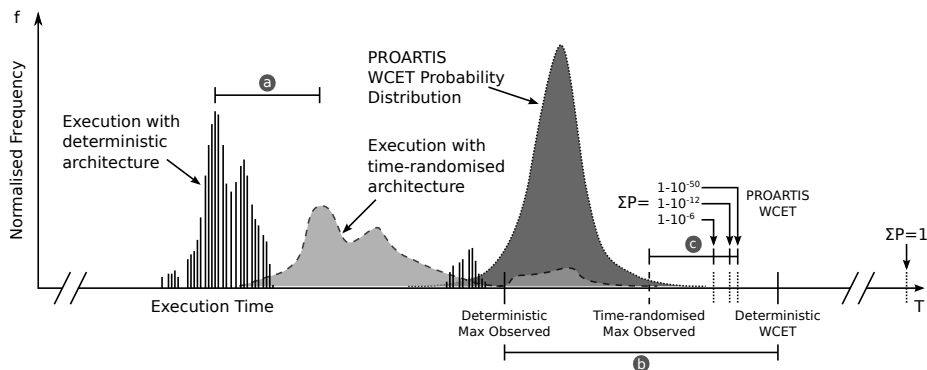

Figure 1: Execution time distributions for conventional deterministic architectures and a proposed time-randomised architecture, superimposed with the PROARTIS worst-case probability distribution

Applying time-randomisation techniques has inevitable consequences for the average-case execution time of a program. Figure 1 illustrates the PROARTIS view of execution time; the execution time shift between deterministic and time-randomised architectures is marked (a). In general we expect the time-randomized profile to shift to the right (to increase) in relation to the deterministic profile, and to spread across a larger range of execution times. However the resulting distribution becomes more predictable: by decoupling timing events (e.g., cache accesses), they compose into a smooth curve, with

a long tail describing execution times which are increasingly unlikely. Dependences between events in deterministic architectures can have an abrupt impact on execution time, producing discontinuities in the possible execution times which are difficult to model with a parametric distribution.

The absolute maximum execution time produced by the PROARTIS analysis will be many times greater than a conventional WCET, as it will correspond with the case where all instructions take the longest possible time to execute (e.g., every cache access is a miss [23]). We expect instead to gain by tightening the gap between observed execution times and worst-case execution bounds using a probabilistic confidence limit. The result of static analysis of deterministic architectures produces a degree of pessimism, where unknown states must be considered to have their worst consequences on the timing of the system. The 'true WCET' lies somewhere in the range marked (b) in Figure 1, between the maximum observed execution time and the WCET bound produced by analysis.

In PROARTIS the consequences of these unknown states can be considered probabilistically, which allows us to reason about the WCET probabilistically. Techniques from Extreme Value Theory (EVT) are used to construct a worst-case probability distribution. We define worst-case bounds with stated confidence levels, which can be chosen to match the degree of uncertainty present in the rest of the system being analysed. This will allow a tighter WCET bound to be considered (c).

WCET estimates are computed by considering the execution time at which the cumulative probability (ΣP) exceeds the required level of confidence. These confidence levels are expressed in Figure 1 in terms of the probability of exceeding the WCET threshold *in any given run*, however this figure should be adjusted based on arrival frequency to determine the probability per hour, or per year as necessary.

1.3   Contribution

In this paper we illustrate our vision, which attacks the limits of current timing analysis techniques using our novel EVT-based probabilistic timing analysis method. We discuss the challenges to be faced at the different layers of the execution stack when adopting our vision, and leave for future work the specification of complete solutions for all of the identified challenges. We argue that our analysis can provide WCET estimations with well-defined confidence bounds for CRTES: this assertion requires that the execution time behaviour of the system has specific mathematical properties, in particular *independence and identical distribution*, which we obtain by introducing randomisation in the timing behaviour at the bottom of the execution stack.

Our main contributions can be summarised as follows:

- We present two approaches to Probabilistic Timing Analysis (PTA). The first, statically derives a-priori probabilities from a model of the system: we call it Static PTA, SPTA. The second, measurement-based, derives probabilities from end-to-end runs on the target hardware of the software under study: we call it MBPTA. From these runs we collect data about the timing behaviour of the application software when executing on our proposed platform with randomised timing behaviour.

- For each such approach, we sketch the mathematical foundations on which PTA is based. We further enumerate and begin to attack the challenges we encounter at hardware and software levels upward the execution stack to pursue our vision.

- We then show how SPTA reduces the amount of knowledge needed to achieve tight WCET estimations, and we discuss the benefits that this reduction brings. SPTA requires the reuse distances of every access, which in fact are much easier to obtain than the addresses of memory operations required for conventional static timing analysis. Tight WCET estimations made with SPTA have less dependence on complete knowledge than conventional methods, and lack of information in the analysis has less impact on the WCET estimations. In fact, WCET estimations provided by SPTA react much more smoothly to the lack of knowledge, with a gradual shift towards a maximum value as knowledge is reduced.

Notably, a probabilistic analysis model is in match with the aging and reliability behaviour of the hardware itself, which common wisdom accepts may fail with a given (though very low) probability. In fact, not only the computing system, but any other mechanical parts that form the embedded system, also have a distinct probability of failure. In that sense, timing failures can be considered just another type of failure that the system may experience. The objective of the probabilistic timing analysis is to provide WCET estimations that are safe enough for the application, and keep the overall failure rate of the system below the domain-specific threshold of acceptability.

Section 2 provides a background description of how current timing analysis methods work and their limitations. Section 3 explains our probabilistic timing analysis approach and in particular how it addresses the main limitations of current timing analysis techniques. Section 4 provides a comparison of conventional static timing analysis and static probabilistic timing analysis for the cache. Section 5 shows how the probabilistic timing analysis approach fits with future embedded systems design. Section 6 presents related work. Section 7 surveys our main lines of future work. Finally, Section 8 presents the main conclusions of this study.

2   Timing Analysis on Deterministic Platforms

Reliable timing analysis in the form of WCET analysis is a key stage in the design and verification process of real-time systems, and becomes paramount for safety critical systems. WCET estimates are needed in the development of hard real-time systems to perform schedulability analysis, to ascertain whether the periodic and sporadic tasks meet their timing requirements, that the latency in the handling of interrupts falls within limits, and that aperiodic tasks are allowed sufficient time to perform useful work within sufficiently short service windows.

Unfortunately, timing analysis is a complex process, as the variations in execution time experienced by programs may be caused by the characteristics of the software itself, as well as by the hardware platform upon which the program is to run. It therefore follows that all salient characteristics of software and hardware must be thoroughly understood in order to provide meaningful WCET estimations. On the software side, execution time variations may arise from multiple sources, such as varying input sets a program can be asked to operate on, or its layout in memory for both code and data. Similarly, on the hardware side, all features that enhance average performance by exploiting properties of execution history, such as caches, pipelines and speculation, are potential sources of variation on the execution time.

In the following section we review current state-of-the-art techniques to perform WCET analysis and identify one of the main sources that threatens to make the cost of acquiring the information needed to perform timing analysis unaffordable: *the dependence of hardware and software timing behaviour on the history of execution*.

2.1 A deterministic approach to timing analysis

Conventional static timing analysis methods require the underlying system to be Time Deterministic (TD) or Time Predictable (TP). In a TD system, for a given input set, the sequence of events which will occur in the system is fully known, and so is the time after which the output of an event will be produced. This form of analysis obviously needs a very accurate and detailed model of the system, fully consistent with the causal dependence in place between all significant events. The safety and tightness of static timing analysis depend directly on the accuracy of the available system model, which in turn depends on the accuracy of information we can obtain about its actual operation and timing. Conventional analysis can also be applied to TP systems, in which the timing of the events that can occur during execution is not known in advance, but can be safely bounded. Hence, when the precise timing of an event cannot be determined, pessimistic assumptions can and must be made about its occurrence. This is, for example, the case with cache accesses, which must be considered misses when the memory address issued is not known beforehand, and hence the time of their occurrence cannot be predetermined. Notably, this assumption only holds in the absence of timing anomalies [19][26], whose presence significantly exacerbates the problem.

Unfortunately, with the increase in complexity of next-generation CRTES – at both hardware and software level – the extent of knowledge needed, as well as the time, effort, cost and complexity required to acquire and comprehend all the relevant details, become unaffordable. That is, the large amount of state that is carried by a modern processor leads to combinatorial explosion when trying to enumerate all possible *execution histories* even for simple pieces of code. Thus, as long as current analysis techniques are unable to scale up to the challenge, increased hardware complexity will cause a significant degradation of the quality of the resulting products.

In summary, industry demands new functionality and higher levels of performance together with reduced cost, weight and power dissipation, which can only be delivered by advanced hardware features. However, the timing behaviour of systems using these advanced hardware features is very hard to deal with by current timing analysis techniques as the amount of information required is becoming unaffordable.

Next we discuss an example of a hardware component, the cache, whose tight and safe timing analysis requires an especially large amount of information.

2.2 WCET Dependence on Execution History

Exploiting the execution history of the program, in the form of either temporal or spatial locality, is one of the most common principles of processor design. A typifying example of this strategy is the cache. The use of caches is widespread in general-purpose processors because they can dramatically improve the average application performance by exploiting locality in memory access patterns, reducing access times by up to two orders of magnitude. However, this strategy has an important downside: the execution time of programs, and their WCET, heavily depend on execution history, which introduces serious complexity, as we discuss below. For caches this dependence shows the level of the cache hierarchy the required data is stored in, and where.

With conventional static analysis approaches, knowledge of the execution history (in addition to knowledge of the target hardware) is required to provide tight WCET estimations. By maintaining the cache state for any point in the program, the analysis can precisely determine whether a memory access will be a hit or a miss. Many research efforts have attempted to obtain predictable cache behaviour by applying cache

locking (present in PowerPC 440, MPC5554, ARM 940) or deterministic replacement policies such as Least Recently Used (LRU) [Rgbw07].

However, to model all possible cache states is extremely costly, as it requires knowledge of all the memory accesses that a program makes. Not having the complete list forces the analysis to make pessimistic assumptions. For example, in a direct-map cache, if a single memory address is not known, the analysis must regard the cache as completely empty, as all cache lines could have potentially been evicted. Memory access times are even harder to predict in the face of complex, multi-level cache hierarchies.

Moreover, deterministic behaviour may result in pathological cases that are extremely difficult to predict or detect by testing. A recent study conducted for the European Space Agency [Bce07] shows the difficulties of using caches in CRTES: small program changes that lead to different memory layouts can trigger pathological cache behaviour which were called cache risk patterns: they are systematic cache misses that lead to large increases in WCET estimation. This may introduce *abrupt changes in the WCET* when the execution conditions when the program is deployed are not exactly the same as when it was analysed.

The difficulty of predicting memory access times in the presence of caches (even though the program may never trigger pathological behaviours) can lead to WCET estimates that are extremely pessimistic, where each unpredictable access is assumed to be a cache miss. This is one of the key sources of pessimism in static WCET analysis.

The level of complexity and the pessimism introduced by cache analysis has led many CRTES to dispense with caches altogether. Hence, despite the fact that the first cache memories appeared in the late sixties, they are still seldom used in CRTES.

3   PROARTIS: a New Approach to Timing Analysis

The central hypothesis of PROARTIS [22] is that new advanced hardware features can be used and analysed effectively in embedded real-time systems with designs that provide time-randomised behaviour. The introduction of randomisation in the timing behaviour of *hard-to-analyse* resources (e.g. caches), has a two-fold objective: (1) it reduces the cost of acquiring the knowledge required to perform trustworthy analysis; (2) it gives the system the properties needed for probabilistic timing analysis by reducing the dependence on execution history. This property, which is discussed in detail in Section 3.1, enables the use of probabilistic analysis techniques in verification arguments for CRTES, proving that the probability of *pathological execution times* is negligible, so that tighter bounds can be safely obtained. An element of the PROARTIS objectives is to develop a probabilistic timing analysis that will prove that *pathological cases* can only arise with quantifiably negligible probability, rather than trying to eliminate them (which is arguably not possible and could be detrimental to performance). This analysis approach will provide a substantial advance over current methods, including conventional static analysis techniques and testing, which are unlikely to scale up to the size and complexity of next-generation CRTES, the advent of which can hardly be stopped.

The techniques developed in PROARTIS will enable probabilistic guarantees of timing correctness to be derived. For example, if the requirements placed on the reliability of a system indicate that the arrival rate of a timing failure must be less than 10^{-9} per hour of operation, then our analysis techniques translate this reliability requirement into a probabilistic worst-case execution time for the system. Probabilistic

analysis provides a continuum of WCETs for different confidence levels. Hence, a system may have a probability of 10^{-9} of exceeding an execution time of $1.5\ ms$ over a one hour period of operation, and arrival rates of 10^{-14}, and 10^{-18} per hour of execution times exceeding $1.6\ ms$ and $1.7\ ms$, respectively. The aim of PROARTIS is that for future CRTES (including the hard real-time systems among them), probabilistic guarantees will offer significant advantages over deterministic approaches which attempt to make absolute guarantees, severely limiting the use of advanced hardware features and inevitably offering significantly lower performance guarantees.

Currently, the overall system reliability is expressed in terms of probabilities for hardware failures, memory failures, software faults and for the system as a whole. PROARTIS extends this notion to timing correctness. We aim to obtain WCET estimations for arbitrarily low probabilities, e.g. in the region of 10^{-50} per time unit of operation: to appreciate how small that probability is, consider that Extinction Level Events, such as an asteroid hitting the Earth, are estimated to happen about once every 100 million years, hence at an arrival rate of 10^{-16} per second, or 10^{-12} per hour.

3.1 A probabilistic approach to timing analysis

In general, a probabilistic analysis estimates the chance of future events based on an a-priori model of the probability. For instance, when throwing a die, assuming that the die is not loaded, a probabilistic analysis of the 'die model' would assume that the probability of it falling on any of its faces is $1/6$.

Statistical analysis searches for a model or properties when studying some (often large) body of data about observed past events. For instance, on the same example as above and assuming each face has a unique number, in order for the statistician to check if the die is loaded, he or she would define the hypothesis that 'the probability of each number appearing is the same'. By analysing a large number of throws, this hypothesis will be accepted or rejected, with a level of confidence based on how closely the observations match the model. In order to avoid confusion, in this paper we will not use the word 'stochastic' which is often associated with non-deterministic behaviour. Although the term is usually described to mean 'based on the theory of probability', it is often used in other contexts with subtly different meanings. Instead, we refer to probabilistic approaches outlining the nature of our idea which goes beyond the indeterminacy: we consider probability distributions to model the knowledge of processes.

Both probabilistic and statistical analyses associate an event A with a random variable[1] $\mathcal{A}$ describing the probability that such an event occurs. For instance if A is the event of throwing a die then $\mathcal{A}$ has the probability function (PF) $f_{\mathcal{A}}(\cdot)$ with $f_{\mathcal{A}}(a) = P(\mathcal{A} = a) = \frac{1}{6}$ being the probability to obtain the value $a \in \{1, 2, \cdots, 6\}$ after throwing the die.

The existing probabilistic and statistical analyses for CRTES consider usually either closed form (usually continuous) to describe the distribution of the random variables [16], or empirical distributions[2] [21],[9]. Within PROARTIS we consider the second option which is more appropriate for measurement-based approaches where the execution traces provide (by grouping the set of values) empirical distributions.

Moreover the existing probabilistic and statistical analyses for CRTES indicate that *they require precise hypotheses about the random variables* [13] to be captured in the

---

[1] In this paper we use a calligraphic typeface to denote random variables. Without loss of generality we consider discrete random variables.

[2] Some authors use the term "density histogram" to design the empirical distributions

analysis model. Among many hypotheses, we identify two of them: *independence* and *identical distribution* of the random variables. We introduce below the definitions of these notions, which we apply in our probabilistic timing analysis.

Definition 1 (Independent Random Variables). *Two random variables $\mathcal{X}$ and $\mathcal{Y}$ are independent if they describe two events such that the occurrence of one event does not have any impact on the occurrence of the other event.*

Definition 1 is equivalent to saying that $\mathcal{X}$ and $\mathcal{Y}$ are independent if and only if:

$$P(\mathcal{X} \leq x \text{ and } \mathcal{Y} \leq y) = P(\mathcal{X} \leq x)P(\mathcal{Y} \leq y) \tag{1}$$

Definition 2 (Identically Distributed Variables). *Two random variables $\mathcal{X}$ and $\mathcal{Y}$ are identically distributed if $\mathcal{X}$ and $\mathcal{Y}$ have the same probability distribution.*

Definition 2 is equivalent to saying that two random variables $\mathcal{X}$ and $\mathcal{Y}$ defined on the same space S^3 are identically distributed if for any subset $A \subseteq S$:

$$P(\mathcal{X} \in A) = P(\mathcal{Y} \in A) \tag{2}$$

We note that Definition 2 does not imply that two random variables $\mathcal{X}$ and $\mathcal{Y}$, which are identically distributed, are equal.

Definition 3 (Independent And Identically Distributed Variables). *A sequence of random variables $(\mathcal{X}_n)_n$ is independent and identically distributed (i.i.d.) when any two random variables $\mathcal{X}_i$ and $\mathcal{X}_j$, belonging to the sequence, are independent, $i \neq j$, and have the same distribution function.*

A sequence of (fair) die rolls where each roll is a random variable is i.i.d. as the random variables describe the same event type and the outcomes are obtained from mutually independent events.

Static[4] *Probabilistic Timing Analysis*: in SPTA, execution time probability distributions for individual operations, or components, are determined 'statically' from a model of the processor or software. The design principles of PROARTIS will ensure that it makes sense to derive execution time distributions for these operations, and to make the assumption that the probabilities for the execution time of each instruction are independent. For example, this means that regardless of whether an instruction is actually a cache hit or a miss when executed, the probabilities for later instructions remain the same.

SPTA is performed by calculating the convolution ($\oplus$) of the discrete probability distributions which describe the execution time for each instruction on a CPU; this provides a probability distribution, or Execution Time Profile (ETP), representing the timing behaviour of the entire sequence of instructions (see Figure 2).

We have used here a very basic model of the system: the CPU will execute each instruction in a fixed number of cycles, and the only source of timing variation comes from the cache. More complicated modeling is possible, generating a greater number of possible timings for each instruction.

---

[3]In our case we suppose S allows the definition of an order relation

[4]We apply the term Static to all analyses that derive the a-priori probabilities from an abstraction of the system.

Figure 2: SPTA uses a convolution to compute a probability distribution for the possible execution times of a sequence of instructions

Applying this technique for all instructions in a trace yields a new execution time distribution for the entire sequence of instructions, which includes all possible execution times for that sequence. By choosing a cut-off probability (e.g. 10^{-9}) and measuring where the inverse cumulative distribution function (the exceedance function) falls below this level, we can read off a WCET with a given confidence that we will not exceed that value.

The convolution of two ETPs requires only one assumption to be made about the input data: that the probability that an instruction causes a cache miss must be *independent* of the sequence of hits and misses which have occurred due to previous instructions. Note that this does not require that the probability distribution for an instruction is independent of the sequence of preceding instructions, only that the actual outcome of an event (a cache hit or miss) will not affect the calculated probability distribution for following instructions. Dependences between instructions may be modelled by creating profiles for short sequences of instructions which represent all the possible timing interactions between them. We aim to reduce the number of such dependences so that the complexity of combining instruction sequences is reduced.

We now consider a theoretical implementation of a cache that fulfills this requirement, which we refer to as a time-randomised cache. It is fully associative and uses an evict-on-access random replacement policy.

The major benefit to be had from this type of cache over conventional ones is that pathological eviction patterns can be avoided by making the probability of any given eviction pattern extremely small, and independent of the sequence of instructions executed. It also reduces the amount of information needed for modeling the cache behaviour to only the number of memory accesses, not their locations.

In our time-randomised cache configuration, the probability of evicting a given line on every access is $\frac{1}{N}$ where N stands for the number of cache entries. The formula used to compute the hit probability of a specific access to such a cache is as shown in equation 3 [30]:

$$P(hit) = \left(\frac{N-1}{N} \right)^{K} \tag{3}$$

In [30], K in Equation 3 corresponds to the number of cache misses between two consecutive accesses to the same cache entry. For the purposes of our WCET analysis, we must assume that every cache access could cause an eviction. Therefore, we define K to be the number of memory accesses between two consecutive accesses to the same cache entry, including the access for which we are computing K. K is referred to in this paper as the 'reuse distance' of an address. The value of K must be computed separately for each memory access, as the reuse distance depends on the address.

Under this cache configuration, the only information that the analysis requires for each memory access is the reuse distance and not the full sequence and addresses of

Address	Reuse Distance (K)	Hit Probability (N=32)
A	∞	0.00
B	∞	0.00
C	∞	0.00
D	∞	0.00
A	4	0.88
B	4	0.88
C	4	0.88
A	3	0.91
B	3	0.91
C	3	0.91

Figure 3: Example of sequence of memory operations and their reuse distance and probability of hit for a cache of 32 entries, i.e., N=32.


Figure 4: Example of SPTA Exceedance Function for simple 10 instruction trace.

the previous memory accesses as required when analysing conventional cache designs such as LRU. Let's assume the sequence of accesses A B C D A B C A B C (see Figure 3). We assume that the the cache is initially empty for this example so the K value for the first reference of each address is ∞ and thus its probability of hit is 0 from Equation (3).

To compute the distribution function for the instruction sequence in Figure 3, the hit probability for each access is converted to a two-value distribution, containing the hit and miss probabilities with their associated latencies.

For example, the hit probability for the second access to A in the example (0.88) is translated to a distribution containing an execution time of 1 with probability 0.88, and 100 with probability 0.12. Accesses with zero hit probability create single valued distributions with an execution time of 100 and probability 1. These distributions are then convolved, and the result can be seen in Figure 4. The chart shows an exceedance function, which is plotted as 1-cumulative probability over the range of execution times. It allows us to read off the probability that the WCET will exceed any given execution time value, or conversely to find a WCET to a given confidence level.

Notice that the maximum execution time is 1000 cycles: this represents the case where each access is a cache miss.

Measurement-Based[5] Probabilistic Timing Analysis: in MBPTA, complete runs of the program under study are made on the target hardware. From these runs we collect data about the timing behaviour of the program when run on low-level software and hardware elements of the PROARTIS platform with randomised timing behaviour. This information is used to determine the timing profile (as an execution time frequency distribution) of individual elements of the system, and then use these as inputs to the timing analysis of the overall system.

For MBPTA the maximum precision for the probability assigned to a given WCET value is $\frac{1}{number\ of\ runs}$. Hence, one would have to run an impossibly large number of tests to be confident of having observed the longest possible execution time. For instance, an execution time with a 1 in 10^9 probability of occurring will need, on average, one billion test runs to be observed. Theories must be developed so that the required number of experiments (runs) on the target platform is significantly decreased.

Our measurement-based timing analysis method is based on probabilistic techniques, and provides an estimation of the WCET of a task or application running as part of a larger system. Using randomisation in a system to defeat dependence on execution history is based on the observation that incurring the worst case is a rare event. Therefore any accurate probabilistic timing analysis needs to use the theory of rare events. This theory focuses on the analysis of tails of probability distribution, and it is classically used to study how random processes deviate from their expected value. More precisely the theory of rare events estimates for an event $\{\mathcal{X} \geq t\}$ the values of t for which the associated probability is very low (e.g. smaller than 10^{-4}).

Two rare events theories fit the WCET estimation problem: the theory of extreme values [12] and the theory of large deviations [17]. Extreme value theory (EVT) provides an estimate for the maximum of a sequence of i.i.d. random variables $(\mathcal{X}_i)_n$. Large Deviation Theory (LDT) instead studies how rare events occur and it provides an estimate for the sum of a sequence of i.i.d. random variables $(\mathcal{X}_i)_n$. LDT relies on the acquired concept that *rare events occur in the most likely way*.

To the best of our knowledge EVT is the only rare event theory previously applied to the WCET estimation problem. We provide here an analytical methodology to avoid the limitations of existing work on EVT (as described in [13]).

Independent and identically distributed variables: since the analysis models the behaviour of one and the same system in the same execution context, the variables are naturally identically distributed. Thus we concentrate our presentation on ensuring that the hypothesis of independence is valid. We identify two means to provide such independence:

- statistical methods: the sequence of dependent variables is transformed into an i.i.d. sequence using sampling techniques ([29]). The obtained i.i.d. sequence is not guaranteed to have the same properties as the original sequence.

- hardware design: the architecture guarantees the minimisation of dependences among traces. The PROARTIS solution is based on this means.

Continuous functions: the common utilisation of a discrete model within CRTES introduces unsafeness if EVT is applied directly on (set of) discrete values. We identify two methods to circumvent the problem:

---

[5]We apply the term Measurement Based to all analyses that derive a-priori probabilities from a measurement based model.

- Block maxima: the values are grouped and EVT is applied only to the maximum of each group. This method, which [14] applied to CRTES, helps EVT provide WCET bounds instead of mere execution time estimation (as it was the case of the EVT-based solution proposed in [10]). Coupling the grouping method to a static probabilistic analysis improves the quality of group size.

- Exceedance threshold: only values that are larger than a given threshold are retained. With respect to the block maxima method, this method compares the values to a given value instead of the maximum value of a group. The difficulty with this method stems from the determination of the threshold.

In addition to the limitations described in [13] we identify the over-estimation incurred by EVT in case of insufficient number of runs. Hence the knowledge of a minimum number of runs for a program is required. This number could be obtained by calibrating the results obtained with measurement-based approaches to the results of the static probabilistic timing analysis: when the former match the latter then the number of measurement runs is sufficient. It follows that measurement-based probabilistic analysis should be built on top of static probabilistic analysis, which thus is the root of the solution. In this paper, we therefore focus on static probabilistic timing analysis and leave the formulation of a measurement-based probabilistic analysis approach as future work.

3.2 Challenges at the platform level

The 'angle of attack' taken by PROARTIS is to develop an execution platform (hardware and software up to and including the middleware) whose timing behaviour has as little dependence as possible on execution history and that enables probabilistic timing analysis, without sacrificing average performance.

3.2.1 Hardware level

At hardware level our approach consists in time-randomising the behaviour of hardware components. In adopting this approach several key questions have to be answered. Do all hardware resources have to be randomised? Can a processor architecture have deterministic and time-randomised resources at the same time? To answer these questions we classify hardware resources into four types. The main characteristic of resources in terms of timing is whether their latency is fixed or not, whether it impacts WCET noticeably and whether it can be easily predicted.

- Many resources have a fixed latency. Estimating their impact on timing is trivial in our processor model. For instance, the latency to read a register or to perform an addition is typically fixed, so estimating the latency of an instruction simply requires knowing whether it uses such resources (which is implicit in the instruction itself) or whether all instructions must incur a time penalty as if they were using such resource (which is fixed by the architecture).

- Some specific instructions may have a variable latency when accessing a given resource. However, the difference between the worst-case and the best-case latencies is rather low. In that case (low variability) the simplest solution would consist of always assuming the worst-case latency to simplify the analysis. The impact in the WCET is negligible. For instance, the latency of a division may

vary by few cycles depending on the value of the input registers to operate and divisions are scarce. Thus, assuming the worst-case latency has very modest impact on the determination of the WCET.

- There are some resources whose latency varies significantly. Hence, simply assuming the worst-case latency is not convenient because the WCET may become extremely pessimistic. However, if accurately predicting the latency of those instructions can be afforded in terms of the effort/cost required for getting the information to have good predictions (highly predictable events), we can make pessimistic assumptions in those cases where full guarantee on their latency cannot be had. The impact on the WCET determination from that pessimism will be rather low.

- Finally, some resources may have significantly variable latency and there is no affordable way to predict it (hard to predict events). Those are the main candidates to be re-designed so that their latency can be characterised probabilistically.

PROARTIS focuses on the latter type of resources, whose impact in timing is high and unpredictable or, if predictable, too costly to reason about. The most prominent examples of this type of resources are cache-like resources, which we analyse in more detail next. The first three types of resources in the above taxonomy can be analysed with static analysis techniques as well as with PROARTIS techniques.

Core design. No core resources have highly variable latency. Two factors influence possible probabilistically analysable designs: timing anomalies and performance. We disallow the use of any hardware component that may introduce timing anomalies [19] [2]. From a hardware-only perspective and according to [2], no timing anomalies can occur in processors that do not allow allocating resources to instructions in a way that could delay the completion of older instructions. The downside of this argument is that we cannot use some performance acceleration features because they break this restriction. However, as shown in the MERASA project [20], the memory and not the core is the performance bottleneck in current CRTES. An in-order core processor design with in-order execution of the instructions and allocation of resources at issue time makes the execution time of an instruction not to depend at all on subsequent instructions. Those processors avoid by design the occurrence of timing anomalies [2]. By the same reasoning, more complex processor pipelines can in principle be used without any problem in our context. Branch prediction, and in particular the instructions executed on a misprediction may affect the state of the processor and hence, WCET estimations. As for the case of static analysis we can simply prevent speculative instructions to change the processor state so that the effect on WCET is bound. Overall, if a pipeline can be analysed with pipeline analysis techniques for Static Timing Analysis, nothing prevents its use with a probabilistic time analysis tool.

Cache design. Caches are among the most difficult resources to analyse for timing, as their timing behaviour, which is inherently jittery, greatly depends on previous execution history. Two main sources of dependence can be singled out: the placement policy, which selects the set to which a memory address maps, and the replacement policy, which selects the cache line to be evicted from a cache set.

- Fully-associative caches with a random replacement policy effectively reduce such dependence by making each cache line eviction independent of previous cache accesses. By doing so, we can probabilistically characterise the outcome (hit or miss) of memory accesses, simply based on the reuse distance (see section 3.1).

- Fully-associative caches are complex to implement and energy-consuming. However, when using a set-associative (with random replacement) or a direct-map cache, the placement policy is deterministic. Deterministic placement policies create dependencies and lead to cache risk patterns and hence to time pathological cases. We devise designs based on set-associative or direct-mapped caches using random replacement algorithms, which exhibit a similar behaviour at a lower cost. At hardware level, we devise a solution where we can randomise the placement function at program start up. To that end we XOR the address accessing the cache with a randomly generated number and use the result as the new address to access the cache. Under every placement setup, the address colliding in the different sets are different. Hence, different runs will lead to different (random) colliding addresses in each set and hence different (random) conflicts and execution times. Alternatively, the compiler can provide software-level support to randomise cache placement as we explain in the next subsection.

Communication elements design. The communication elements between hardware components, such as buses or interconnection networks, can also be a source of dependence, especially those which use access policies based on execution history such as Round-Robin or Least-Recently Accessed. We devise using random access-granting policies to remove such dependencies (e.g., lottery bus [15]).

3.2.2 Software level

Random-replacement caches with deterministic placement policies can be complemented with compiler support, as the timing behaviour of the placement policy depends directly on how objects are mapped in memory, i.e. on the memory layout. Thus, compiler techniques to randomise the allocation of objects in memory, such as the code, the stack or the heap, can be applied [3]. By doing so, the effective memory addresses of objects and so the cache set to which they are mapped will periodically change, thereby "randomising" cache placement and consequently enabling probabilistic analysis for cache memories with deterministic placement.

At the run-time level, it is known that not all operating system calls have constant timing behaviour; some system calls may in fact exhibit timing behaviour that depends on the invocation time of the call in relation to the internal operations of the operating system, and on which applications have executed and when and for how long, since the previous system call. An operating system with constant-time services also capable of not causing execution history dependence in the subsequent run time of the application would warrant time composability between the operating system itself and the application. Working towards this goal is also part of our future work.

4 An illustrative Example

In this section we discuss results obtained from the application of Static Probabilistic Timing Analysis. The development of techniques based on measurement-based probabilistic timing analysis is left for future work.

4.1 Experimental Set-up

For the experiments shown in this section we focused on the data cache and assumed a perfect instruction cache (so that all its accesses were hits). The analysis we conducted

```
:LOOP_START        load @1
                   load @2
                   ...
                   load @100
                   iter = iter + 1
                   compare iter, 100
                   jump LOOP_START if smaller
```

Figure 5: Synthetic benchmark used in this section. The loop performs 100 iterations.

can equally be applied to the instruction cache removing the all-hits hypothesis. All core operations have a fixed latency of 1 cycle. Latencies for the data cache are 1 cycle in case of hit and 100 cycles in case of miss.

To study the potential of probabilistic timing analysis, we used a synthetic benchmark consisting of a loop with 100 distinct data memory accesses (see Figure 5). Apart from memory operations, we included 3 core operations with a fixed latency of 1 cycle for the loop control. The only relevant event in this experiment is the latency of data accesses, as the fixed latency control operations simply add 3 cycles per iteration.

Results are shown for 100 consecutive iterations after executing the loop at least once. This way we avoid compulsory misses that would simply add 10,000 cycles (100 misses with a 100-cycle latency each) to the WCET estimation regardless of the cache and analysis used.

For the conventional static timing analysis we consider 2-way, 4-way and 8-way LRU replacement caches. For the probabilistic timing analysis we consider a fully-associative cache in which elements can be allocated randomly in any position. In our particular example, $K = 100$ for all accesses.

While time-deterministic caches (e.g. those based on LRU) produce a single WCET, time-probabilistic caches produce a distribution of execution times. In all experiments we use the execution time whose exceedance probability is at most 10^{-9} for the WCET, unless otherwise stated.

We assume that addresses do not conflict in the conventional LRU-based caches. So all accesses are hits in each LRU cache configuration if data fit in cache.

## 4.2 Static Probabilistic Timing Analysis Results

With SPTA, the distribution function is computed by convolving the discrete probability distribution calculated for each instruction. These distributions describe the probability that each instruction will take a given execution time. In our example, the distributions for memory operations contain only two values: the execution time when the instruction hits an existing cache entry (1 cycle), with probability $p(hit)$, and the execution time when the instruction causes a cache miss (100 cycles), with probability $1 - p(hit)$. The distribution for core operations will contain a single value of 1 cycle with a probability of 1.

Figure 4 shows an exceedance plot of the distribution function for a very short, 10 access sequence. With longer sequences of instructions, we can obtain WCET estimations for arbitrarily low probabilities (see Figure 6). Probabilities can be in the region of e.g. 10^{-50} or lower.

The horizontal 'shelf' which can be seen in Figure 6 corresponds to the level of precision used in the calculation of the probabilities: for probabilities lower than those representable by the level of precision selected, the value is not representable and results in the shown artefact . The 'real' distribution continues to follow the shape of
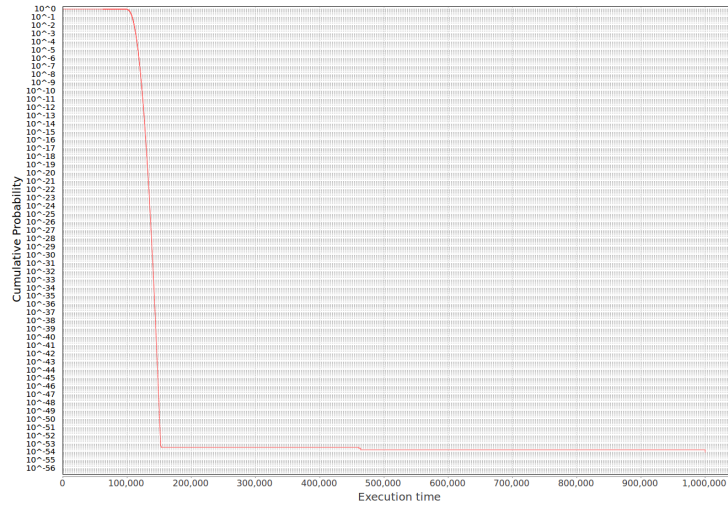
(a) Overall execution time distribution
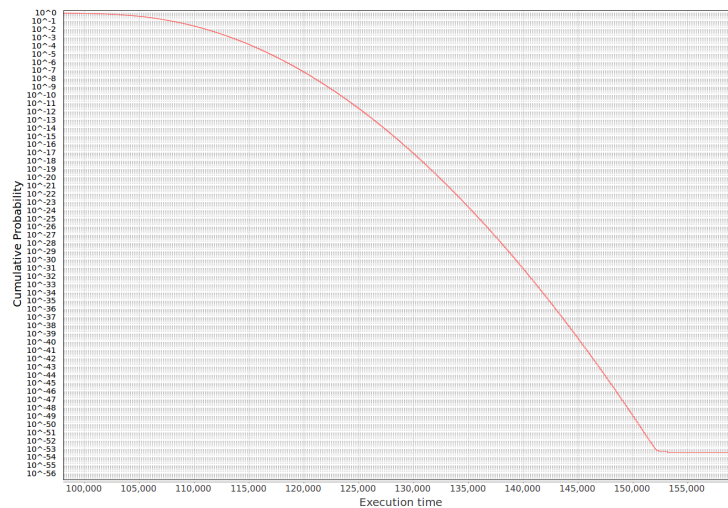
(b) Zoomed region of the function focussing on the area of the large drop in probability.

Figure 6: WCET exceedance function plotted on a logarithmic scale. Results are for the synthetic benchmark described in this paper with N=1000 and no unknown addresses.

a normal distribution, which puts the probability of the maximum value (the far right of Figure 6(a), 1,000,000 cycles) at an extremely small probability: 100 iterations of a 100 memory access loop in which every memory access misses in a cache of 1,000 entries (N=1,000) has a probability of $1 \times 10^{-30,000}$.

4.3 Comparing Conventional Static Timing Analysis and Static Probabilistic Timing Analysis

We have shown that static probabilistic timing analysis requires less information than conventional static timing analysis. Moreover, the dependence between different parts

of an application can be reduced when using probabilistic techniques. In this section we provide examples of how the probabilistic and conventional analyses react as the amount of information available about execution history decreases.

Making a fair comparison between two timing analysis techniques is difficult. There are many factors that affect the accuracy of the provided WCET estimations in each case. As an illustrative example, in this section we focus on the timing analysis of the cache only. To that end, we propose an evaluation framework in which the cache is the only source of WCET estimation inaccuracy. In order to prevent the processor pipeline affecting timing variability, which can be significant if the processor suffers from timing anomalies or domino effects, we assume a simple in-order pipeline in which each instruction that is not a memory operation takes a fixed latency which is known a-priori. Analogously, in order to discount the effect of other phases of the analysis, such as path analysis or loop bound analysis, we take programs with one only possible path, i.e., the input data is fixed so that all loop bounds are fixed and any branches are chosen deterministically. In order to extend this approach to software with multiple paths, we intend initially to compute the SPTA distribution for each path under analysis, and then to combine them using the maximum envelope of the inverse cumulative distribution functions. This can be shown to be a pessimistic upper bound on the execution time of the system, regardless of any relative weighting which could be given to different paths in the system.

Addressing the feasibility of analysing every path in a program, the sequences of instructions which make up a basic block or a single path through a function could be considered a smaller-scale version of our single path problem. By using a tree-based approach to combine the blocks using the envelope method at branch locations and convolutions for sequential segments, it should be possible to construct an SPTA profile without the need to evaluate all paths in the software separately. This is part of our future work.

- Conventional Static Timing Analysis of the cache distinguishes between certain cache hits/misses and unclassified accesses. An exact static data cache analysis requires knowledge of the target addresses of all memory access operations, which is often not possible to calculate statically. As a reference configuration in our experimental environment, we assume that the cache is the only source of inaccuracy for conventional WCET estimations. We further assume that all addresses are known, so in this configuration we have the tightest WCET estimation possible, which is, in fact, the actual WCET of the program.

- For the Static Probabilistic Timing Analysis of the cache we assume that the reuse distance is known. Note that knowing the reuse distance requires much less information about the memory operations than knowing the exact sequence of memory operations, which is the information required for tight conventional WCET estimations. For example, we may not know until run time the address of a variable on the stack, or the target of a function pointer, but we *can* compute the reuse distance for accesses to such a variable, and similarly for pointers if the pointer itself is not modified between the accesses.

Reducing the amount of required information   In addition to the idealised case in which all the information required by the analysis is available, we show how different WCET techniques react when part of this information is missing. In particular, for the conventional analysis we assume that a given percentage of the accesses to memory

(a) initial state   (b) 1 unknown address processed   (c) 2 unknown addresses processed

Figure 7: Cache state kept by the conventional static analysis after several unknown addresses are processed

have an unknown address. For the conventional WCET analysis *and in our experimental set-up* this lack of information translates into assuming that unknown accesses are misses and that in the cache state kept by the conventional analysis, every unknown address that accesses the cache moves all the data one position towards the LRU position.

For example, Figure 7(a) shows a two-way 8-set cache with a given initial state. Each cell represents a line in the cache and in each cell we have a piece of data, represented by a letter. The subscript is the LRU-stack position. In this case, in each set a line can be either the Most-Recently used line (MRU) or the Least-Recently Used (LRU) line. With exact cache analysis and full information on the accesses, this would be the state of the cache analysis. Figure 7(b) shows the effect on the cache state of one unknown address. Even if the address will affect only one set, the particular set that is affected cannot be determined, meaning that the unknown access can hit any set. Hence, all lines in all sets are moved to the LRU position: the data in the LRU position is assumed to be evicted and the data in the MRU position in our 2-way cache moves to the LRU position. Figure 7(c) shows the cache state after two unknown addresses are processed. All data have been virtually evicted from the cache because we cannot guarantee any particular piece of data to be in cache after those two accesses with unknown addresses.

In the case of Static Probabilistic Timing Analysis, we assume that the reuse distance of some of the addresses is not known. To track reuse distance we only need to know the number of different accesses between two accesses to the same cache entry. This can be calculated without knowing the actual addresses.

One of the properties of this probabilistic model is that every unknown address only has effect on previous instances of the same address. Let's illustrate this concept with the same simple example used before. With our random cache the probability of evicting each line is $\left(\frac{N-1}{N}\right)^K$, where K is the reuse distance. Let us again assume the sequence of accesses $A\ B\ C\ D\ A\ B\ C\ A\ B\ C$ discussed in Section 3.1. If the second instance of A is unknown, the K value of the third instance of A increases from 3 to 7 (see Figure 8). Obviously, the K value for the access with an unknown address is ∞.

We observe how the lack of information for a particular access affects two accesses at most. So the effect of lack of information at analysis time about some addresses has a small effect on WCET.

Address	Reuse Distance (K)	Hit Probability (N=32)
A	∞	0.00
B	∞	0.00
C	∞	0.00
D	∞	0.00
?	∞	**0.00**
B	4	0.88
C	4	0.88
A	7	0.80
B	3	0.91
C	3	0.91

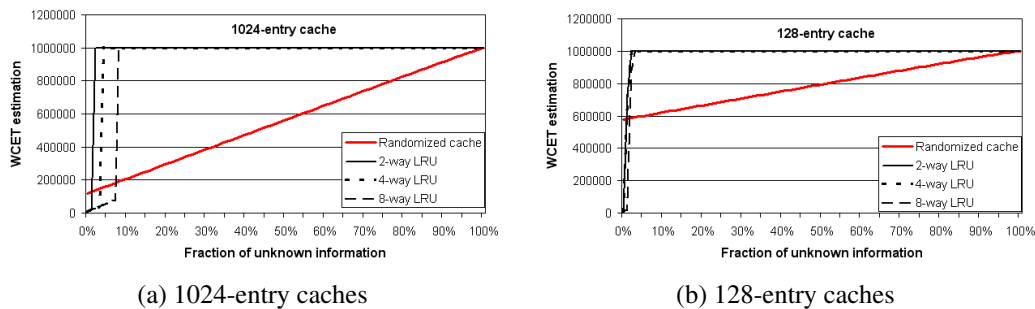Figure 8: Example of sequence of memory operations and their reuse distance and probability of hit for a cache of 32 entries, when the reuse distance of one access is missing

(a) 1024-entry caches

(b) 128-entry caches

Figure 9: WCET estimations for caches with different size, expressed in number of cache entries, when varying the fraction of unknown addresses.

4.4 Results

Effect of lack of information on WCET estimations   The objective of this experiment is to show the dependence of each type of analysis on the amount of available knowledge. To that end, in this experiment we show the WCET bounds obtained for different cache configurations with varying extents of unknown information.

In terms of WCET estimation, those accesses where the address is unknown are assumed to miss in cache and to produce the worst potential impact, as considered by the timing analysis methods. In particular, any element in the LRU position of each set cannot be assumed to be in cache after an unknown access. Conversely, an unknown access does not vary the K for all remaining accesses in the loop for the randomised caches.

Figures 9(a) and 9(b) show WCET bounds for the benchmark example for two different cache sizes: 1024 and 128 entries for different percentages of unknown information; that is, the number of accesses for which the address is unknown in the conventional analysis and the number of accesses where the reuse distance is unknown in the case of probabilistic analysis. When the number of entries is significantly larger than the working set ($N = 1024$) LRU-based caches provide the best WCET if all or almost all addresses are known. However, as the amount of uncertainty grows (which is the common case), the WCET for LRU-based caches becomes rapidly pessimistic. Conversely, the randomised cache still provides low WCET estimates, and degrades smoothly as the amount of uncertainty grows. This is a desirable property, since small

Figure 10: WCET estimations for a time-randomised cache and different fractions of unknown addresses when varying the cache size expressed in number of cache entries.

changes in the deployment conditions with respect to the analysis conditions do not cause abrupt changes in the WCET.

Figure 9(b) shows results for a cache where data occupies most of the space ($N = 128$). LRU-based caches show an abrupt increase in WCET when only a few addresses are unknown. With SPTA, caches experience higher replacement probabilities because of the lower N (see Equation 3), but the WCET increases smoothly and is significantly lower than that of LRU-based caches when at least 3% of the addresses are unknown. Although not shown, when the cache size is lower than the working set LRU-based caches are unable to provide any cache benefit, even if all addresses are known. Time-randomised caches still provide some hits because any piece of data has some probability to remain in cache after an arbitrary number of evictions.

Sensitivity analysis The second set of experiments studies the sensitivity of the different cache designs to the cache size for different levels of unknown information (0%, 20%, 40%, 60% and 80%). In Figure 9 we can observe that WCET estimations for LRU-based caches are good only if the working set fits in cache and all (or almost all, i.e. 10%) addresses are known. However, if some addresses are unknown or the working set does not fit in cache the WCET becomes very pessimistic. Conversely, the time-randomised cache provides low WCET estimates that degrade smoothly as we reduce the cache size or increase the fraction of unknown addresses, as shown in Figure 10.

Note that in our experimental set-up, we made the cache be the only source of WCET overestimation. Under this scenario, when the amount of unknown information is 0%, the WCET estimation matches the actual execution time. What we change in each experiment is the amount of information the analysis method knows, which is the source of pessimism of the WCET with respect to the actual execution time. In the particular case of no unknown information, the WCET estimation of Static Timing Analysis with LRU caches is 5-7x better than Static Probabilistic Timing Analysis with random caches. However, this scenario is unrealistic, as in most cases a significant fraction of the information needed for the timing analysis is unavailable. This is what we show in Figure 9 as we move right on the x axis. Some examples of lack of information are (i) matrices indexed with input-dependent variables, (ii) pointers, (iii) stack size and location for functions with input-dependent parameters, etc.

(a) 1024-entry cache

(b) 128-entry cache

Figure 11: WCET estimations for a time-randomised cache of varying size, expressed in number of cache entries, when varying the exceedance probability.

Sensitivity to exceedance probability   The third set of experiments studies the sensitivity of the time-randomised cache to the exceedance probability used as threshold to determine the WCET. Results are shown for 4 particular combinations of cache sizes (1024 and 128 entries) and unknown information fractions (0% and 20%), although trends are very similar for other scenarios. The exceedance probabilities we studied range between 10^{-3} and 10^{-30}. Figures 11(a) and 11(b) show results for 1024 and 128 entries caches respectively. Decreasing the exceedance probability to increase the safety level of the system does not have a significant impact in the WCET obtained. Decreasing the exceedance probability by 3 orders of magnitude requires increasing the WCET by 0.9% on average across configurations. Hence, the WCET estimations can be used in systems with the highest timing constraints.

5   Contextualising PROARTIS

Cost pressure, flexibility, extensibility and the demand for increased functional complexity are changing the fundamental paradigms for the definition of automotive and aeronautics architectures. In this section we motivate the applicability of the PROARTIS approach in the avionics domain. Similar reasoning can be applied to the automotive domain.

In the past, conventional avionics systems were based on the *federated architecture* paradigm, in which each computer is a fully dedicated unit, which may have local optimisations. However, since most of the federated computers perform essentially the same functions (input acquisitions, processing and computation, and output generation), a natural optimisation of resources is to share the development effort by identifying common subsystems, standardising interfaces and encapsulating services; in other words, adopting a modular approach. This is the purpose of the Integrated Modular Avionics (IMA) concept, where 'integrated' refers to the sharing of these resources by several systems. To that end, federated architectures have been replaced by *integrated architectures* in which the same computer can execute one or more applications, or 'partitions', potentially of different criticality levels.

One fundamental requirement of IMA is to ensure the possibility of *incremental qualification*, whereby each partition can be subject to verification and validation – including timing analysis – in isolation, independent of the other partitions, with obvious benefits for cost, time and effort. From the perspective of timing analysis, incremental qualification relies on each hardware and software component that is part of the system, such as the cache at hardware level or the linker at software level, exhibiting the

property of *time composability*. In the most general definition, composability ensures that, given a certain property of each item of a collection, that property can be determined for each item taken in isolation and does not change when that item is brought together with other items. Time composability refers to the fact that the execution time of a software partition, observed in isolation by the timing analysis, is not affected by the presence of other partitions on the same system.

The PROARTIS approach attacks the root of this problem: firstly by reducing, or even completely eliminating, the information required to calculate WCET estimations, which reduces the cost of acquiring the required knowledge to model the timing behaviour of the system; secondly, by reducing the dependence of the WCET estimations on execution history. In this way, software execution times are less dependent on previous executed software components, which can belong to other partitions. As a result the system integration can be easily achieved. In the extreme case, if all the execution history dependence is completely eliminated, each individual resource will be timing composable, allowing partitions to be replaced without requiring the timing of other components to be re-analysed.

6   Background

The first papers in the real-time community related to the contribution of this paper used the words 'stochastic analysis' [11], 'probabilistic analysis' [27], 'statistical analysis' [1] and 'real-time queuing theory' [16]. Since the paper of Diaz et al. [9], the 'stochastic analysis' of real-time systems has been used regularly by the community regardless of the approach (probabilistic or statistical). We note that the terminology used in the past is somewhat inconsistent; one of the objectives of the paper is to define precisely those terms so that they can be widely used.

The literature on probabilistic methods is vast. It is difficult to draw a cohesive view of the field. We rather describe here the precise application of these techniques to the timing analysis of real-time systems. Probabilistic approaches for real-time systems are promising because they answer questions that cannot be addressed in a deterministic manner (e.g., distribution of different task parameters) and consider more realistic models of task execution time or the expression of soft real-time constraints. In general, a probabilistic approach provides the distribution function of a parameter.

Probabilistic approaches for real-time systems can be divided into two main classes:

- One class consists of extracting quantitative information for one or more parameters (e.g., distribution of execution times) from samples of observations collected by monitoring the system. The first paper considering such statistical estimation of worst-case execution times is [10] in which extreme value statistics are used to model the behaviour of caches; more recently [14] improved the analysis, addressed flaws in the previous work and produced a refined method also based on extreme value statistics. This line of work is promising if one can show that hypotheses like independence (for instance) are met, which unfortunately is not the general case [13]. Such hypotheses can be met for instance by designing platform (hardware and software) support as we suggest in Section 3.2.

  The RapiTime tool [24] from RAPITA is the only timing analysis tool that captures execution time distributions, also called Execution Time Profiles (ETP), and implements a method based on the theory of copulas [5] to compute the dis-

tribution of execution times of the worst-case path. The results in [21, 6] are the only response time analysis that belongs to this class of probabilistic approaches.

- The second class of approaches concerns the temporal analysis of systems with at least one parameter described by a random variable. We present here the papers that consider the temporal analysis of systems with WCET described by random variables. In this case the probabilistic analysis specifically decides the schedulability of the system, i.e., answers the question *are the deadlines met, and with what probability, when the elementary execution times and thus the WCET are random variables?*. Therefore such analysis takes WCET estimations for each program in input (as provided as random variables) and returns the probability of missing the deadline for the set of programs. The first results on probabilistic schedulability analysis were proposed in [16] where the author extends queuing theory under real-time hypotheses. Even if this work was subsequently improved by results such as [31], the main limitation concerns the use of the same probability law for the execution times of all programs. This latter hypothesis is not always realistic, e.g., the traces of execution of two different programs could fit (after a WCET estimation) to different probability laws.

  Another relevant work [1] proposes an interesting definition of probabilistic deadline, but the authors consider a special scheduling model providing isolation between tasks. The results presented in [11, 27] consider also execution times as random variables and special assumptions are made on the critical instant.

  The main result of this class is proposed in [18], but the analysis is difficult to use in practice for computational reasons. An improvement based on re-sampling was proposed in [25] but this work fits better in the first class of approaches for the statistical analysis of real-time systems.

The main difficulty when using both approaches is the consideration of worst-case behaviours that are known to be rare events. In this case, Rare Events Theory needs to be used in order to ensure that the worst-case behaviour is considered. These approaches have the advantage of predicting unexpected behaviours, but they also have the disadvantage of offering incomplete information (the distribution function is not entirely defined). Except for [4, 5], all existing approaches consider that the parameters given by random variables are independent. This hypothesis is due to theoretical limitations, but also to the difficulty of describing and capturing the relations between different parameters like the relation between the execution times of two programs for instance.

Besides some known probabilistic results (such as the stability of Markov chains) the results used for the schedulability of real-time systems are mainly based on operations with random variables which extend deterministic calculation. Queuing theory is a notable exception, but it is usually classified as operations research.

As regards the three distributions of Extreme Values Theory (Gumbel, Frechet, and Weibull) [12], we are interested in using Gumbel's which provides information on the right-tail part of a distribution (thus the interesting part for a WCET estimation). A first step toward fitting the data to this particular distribution is needed and errors could be introduced during this step (thus different levels of confidence might be proposed). One solution to decrease the errors is proposed in [14] and it is based on block maxima (grouping different values within a distribution into the maximal one). Another rare events theory that interests us is Large Deviation Theory [17]. This theory also deals with tail events, and extends Central Limit Theory. Thus it has the feature that it can

only be applied to a sum of i.i.d. random variables. This property could be useful for instance if we wanted to reason on the combination of execution traces taken from different "parts" (e.g., paths, units) of a program.

7   Ongoing and Future Work

In this paper we have shown the feasibility and the requirements to properly implement a probabilistic timing analysis approach. Given the seminal nature of this work, we consider several lines of future work.

Firstly, providing more measurement-based and static probabilistic analysis techniques. A key aspect of this line, is that these techniques have to be designed hand-in-hand with the platform, which has to ensure that the hypotheses used by the technique, e.g., i.i.d, are demonstrated by the hardware and low-level software.

Secondly, at the hardware level our aim is to find designs which facilitate analysis in general and improve the behaviour of both static and measurement based probabilistic timing analysis. We want to extend our solutions to hardware resources that are hard-to-analyse with current analysis techniques, e.g. the Translation Lookaside Buffer (TLB).

Thirdly, at the software level our objective is to determine how software support for randomisation can be used in conjunction with the hardware techniques to further facilitate probabilistic timing analysis.

And finally, there is a shift towards multicore processors in the real-time market. The main impediment of the use of multicores in real-time systems is the difficulty of predicting how tasks will interact in shared hardware resources. We believe that probabilistic timing analysis is a promising solution to bound (in a probabilistic manner) inter-task interactions, so that we know the effect that one task may introduce in the WCET estimations for another.

8   Conclusions

The industrial demand for new functionality and higher levels of performance in addition to reduced cost, weight and power dissipation calls for the adoption of advanced hardware acceleration features, of which cache memories are the epitome. However, the timing behaviour of systems using these advanced hardware features is very hard to analyse with current timing analysis techniques: the level of knowledge needed to compute tight WCET estimations, as well as the time, effort, cost and complexity required to acquire and comprehend all down to the finest the system details become unaffordable. This is so because the large amount of state information that is carried by a modern processor incurs a combinatorial explosion when one tries to enumerate all possible histories of execution even for simple programs. Thus, as long as current analysis techniques are unable to scale up to the challenge, increased hardware complexity will lead to a significant degradation of the quality of the resulting products.

We call these problems the Time Analysis Walls. Probabilistic Timing Analysis attacks these walls and has potential for tearing them down. It is our contention that the cost of acquiring the required knowledge to perform trustworthy analysis can be significantly reduced by adopting a hardware/software architecture whose execution timing behaviour eradicates dependence on execution history. One way to achieve this independence is via introducing randomisation in the timing behaviour of the hardware

and software (while the functional behaviour is left unchanged), together with new probabilistic analysis techniques.

We have devised two different flavours of Probabilistic Timing Analysis. The first, SPTA, is static in the sense that it derives a-priori probabilities from a model of the processor or software. The second, MBPTA, derives probabilities from complete runs of the program under study that are made on the target hardware. From these runs we collect data about the timing behaviour of the program when run on low-level software and hardware elements of the PROARTIS platform with randomised timing behaviour.

In each case we have shown the main challenges and opportunities in implementing that approach. Our initial results for the cache using the SPTA approach show that (1) SPTA requires knowing the reuse distances of every access, which in fact are much easier to obtain than the addresses of memory operations required for Static Timing Analysis. (2) SPTA has minimal dependence on the amount of knowledge required to provide tight WCET estimations. Moreover, with SPTA, any lack of information in the analysis has modest negative effect on the WCET determination and the tightness of the resulting WCET estimations degrades smoothly with increasing lack of knowledge.

Finally, we have outlined the direction of our current and future work in the regard of the subject matter.

9 Acknowledgments

References

[1] L. Abeni and G. Buttazzo. Integrating multimedia applications in hard real-time systems. In *the 19th IEEE Real-Time Systems Symposium (RTSS98)*, pages 4–13, 1998.

[2] Peter P. Puschner. Albrecht Kadlec, Raimund Kirner. Avoiding timing anomalies using code transformations. In *ISORC*, 2010.

[3] Emery D. Berger and Benjamin G. Zorn. DieHard: Probabilistic memory safety for unsafe languages. In *In Proceedings of the ACM SIGPLAN 2006 Conference on Programming Language Design and Implementation*, pages 158–168. ACM Press, 2006.

[4] G. Bernat, A. Colin, and S.M. Petters. WCET analysis of probabilistic hard real-time systems. In *RTSS*, 2002.

[5] G. Bernat and M. Newby. Probabilistic WCET analysis, an approach using copulas. *Journal of Embedded Computing*, 2006.

[6] P. Binns. Statistical estimation of aperiodic response times when scheduled on top of static timelines. In *1st Int Workshop on Probabilistic Analysis Techniques for Real-Time and Embedded Systems(PARTES04)*, 2004.

[7] R.N. Charette. This car runs on code. In *IEEE Spectrum online*, 2009.

[8] P. Clarke. Automotive chip content growing fast, says gartner. In *http://www.eetimes.com/electronics-news/4207377/Automotive-chip-content-growing-fast*, 2011.

[9] J.L Díaz, D.F. Garcia, K. Kim, C.G. Lee, L.L. Bello, López J.M., and O. Mirabella. Stochastic analysis of periodic real-time systems. In *the 23rd IEEE Real-Time Systems Symposium (RTSS02)*, pages 289–300, 2002.

[10] S. Edgar and A. Burns. Statistical analysis of WCET for scheduling. In *the 22nd IEEE Real-Time Systems Symposium (RTSS01)*, pages 215–225, 2001.

[11] M.K. Gardner and J.W. Lui. Analyzing stochastic fixed-priority real-time systems. In *the 5th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS99)*, pages 44–58, 1999.

[12] B.V. Gnedenko. Sur la distribution limite du terme maximum d'une seris aleatoire. *Annals of Mathematics*, 44:423–453, 1943.

[13] David Griffin and Alan Burns. Realism in Statistical Analysis of Worst Case Execution Times. In *the 10th International Workshop on Worst-Case Execution Time Analysis (WCET 2010)*, pages 44–53, 2010.

[14] J. Hansen, S. Hissam, and G. A. Moreno. Statistical-based wcet estimation and validation. In *the 9th International Workshop on Worst-Case Execution Time (WCET) Analysis*, 2009.

[15] K. Lahiri, A. Raghunathan, and G. Lakshminarayana. LOTTERYBUS: a new high-performance communication architecture for system-on-chip designs. In *Proceedings of the 38th annual Design Automation Conference*, DAC '01, pages 15–20, 2001.

[16] J.P. Lehoczky. Real-time queueing theory. In *the 10th IEEE Real-Time Systems Symposium (RTSS96)*, pages 186–195, 1996.

[17] J. T. Lewis and R. Russel. An introduction to large deviations for teletraffic engineers, 1997. available at http://www.statslab.cam.ac.uk/%7Errw1/ld/LD-tutorial.ps.

[18] J.M. Lopez, J. L. Diaz, J. E., and D. Garcia. Stochastic analysis of real-time systems under preemptive priority-driven scheduling. *Real-time Systems*, 40(2):180–207, 2008.

[19] T. Lundqvist and P. Stenstrom. Timing anomalies in dynamically scheduled microprocessors. In *RTSS*, 1999.

[20] MerasaD2.2. `http://www.merasa.org/downloads/Deliverable_2_2.pdf`. 2008.

[21] N. Navet, L. Cucu, and R. Schott. Probabilistic estimation of response times through large deviations. In *WIP session of the 28th IEEE Real-Time Systems Symposium (RTSS'07)*, 2007.

[22] PROARTIS. Probabilistically analyzable real-time systems. feb 2010. http://www.proartis-project.eu/.

[23] Eduardo Quinones, Emery D. Berger, Guillem Bernat, and Francisco J. Cazorla. Using Randomized Caches in Probabilistic Real-Time Systems. In *22nd Euromicro Conference on Real-Time Systems (ECRTS)*, pages 129–138. IEEE, 2009.

[24] RapiTime. *www.rapitasystems.com*, 2008.

[25] Khaled S. Refaat and Pierre-Emmanuel Hladik. Efficient stochastic analysis of real-time systems via random sampling. pages 175–183, 2010.

[26] J. Reineke et al. A definition and classification of timing anomalies. In *WCET*, 2006.

[27] T.S. Tia, Z. Deng, M. Shankar, M. Storch, J. Sun, L.C. Wu, and J.S Liu. Probabilistic performance guarantee for real-time tasks with varying computation times. In *the 2nd IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS95)*, pages 164–174, 1995.

[28] Wilhelm R. et al. The worst-case execution-time problem overview of methods and survey of tools. *ACM Transactions on Embedded Computing Systems*, 7:1–53, May 2008.

[29] Lu Yue, Iain Bate, Thomas Nolte, and Liliana Cucu-Grosjean. A new way about using statistical analysis of worst-case execution times. September 2011.

[30] S. Zhou. An efficient simulation algorithm for cache of random replacement policy. In *Proceedings of the 2010 IFIP international conference on Network and parallel computing*, NPC'10, 2010.

[31] H. Zhu, J.P. Hansen, J.P. Lehoczky, and R. Rajkumar. Optimal partitioning for quantized EDF scheduling. pages 202 – 213, 2002.