







# Fitting Software Execution-Time Exceedance into a Residual Random Fault in ISO-26262

Irun Agirre , Francisco J. Cazorla , Jaume Abella , Carles Hernandez , Enrico Mezzetti ,  
Mikel Azkarate-askatsua, and Tullio Vardanega 

**Abstract**—Car manufacturers relentlessly replace or augment the functionality of mechanical subsystems with electronic components. Most such subsystems (e.g., steer-by-wire) are safety related, hence, subject to regulation. ISO-26262, the dominant standard for road vehicles, regards software faults as *systematic*, while differentiating hardware faults between *systematic* and *random*. The analysis of systematic faults entails rigorous processes and qualitative considerations. The increasing complexity of modern on-board computers, however, questions the very notion of treating the violation of execution-time envelopes for software programs as a systematic fault. Modern hardware in fact reduces the user's ability to delve deep enough into the fabric of hardware–software interaction to gage its extent of contribution to the worst-case execution time (WCET). Changing the nature of the WCET-analysis problem may help address that challenge effectively. To this end, we propose a solution that should allow ISO-26262 to quantify the likelihood of execution-time exceedance events, relating it to target failure metrics employed in support of certification arguments, similarly to random faults in hardware. To this end, we inject randomization in the timing behavior of the computer hardware to relieve the user from the need to control hard-to-reach low-level parts, and use measurement-based probabilistic timing analysis to quantify, constructively, the failure rates resulting from the likelihood of execution-time exceedance events.

**Index Terms**—Automotive real-time systems, execution-time exceedance, measurement-based probabilistic timing analysis (MBPTA), safety certification.

## I. INTRODUCTION AND MOTIVATION

**A**N INCREASING variety of functions in modern cars are controlled by electrical and/or electronic (E/E) subsystems; for instance, active/passive safety and driver assistance. For quantity, complexity, and use, those functions make the safety of E/E systems an increasingly important and complex matter.

ISO-26262 [28] is the functional safety standard of reference for the automotive domain. It is an adaptation of the broader IEC-61508 safety standard, which has been similarly adapted to nuclear plants, industrial machinery, railway, and other application domains (see Fig. 1).

ISO-26262 seeks to preserve systems' safety by sustaining *safety goals* (SGs) that prevent hazardous situations due to E/E malfunction. To this end, ISO-26262 (much like the parent IEC-61508) defines procedures for the management of deterministic design faults (i.e., *systematic* faults) and unpredictable hardware faults (i.e., *random* faults). The ISO-26262 tenet is that systematic faults can be either avoided by adopting prevention measures throughout the development process, or controlled at run time by safety mechanisms such as diverse redundancy. ISO-26262 uses cognizant assessment, based on judgment from practical experience, to guarantee that the contribution of systematic faults to SGs violation is kept acceptably low by assuring coverage of all requirements of the standard. Conversely, random faults can only be controlled at run time: ISO-26262 requires their likelihood of occurrence to be quantified and assessed against reference values, asserting with sufficiently high confidence that the *residual risk* of SGs violation falls below tolerable rates.

**Motivation.** While for hardware parts,<sup>1</sup> the standard contemplates both systematic and random hardware faults, software faults are all deemed systematic. Yet, software has functional and nonfunctional traits, which may give rise to different fault trees, ill-fit for the homogeneous treatment prescribed by ISO-26262. This problem becomes apparent for *execution-time exceedance* events (i.e., the violation of the worst-case execution-time (WCET) boundaries), which is a nonfunctional

<sup>1</sup>In this paper we focus on the functional safety of the computer subsystems in cars, using the term *hardware* to refer to embedded computers within the automotive E/E; likewise we use *software* to refer to applications.

Manuscript received July 31, 2017; revised November 16, 2017 and February 20, 2018; accepted April 14, 2018. Date of publication May 24, 2018; date of current version August 30, 2018. This work was supported in part by the Spanish Ministry of Science and Innovation under Grant TIN2015-65316-P, in part by the HiPEAC Network of Excellence, and in part by the CONCERTO project (ARTEMIS-JU Grant 333053). The work of J. Abella was supported in part by the Ministry of Economy and Competitiveness under Ramon y Cajal postdoctoral fellowship number RYC-2013-14717. The work of C. Hernández was supported in part by the Spanish Ministry of Economy and Competitiveness and in part by FEDER funds through Grant TIN2014-60404-JIN. The work of E. Mezzetti was supported in part by the Spanish Ministry of Economy and Competitiveness under Juan de la Cierva-Incorporación Postdoctoral Fellowship number IJCI-2016-27396. Associate Editor: A. Romanovsky. (*Corresponding author: Irun Agirre.*)

I. Agirre and M. Azkarate-askatsua are with the Department of Dependable Embedded Systems, IK4-IKERLAN, Mondragón 20500, Spain (e-mail: iagirre@ikerlan.es; mazkarateaskasua@ikerlan.es).

F. J. Cazorla is with the Computer Architecture—Operating Systems group, Barcelona Supercomputing Center, Barcelona 08034, Spain, and also with Artificial Intelligence Research Institute, Spanish Council for Scientific Research (IIIA-CSIC), Barcelona 08193, Spain (e-mail: francisco.cazorla@bsc.es).

J. Abella, C. Hernandez, and E. Mezzetti are with the Computer Architecture—Operating Systems group, Barcelona Supercomputing Center, Barcelona 08034, Spain (e-mail: jaume.abella@bsc.es; carles.hernandez@bsc.es; enrico.mezzetti@bsc.es).

T. Vardanega is with the Department of Mathematics, Università degli Studi di Padova, Padova 35122, Italy (e-mail: tullio.vardanega@math.unipd.it).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TR.2018.2828222

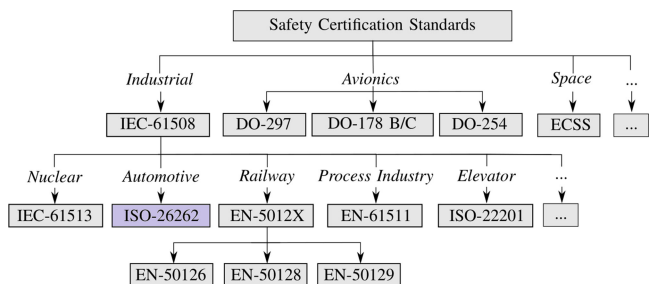


Fig. 1. Safety standards in different application domains and those inheriting from IEC-61508 (including ISO-26262).

trait, evidently involving software *and* hardware concerns. An incorrect (optimistic) WCET estimation may be the root cause of a possible deadline violation, and thus, of a timing failure. For instance, the system design may assign a task an insufficient execution-time allowance, and this underprovision may go unnoticed because the established boundary value is only exceeded when rare circumstances of hardware/software interaction happen, either undocumented or unknown, or exceedingly hard for the user to reproduce during the WCET analysis. Determining the WCET of a software program is a very difficult task indeed, as the programs' execution time varies much beyond user control (and, sometimes, also comprehension). This strenuous task is being made significantly harder by the massive increase in complexity of the hardware and software of modern automotive systems. Postulating that such execution-time violations can all be prevented by standard procedures defined for systematic faults is becoming increasingly prohibitive, yielding unsatisfactory ratios of effort versus quality of outcome. The following observations manifest the magnitude of the problem.

- 1) While the software embedded in cars already totals hundreds of millions lines of code [18], the computational needs of novel functionalities such as advanced driver assistance are projected to increase by 100x in the next decade [11]. This trend reflects the centrality of software to a rising proportion of the competitive value of the vehicle.
- 2) Those performance needs can only be met with high-performance processors that include multi- and many-core components, with deep cache hierarchies and high-end GPUs (like in the NVIDIA DrivePX [2], RENESAS R-Car H3 [4], QUALCOMM Snapdragon 820 Processor [3], and the Intel Go [1]), with massively increasing hardware complexity.

A string of increasingly powerful WCET analysis techniques for safety-critical systems has been put forward by the research community over the last two decades [7], [45], and commercial tooling exists that implements (some of) them [9], [42]. Yet, the most part of those techniques only really applies to a small subset of relatively simple and highly critical (automotive safety integrity level (ASIL)-D) subsystems running on simple and well-understood processor architectures, thus only covering a fraction of the needs. For most subsystems, therefore, the common industrial practice to upper bound the execution time of real-time software programs uses high-water mark (HWM)

measurements and adds a safety margin to them to account for unobserved behavior. With this practice, the confidence in the resulting estimates rests on the user ability to: 1) understand the hardware internals well enough to capture the major sources of execution-time variability and 2) construct test cases that serve for the WCET determination effectively. This knowledge, together with the addition of a conservative margin, sustains the argument that the risk of missing out relevant situations in the analysis is sufficiently low. While this approach may seem inadequate in comparison to state-of-the-art static analysis solutions for other than low-criticality parts, evidence of (cautious) use of measurement-based methods exists for DO-178C-certified avionics software at the highest criticality [35].

Emerging systems imperil the current timing analysis practices, by challenging the user ability to understand deep enough the sources of the jitter in the hardware internals, and to control them. The former weakness hinders the determination of how hardware–software interactions affect timing; the latter impairs the creation of effective analysis scenarios. In those circumstances, the risk of execution-time exceedance events can be made “sufficiently” low by using either inordinately large margins (hence, renouncing resource efficiency) or lower margins with less support evidence (hence, increasing risk). Either prospect faces the user with a dire conundrum.

Measurement-based probabilistic timing analysis (MBPTA) [8] proposes a set of techniques that require applying small and sustainable changes in the hardware design (or alternatively in dedicated runtime libraries) to cause the system to exhibit a probabilistic—hence, probabilistically analyzable—timing behavior. In this way, MBPTA provides by-construction evidence to quantify the probability of execution-time exceedance events. Earlier work describes how to design *MBPTA-friendly* hardware and software platforms [31], [32] such that execution-time exceedance occurs with an (arbitrarily low) probability. Both hardware [31] and software [32] implementations of the MBPTA support have proven viable even with complex hardware designs (i.e., multicore processors with multilevel cache hierarchies), in space [27] and automotive platforms [29], with successful evaluation in industrial case studies [20], [25], [29]. So far, however, there is lack of understanding of how the probabilistic treatment of execution-time exceedance events can be understood by safety certification standards in general, and ISO-26262, in particular. In this regard, this paper seeks to answer the following research question.

*How does the approach of quantifying the probability of occurrence of execution-time exceedance events fit the scope and intent of ISO-26262?*

*Contribution.* To address this research question, this paper analyzes ISO-26262 and its treatment of faults, describing how probabilistic execution-time analysis can satisfy ISO-26262 prescriptions and how quantitative evidence can be obtained to support certification arguments. Our contention is that to tackle this challenge satisfactorily, we should change the nature of the WCET-analysis problem: standards should be enabled to allow sound quantification of the execution-time exceedance rate (or its likelihood of occurrence), in relation with target failure metrics associated with SGs. This approach would be akin to

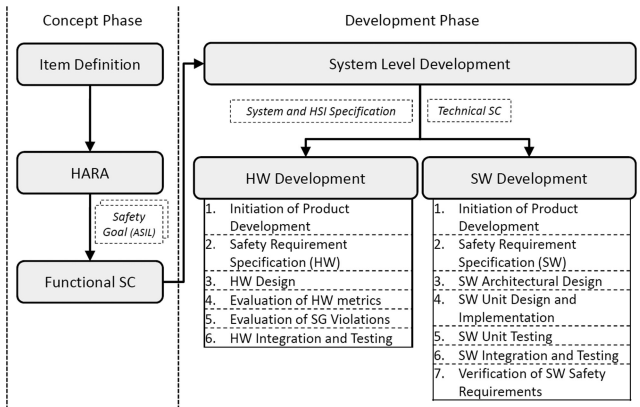


Fig. 2. Schematic view of ISO-26262 concept and development phases.

established practice for hardware random faults in ISO-26262, and would no longer require leaning on qualitative cognizant experience, scarcely available with new hardware, as in current practice for the treatment of systematic faults.

Accordingly, in the remainder of this paper: we survey the management of systematic and random faults in ISO-26262 (see Section II); we show how an asymmetric treatment of software faults that addresses execution-time exceedance probabilistically, fits in the safety life cycle defined in ISO-26262 and how it can be extended to IEC-61508 (see Section III); we present the concept of the MBPTA as a solution to quantify execution-time exceedance rates, and examine the feasibility of applying it to ISO-26262 compliant automotive applications (see Section IV); we provide evidence of the viability of the proposed approach with an automotive case study targeting the AURIX [44], a multicore processor candidate for use in automotive systems (see Section V). Finally, in Section VI, we draw the main conclusions from this study.

## II. HARDWARE AND SOFTWARE FAULTS IN ISO-26262

ISO-26262 requires the user to provide evidence of the absence of unreasonable risk due to hazards caused by the malfunction of E/E systems. For the management of functional safety, ISO-26262 includes a concept definition phase, system, hardware and software development processes, and production and operation measures. Fig. 2 depicts the ISO-26262 workflow for the concept and development phases, which are the focus of this study.

*Concept Phase:* For each item to be developed and certified, the *hazard analysis and risk assessment* step defines the set of hazardous events caused by item's malfunction under specific operational situations. Safety experts classify the hazardous events at different integrity levels—called ASIL—based on their severity, probability of exposure, and controllability. The ASIL levels range from A to D, with D being the most restrictive. Overall, this step formulates the safety goals and associated ASILs for each hazardous event.

For each safety goal, the *functional safety concept* (functional SC) defines the *safety measures* to be implemented in the item. Rather than the technical implementation details, the functional

SC describes the functional *safety requirements* to achieve the safety goal. *Safety measures* include activities for the avoidance of systematic faults and technical *safety mechanisms* to detect and control errors caused by systematic and random hardware faults. Whenever a safety mechanism detects an error, an action shall be taken as defined in the functional SC. In the application domain of ISO-26262, this action typically seeks to achieve or maintain a *safe state*, in which no unreasonable level of risk is known to exist. If the system has a safe state, then it is categorized as *fail safe*.

*Development Phase.* Here, the *technical SC* elicits technical requirements from the functional SC requirements, which determine how the hardware and software parts should implement the functional SC to achieve the stated SG. The ASIL of each SG determines the set of safety requirements assigned to each part. In this way, the stringency of the design is determined by the properties of the possible hazardous events that the parts may influence. At that point, the hardware and the software parts of the system are developed in accord with the technical SC.

### A. Hardware Faults in ISO-26262

ISO-26262 provides quantitative techniques for assessing the safety mechanisms, and the residual risk of violating SGs.

The hardware development process (see Fig. 2) involves the following steps:

- 1) determining and planning functional safety activities in the product initiation phase;
- 2) deriving the hardware safety requirements from the technical SC;
- 3) designing hardware components;
- 4) designing their interconnect at architectural level, and each component in detail, factoring safety requirements in them (i.e., with provisions for fault tolerance);
- 5) evaluating the hardware mechanisms designated to handle faults;
- 6) integrating and verifying the hardware architecture against the system specification.

Steps 4 and 5 include a quantitative analysis of safety mechanisms and residual risk: the hardware architectural metrics defined in step 4 evaluate the effectiveness of the hardware architecture and the implemented safety mechanisms against the fault handling requirements; and step 5 requires evaluating whether the residual risk of safety goal violations is acceptable (i.e., sufficiently low).

ISO-26262 acknowledges that safety techniques cannot achieve full coverage for all types of faults and allows diagnostic coverages even below 90% for the highest-criticality applications. The system may, therefore, be exposed to uncovered faults, which results in residual risk that needs to be assessed. Faults can be classified into following two types: 1) safely ignorable faults (i.e., multiple-point perceived or detected faults) that are regarded as irrelevant since their effects become “visible” before they can do harm, or they are simply harmless; and 2) nonsafely ignorable faults (i.e., single-point faults that are not covered by safety mechanisms, residual faults that may escape safety mechanisms and multiple-point



TABLE I  
TARGET VALUES FOR HARDWARE QUANTIFICATION METRICS [28]

ASIL	Single-Point Fault Metric	Latent Fault Metric	Diagnostic Coverage (Residual Faults)			
			≥ 99,9 %	≥ 99 %	≥ 90 %	< 90 %
D	≥ 99 %	≥ 90 %	FRC 4 (10 <sup>-7</sup> )	FRC 3 (10 <sup>-8</sup> )	FRC 2 (10 <sup>-9</sup> )	FRC 1 (10 <sup>-10</sup> )*
C	≥ 97 %	≥ 80 %	FRC 5 (10 <sup>-6</sup> )	FRC 4 (10 <sup>-7</sup> )	FRC 3 (10 <sup>-8</sup> )	FRC 2 (10 <sup>-9</sup> )*
B	≥ 90 %	≥ 60 %	FRC 5 (10 <sup>-6</sup> )	FRC 4 (10 <sup>-7</sup> )	FRC 3 (10 <sup>-8</sup> )	FRC 2 (10 <sup>-9</sup> )*

(a)

(b)

(\*) + Dedicated Measures

latent faults) that are critical, as they may lead to SGs violation.

ISO-26262 addresses nonsafely-ignorable faults by defining the *single-point fault metric* that determines the item's robustness to single-point and residual faults by either design or safety mechanisms, and the *latent fault metric* that determines the item's robustness to latent faults by either design or safety mechanisms or driver logic diagnosing the fault before SGs violation. The pass/fail reference figures [see Table I(a)] defined for these metrics range between 90% and 99% for single-point faults and between 60% and 90% for latent ones, depending on the target ASIL level.

To assess whether the residual risk is acceptable, strict values are imposed on the allowed failure rates. In one of the methods described in ISO-26262, *failure rate classes* (FRCs) 1 to 5 are defined with different target rates. Table I(b) describes the maximum FRC for hardware parts depending on the *diagnostic coverage* achieved for the hardware faults and the target ASIL. For instance, an ASIL-D SG requires proving residual failure rate  $\leq 10^{-7}$  (FRC 4) when the diagnostic coverage is above 99.9%. Lower failure rates are required if the diagnostic coverage is lower, with higher failure rates allowed if the ASIL level is lower (e.g., C or B).

Overall, the quantitative assessment of random hardware faults provides evidence of whether the resulting design meets its assigned safety requirements.

### B. Software Faults in ISO-26262

ISO-26262 holds a deterministic view of software faults and classifies them all as systematic. Moreover, it assumes that all systematic faults have to be prevented, tolerated, or removed at some stage of the development process. It is for this reason that their contribution to the residual risk is not contemplated. The software development process is similar to the hardware one (see Fig. 2), except that it does not include quantitative analysis.

In practice, however, process-oriented solutions cannot provide positive evidence of the lack of residual faults, especially in the face of the increasing complexity of modern software functions, and the intricate interactions that they may have with advanced hardware. Interestingly, some authors [43] argue that the software complexity combined with that of the associated development process cause faults to be randomly scattered across the program code.

Qualitative analysis is meant to prevent faults in the development phase, not to predict their occurrence during operation.

For qualitative assessment, software variants exist [36], [39] of state-of-the-art techniques that apply to hardware components, such as fault tree analysis and failure mode and effects analysis. The main focus at this level would be on process-level issues, to make sure that all discrepancies between program behavior and functional specification are intercepted. Proactive techniques, such as software fault injection or workload generators can be leveraged to further increase the test coverage and reduce the risk of residual faults. All the aforementioned techniques, however, suffer from the limitations that the quality of their outcomes depends on the user's ability to achieve the sufficient test coverage.<sup>2</sup>

The objective of quantitative analysis, instead, is to predict the occurrence of residual faults. ISO-26262 introduces the quantitative assessment for random hardware faults, to quantify the risk of residual faults and to determine whether it is below the assigned threshold. Our contention here is that the same should be done for software: means should be provided to reason on the probability of residual software faults (whose presence is bound to stem from the increasing complexity of the system), and to relate that probability to given thresholds.

The metric to be used to quantify the risk of residual software faults depends on the specific property, either functional or non-functional, for which the risk needs to be quantified. From the functional/implementation standpoint, a lot of effort has been devoted to study and predict the occurrence of software faults as a ground for reasoning on software reliability. Both deterministic or probabilistic models have been proposed. Deterministic models build on characteristics of the program's code (e.g., Halstead's delivered bugs metric [26] or McCabe's cyclomatic complexity [37]) complementary to those suggested by best practice and guidelines for the software implementation. Probabilistic models instead relate the occurrence of faults in a function to its frequency of execution or, inversely, to the number of tests executed on it [40]. Probabilistic models generally extrapolate the information collected during the test campaign to predict the occurrence of faults during operation. These models derive reliability predictions from trends observed in failure data. Relevant techniques include failure rate, fault count models, or the software reliability growth models [23].

For nonfunctional properties, such as, e.g., the program's timing behavior, the metrics of interest tend to relate to the test quality and the (test) coverage achieved during the development. While a quantitative approach may be needed to assess the residual risk of various types of software faults, in the sequel, we focus on execution-time exceedance events, where a software unit exceeds its assigned budget during operation.

### III. CASE FOR EXECUTION-TIME EXCEEDANCE RATES

ISO-26262 requires establishing upper bounds on the execution time of real-time tasks. The resulting WCET estimates allow deciding how to schedule tasks at run time, thereby, assuring the overall feasibility of system's execution. The provided WCET values should be tight, to avoid waste of processor

<sup>2</sup>Note that the analysis of nonfunctional failures has its own metrics and analysis techniques (including timing and schedulability analysis).

resources. The WCET values should also be consistent with the SG requirements. It is commonly held that any WCET estimate overrun may cause a system-level failure. Yet, this is a misconception since existing safety mechanisms may factor in the execution-time exceedance's impact on the SG, and prevent its escalation into a *timing failure*.

Whereas an execution-time exceedance may not compromise system safety, the timing behavior of software functions should still be characterized to assure proper functioning of the system. It is, therefore, crucial to assess the quality of the provided WCET estimates to assure that they tightly upper bound the application's timing behavior under any possible execution scenario. Unfortunately, as noted, the increasing complexity of modern computing platforms threatens the soundness of *qualitative* assessment of timing correctness, and may allow execution-time exceedance situations to escape prevention.

Execution-time exceedance may result, for instance, from the combination of specific task interleaving, initial cache states, interrupt arrival patterns, and DRAM refresh operations, whose sources are often too remote from the user's reach and too difficult to control and prevent. Accordingly, we contend that execution-time exceedance events should be treated by ISO-26262 similarly to random hardware faults, and the concept of the residual risk should apply to the former too, in conjunction with quantification means as proposed in this paper.

#### A. Timing Analysis Challenges on Complex Systems

With increasingly complex hardware and software, the WCET bounds obtained with traditional means are subject to unquantifiable risk arising from the limitations of the analysis process and the exceeding hardness of the verification procedures. Two main WCET-analysis paradigms have been used so far in industry [45]: static timing analysis (STA) and measurement-based timing analysis (MBTA). Those paradigms and their hybrid variants have been reviewed critically in [7], concluding that, in spite of occasional successes in industrial applications, none of them can be claimed to be effective in the general case and even less so against the relentless increase in complexity of new-generation systems.

While STA is generally held as scientifically sound, confidence in the results of it critically depends on the availability of a detailed and trustworthy timing model of the computing platform underneath the application. Sadly, the latter is increasingly rare, as IP restrictions severely limit the level of detail in public documentation. Hence, for future complex hardware and software systems, STA may become untenable, as obtaining the information needed for it may become too hard or altogether impossible. Evidence of this trend emerges from recent avionics and automotive reports, where the industrial teams and their STA tool providers have been compelled to resort to measurement-based analysis to derive timing bounds for multicore processor architectures like the NXP P4080 [38], Texas Instrument TMS320C6678 [34], and ARM-based SABRE Lite [14]. Industrial pragmatism, therefore, continues to regard the MBTA as the most practicable timing analysis approach even for safety-

related real-time systems, which explains STA's weaker penetration [45].

The MBTA requires identifying the main sources of execution-time jitter, to activate them during analysis. While being far from trivial, this identification is a much easier job than building or acquiring the detailed timing model required by STA, and can be performed by first reviewing processor specifications to identify those resources, and then, using specialized programs called microkernels [38][41] that place a predetermined load on the desired processor resource(s) to quantify their impact on timing. For the MBTA, uncertainty stems from the inherent difficulty in mimicking, during analysis, all of the execution conditions—especially those of jittery processor resources—that can arise during operation. Deriving reliable WCET estimates on the complex hardware requires that low-level architectural features, which can contribute to significant execution-time variations (e.g., cache placement), are factored in the measurement runs taken during analysis so that the observed execution times can be considered *representative* of those that can arise during operation. As complex hardware architectures may have a huge number of potential states with bearing on execution-time jitter, it is not realistically possible to fully explore them during analysis. Hence, by construction, MBTA cannot exclude that residual execution-time exceedance events may occur during operation, reflecting circumstances not covered during analysis.

Common industrial practice to address uncertainties in the WCET analysis requires adding conservative safety margins (often starting at 20%) to the computed WCET value. Any such number, however, evidently lacks scientific grounding and simply rests on engineering judgment. Consequently, this practice may yield either ineffective use of the available resources (due to WCET overestimation) or higher risk of execution-time exceedance events (owing to WCET underestimation), as a result of insufficient quality in the computed bound. Moreover, this practice does not scale to more complex hardware and software. Already on a relatively simple four-core processor, in fact, small variations in execution conditions have been shown to cause either tiny (e.g., below 10%) or huge slowdowns (e.g., up to 20x) [24]. Appropriate means are, therefore, needed to produce tight WCET estimates that can be related to a quantified (and arbitrarily low) risk of execution-time exceedance.

#### B. Probabilistic WCET Distribution

To address this challenge, we build on timing analysis solutions that yield probabilistic distributions of the execution-time behavior of application tasks (nicked probabilistic WCET, pWCET), instead of a single-valued WCET. The pWCET distribution, illustrated in the right side of Fig. 3, represents the probability that a task may exceed the assigned budget envelope at run time. Cutting the tail of it at the desired probability of exceedance ( $10^{-10}$  on the Y-axis, per run or hour of operation) projects onto an execution-time value (7 on the X-axis) that may serve as the WCET budget at that level of assurance. Hence, the pWCET provides means to statistically quantify the likelihood of execution-time exceedance accurately.

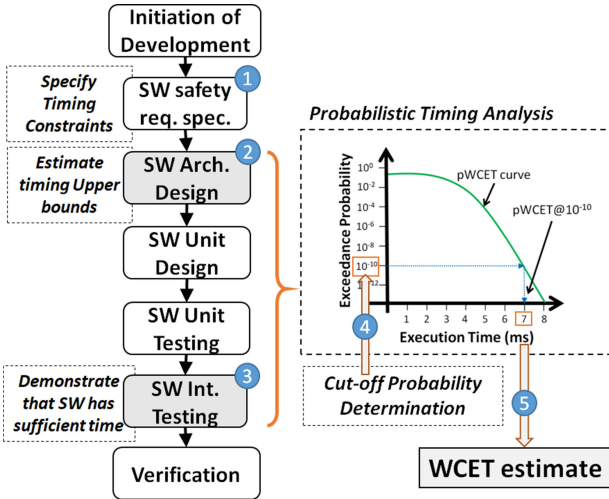


Fig. 3. Sketch of how pWCET fits in the ISO-26262 software development process.

### C. Fitting pWCET into the Safety Life Cycle

Interestingly, the notion of the pWCET distribution follows ISO-26262’s philosophy for the handling of random hardware faults and applies it to the timing domain. Returning to the ISO-26262 life cycle depicted in Fig. 2, with focus on timing-related requirements, we now describe how an approach delivering a pWCET curve can fit in the software development process defined in the standard, which we illustrate in Fig. 3.

In the *concept phase* (not shown in Fig. 3), the functional SC should be extended to also consider the possibility of execution-time exceedance events that can propagate into timing failures, and accordingly, define adequate safety protection measures against them (e.g., watchdog timer).

In the software *development phase*, ISO-26262 includes timing-related requirements in three different phases of the software V-model (sketched in Fig. 3). First, during software safety specification phase ①, it requires system designers to specify the time budgets of critical software. Then, the software architectural design ② shall consider the time upper bounds to dimension the system. If an approach delivering a pWCET distribution instead of a single-valued WCET is used, the designer needs to identify the appropriate probability of exceedance ④ to determine the corresponding WCET from the pWCET curve ③. To this end, the cutoff exceedance probability (or allowed execution-time exceedance rate) shall be evaluated together with the diagnostic coverage for timing errors and the ASIL of the SG. In other words, the standard should provide target metrics for the combination of these three factors as done for random hardware faults in Table I.

For integration testing ③, ISO-26262 requires providing evidence that the software is allocated enough time to complete its functionality. The pWCET distribution allows associating the assigned budget envelope to the corresponding probability of exceedance.

This approach advocates abandoning the current practice of adding a safety margin to the WCET estimate and assuming—on expert judgment only—that it will never be exceeded, and

therefore, exposing to an unquantified risk of execution-time exceedance. In contrast with that, the pWCET improves the soundness of the verification process by providing a quantitative upper bound to the risk of execution-time exceedance estimated with a sound approach. To this end, however, it is of vital importance that the timing analysis technique meets the property of guaranteeing that the delivered pWCET distribution is representative of the worst-case timing behavior that may occur during operation.

### D. Software FRCs and Diagnosis Coverage

We now describe FRCs and diagnostic coverage for execution-time exceedance events so that they can be used as for hardware random faults.

1) *Failure Rate Classes*: The pWCET distribution allows selecting the acceptable rate of execution-time exceedance, normally associated with a *single run* of the task. By multiplying this value by the task’s execution frequency per hour, we determine the execution-time exceedance rate per hour for the task. For instance, in order to assure an execution-time exceedance rate per hour of, e.g.,  $10^{-9}$ , for a program executed  $10^3$  times per hour, the user should cut the pWCET tail at the  $10^{-12}$  exceedance threshold, which would yield a 7.7-ms WCET value in Fig. 3. In this way, it is probabilistically guaranteed that the accumulated execution-time exceedance rate of all instances of the program executed per hour is below  $10^{-9}$ . This reasoning matches random hardware metrics as defined in Table I. Similarly to the hardware case, the particular probability to choose comes from the ASIL level assigned to the software element.

2) *Diagnostic Coverage*: The standard suggests the usage of watchdog timers to detect the consequences that a fault in a hardware component may have in the program schedule (e.g., missed, delayed, or too close activations of the program). In this scenario, the standards of interest categorize the diagnostic coverage achievable by watchdogs for errors in the control logic of processing units as either low (60%) or medium (90%). Accordingly, watchdogs can also detect (possibly with a high,  $>99\%$ , diagnostic coverage) execution-time exceedance events in the operational system. On the occurrence of such an event, the safety mechanisms in place detect the error and instigate action to remove the residual risk of the SG violation. While advising the usage of an external monitoring facility (e.g., watchdog) for the error detection at the software architectural level (which correlates to software faults categorized as systematic), ISO-26262 does not explicitly allude to the achievable diagnostic coverage of mechanisms against execution-time exceedance.

For fail-safe systems, the system should be moved to a safe state every time a diagnostic mechanism detects an execution-time exceedance. In that manner, the SG would be preserved at the expense of making some functionality (or the entire system) unavailable. As a result, the degree of diagnostic coverage that the safety mechanisms provide for timing faults should be taken into account when quantifying the residual failure rate [as it is the case for hardware, see Table I(b)]. Whereas safety is not affected in fail-safe systems in the event of an execution-time exceedance (assuming that high diagnostic coverage mechanisms



are in place), system availability is instrumental for the end user since an unavailable system does not deliver the expected functionality. Arguably, therefore, solutions that yield reliable pWCET distributions can improve the design process by allowing the user to assess the availability of fail-safe systems from a timing perspective.

For fail-operational systems, which need to stay operational to preserve safety, the event of an execution-time exceedance should activate the use of appropriate forms of redundancy or diversity so that the occasional failure of one unit does not stop the (safe) operation of the entire system. Whenever this solution is not possible, having high diagnostic coverage against timing faults is not sufficient to preserve safety and a sufficiently low cutoff probability needs to be chosen to ensure that the contribution of execution-time exceedance to the residual risk is kept correspondingly low.

#### E. From ISO-26262 to IEC-61508

The IEC-61508 meta- (or parent-) standard differs from ISO-26262 only slightly. The latter refines some definitions for the life-cycle phases and provides additional requirements for the safety requirement specification of the hardware. IEC-61508 does not organize the life cycle around concept and development phases explicitly, but rather fragments it into smaller units that match the activities defined within the ISO-26262 workflow. Those activities include, for instance, the *hazard and risk analysis* (which ISO-26262 names *hazard analysis and risk assessment*) and the *overall safety requirements and allocation* (which ISO-26262 places in the *functional safety concept*), where the SIL level, ranging 1–4, is computed. As a rule of thumb, the highest ASIL level in ISO-26262 (ASIL-D) matches on certification ambition an SIL-3 in IEC-61508. The safety requirement specification concerning random hardware faults is lighter in IEC-61508 than in ISO-26262. The hardware concept and development involve the same steps in the metastandard, but the derivation of hardware faults neither includes latent faults nor multiple-point faults, which simplifies the calculations. Regarding software faults, the approach is identical in both standards: software faults are considered systematic and qualitative measures are recommended for fault avoidance, such as WCET analysis to assure temporal independence among software elements.

Like ISO-26262, IEC-61508 determines the requirements for avoiding or controlling systematic faults based on expert judgment from practical experience. IEC-61508 states that “*the probability of occurrence of systematic faults cannot in general be quantified.*” To exemplify this difficulty, IEC-61508 reasoning observes that the effects of systematic faults manifesting at run time, depend on the moment of the life cycle in which they were introduced, and the effectiveness of the prevention measures (e.g., structured programming) in place, which are both difficult to quantify sensibly. However, IEC-61508 allows considering that the target failure reduction for a safety function is achieved by demonstrating compliance to all requirements of the standard. In this regard, the standard introduces the concept of *systematic capability*, which is equivalent to the SIL, but

only considers systematic faults. In addition to systematic fault reduction or prevention in the design, the standard does also define mechanisms to control the run-time errors arising from systematic faults (e.g., diverse software redundancy).

Overall, IEC-61508 retains the notion of random hardware faults and proposes a qualitative approach for software faults that may not scale well against increasingly complex systems. Arguably, therefore, all the application domains covered by the IEC-61508 umbrella might equally benefit from incorporating an execution-time exceedance quantification approach, much like the automotive domain would do via ISO-26262 following the solution presented in this paper.

#### IV. MBPTA: CONCEPT AND APPLICATION

As a particular probabilistic timing analysis solution, we build on the MBTA variant proposed in [8], [30], [31], called MBPTA. MBPTA yields a reliable pWCET distribution while guaranteeing, by construction, that the delivered pWCET is an upper bound of the execution conditions that may occur at system operation: it, therefore, fits the ISO-26262 exceedance rate quantification approach presented in Section III.

MBPTA acknowledges that the control that the user can exercise on the application’s timing behavior during analysis necessarily leverages high-level metrics such as software code coverage, but has increasingly less means to address low-level hardware aspects (e.g., bus occupancies, placement of program’s code/data in cache) comprehensively. Hence, MBPTA relieves the user from the latter burden by introducing some platform modifications.

The application of the MBPTA rests on the premise that the computing platforms that enable its use [31], [32] modify the timing behavior of selected jittery resources so that the execution-time measurements collected during analysis either match or upper bound probabilistically the timing behavior that may occur during operation. In that manner, the obtained pWCET distribution is warranted to capture any extreme behavior that may occur at operation, and it is produced without burdening the user with the need to comprehend all system states relevant to execution-time analysis.

If hardware support is provided to enable the use of the MBPTA, the processor vendor is the party in charge of singling out jittery resources and of designing MBPTA compliance around them appropriately. Interestingly, using MBPTA, the processor vendor would not need to build a timing model of its processor, or granting access to all details of the hardware design as STA requires. All it would be required of the vendor is to design processor resources that can be explicitly and individually configured to feature the desired forms of the MBPTA conformance, and to document them in public user manuals.

Conversely, if hardware support for MBPTA were scarce or inexistent, the user would have to identify the sources of execution-time variation building on the processor specifications, and apply software solutions to reach MBPTA compliance. In general, the resources that cause the largest jitter (e.g., cache memories, interconnection networks, and memory controllers) are easy to identify. Missing out some sources of jitter,

while not desirable, is not particularly harmful as long as the jitter that they may produce is not larger than the cumulative effect of the execution-time variation produced by the other known sources.

To handle jittery resources, MBPTA defines two main techniques, implemented in either hardware [31] or software [32], which we present in the following subsections.

#### A. Time Upper Bounding

This technique forces selected jittery hardware resources to work at their highest latency during analysis. In that manner, the operation conditions cannot lead to higher execution times from them, and hence, a single run suffices to capture their worst-case operation-time behavior. The hardware resources that best fit the use of this technique are those whose extended variation in timing behavior depends on elements that the hardware cannot discriminate efficiently [19], [27].

The floating-point unit (FPU) provides an illustrative example of this kind of resources. The latency of floating-point (FP) operations depends on the operands, outside of the hardware's own control. For instance, multiplying any value by 0.0 may incur shorter latency than multiplying any pair of not-null parameters. Hence, for the analysis of even the simplest sequential program that included FP operations, capturing the full extent of latencies that it might incur would require enumerating all of the executed FP operations and their respective operands, which is unduly onerous and likely to involve laborious debugging. On top of that, the user would also need to determine whether the distribution of the FP operations and operands observed during analysis is representative of what may occur at operation, which is even harder, if at all possible. Instead, MBPTA's prescription to force the FPU to work at its highest latency (per operation type) during analysis relieves the user from the burden of controlling the impact that each FP operation incurs on the program execution time.

Original FPUs allow serving the result and releasing as soon as the current operation finalizes. To implement the said technique, the hardware default is modified by deactivating the immediate-release check so that all operations take maximum latency regardless of the input operands. The hardware feature that allows enforcing the highest latency can be enabled or disabled by setting the corresponding configuration register accordingly, so that it can be kept enabled during analysis and disabled during operation. In that manner, operation-time behavior may experience shorter, but never longer, latency.

Time upper bounding also applies to other resources such as, e.g., the number of arbitrated contenders on the shared bus that connect cores to a shared L2 cache [19], [27]. For a program running on a core, the contentions suffered during operation depend on the software being run on the other cores. This information is exceedingly difficult to determine during analysis even for the strictest of static scheduling scenarios, since the arrival time of bus requests from contenders may change across different execution paths and cache hit/miss patterns. To address this challenge, a simple modification to the hardware arbiter is applied [27] to cause arbitration to occur across all potential

contenders regardless of whether they have pending requests or not, keeping the bus busy for the longest request latency after selection. Selectively disabling this feature during operation allows the program to experience fewer stalls than contemplated for WCET analysis.

#### B. Time Randomization

This technique causes the response time of some jittery resources to exhibit a probabilistic behavior that also holds during operation. Accordingly, a representative distribution of the impact that jittery resources may cause on execution time can emerge after a statistically significant number of observation runs. For instance, randomizing the placement and replacement of objects in cache memories, allows using execution-time measurements to model cache behavior probabilistically. Such randomization makes cache conflicts independent of the memory location of program objects, which relieves the user from the need to control memory placement. In integrated modular architecture systems as used in avionics [5] and automotive [13], individual software applications are often subcontracted to different providers. As a result, the integration of the system progresses incrementally, requiring to assess at every step of integration that the new build conforms with the specification, for functional and nonfunctional requirements. However, as applications get integrated into the system binary, their memory placement and cache layout may vary [22], invalidating the WCET estimates computed previously. This phenomenon defers timing verification to the latest stages of integration, where the (binary) image is near final, which in turn makes timing faults much more costly to handle than during earlier phases of the development.

Randomized caches mimic the behavior of multiple software integrations, which allows the WCET estimates computed in earlier development stages to hold across the whole process of integration as well as during operation.

To date, time randomization has been implemented in processor resources whose timing behavior depends on the structural dependencies created by the hardware design. Randomization helps remove those dependencies, which have no bearing on the program semantics. For instance, whether two addresses compete for the same cache space depends on how they are mapped to cache lines. And cache mapping can be randomized to make conflicts occur probabilistically. To ensure that the observations made during analysis represent (probabilistically) the timing events that may occur during operation, such randomization must be kept enabled at all times, with no distinction between analysis and operation. Random placement and random replacement have been successfully implemented in the hardware [19], [27].

This technique is evidently superior to the "time upper bounding" alternative of modifying the cache hardware to respond with the highest (miss) latency during analysis, owing to the massive performance decay incurred by the latter.

With little difficulty, time randomization has also been applied to the bus arbiter, changing the way it chooses which core is granted access to the shared L2 cache. Bus arbiters, there-






Fig. 4. Example results of MBPTA application.

fore, have two types of modifications: time upper bounding to determine the number of contenders and the latency with which the bus is released; and time randomization to choose which core is granted access to the shared L2 cache.

All hardware parts that use time randomization require a hardware source of randomness. An exemplary pseudorandom number generator (PRNG) has been implemented to that end [10], with a degree of randomness that has passed the most stringent cryptographic tests. The cited publication shows that the hardware cost of its implementation is low, also because a single PRNG can be shared across multiple resources. The PRNG has also been proven compatible with high safety integrity levels. If time randomization is to be implemented in software, a software implementation of the very same PRNG algorithm can also be used.

### C. Probabilistic Analysis

The execution time of the program “inflated” by time upper bounding and randomization results in an analysis-time distribution (ATD) that upper bounds the operation-time distribution (OTD) by construction. Fig. 4 illustrates this notion. The dotted line depicts the empirical complementary cumulative distribution (ECCDF) of the OTD, and the dashed line the ECCDF of the ATD. A sound use of *probabilistic analysis* [such as, e.g., extreme value theory (EVT)] uses a sample of ATD values—no less than a hundred, and typically up to two thousands, which keeps the MBPTA overhead low—to derive a high-quality pWCET distribution that upper bounds the ATD (and hence, the OTD) [8]. For the EVT to be applicable, the observed execution times must correspond to independent and identically distributed (i.i.d.) random variables, which means that each measurement observation must belong to the same execution-time distribution. Satisfying this requirement has been proven doable with simple-enough procedures [15].

The MBPTA process collects execution-time samples from the ATD, earning MBPTA conformance thanks to the hardware modifications discussed earlier, and to a measurement collection process that controls the initial conditions of the experiment [15]. The analysis procedure may determine that the sample fails to meet the eligibility criteria for the application of the EVT or detect that it cannot be upper bounded by exponential tail distributions, which is required to ensure tightness. These situations are addressed by enlarging the sample size. It is known, in fact, that increasingly larger samples from an i.i.d. random variable with a guarantee finite bound will eventually

be proven statistically i.i.d., and also converge, more tightly, to either exponential or light tails, the former always upper bounding the latter. In our analogy, the program’s execution-time observations are the random outcomes of that variable, and the program itself has bounded duration in conformance with the well-established real-time coding practice. At that point, the larger the sample size, the tighter the pWCET. Accordingly, MBPTA users should collect large—yet affordable—samples, below 2000 measurements on average [8].

MBPTA promotes a paradigm shift with respect to the traditional, deterministic (i.e., single-valued) WCET analysis. The relation of the MBPTA with its deterministic counterpart is straightforward: MBPTA’s main constituents (*time upper bounding* and *randomization*) specifically address the representativity concerns that afflict standard measurement-based approaches, and threaten to become insurmountable with increasingly complex systems. Relating MBPTA to STA is much harder instead, as those two techniques build on largely different (and mostly incompatible) assumptions [6]. The correctness and the precision of either of them depend on whether and to what extent their assumptions are guaranteed to hold; see [7] for a detailed analysis of those assumptions and how they relate to hardware and software complexity.

### D. MBPTA: Industrial Viability

MBPTA’s viability for industrial use in safety-related systems relates to the cost of the required hardware or software changes, and how the approach can be fitted in the overall ISO-26262 safety life cycle as discussed in Section III-C.

The latter question leverages the need to step up the guidelines of current safety standards to increasing complexity of new-generation processors. This has been done, for instance, in the avionics domain, where CAST32 [16] and the accompanying CAST-32A [17] address the use of multicore processors. Arguably, this game-changing scenario should ease the task of incorporating MBPTA-related changes.

The MBPTA requirements on the computing platform, if implemented at the hardware level, have been shown affordable, first by implementation in architectural simulators, then at the register-transfer level (RTL) in a field-programmable gate array, and finally, in off-the-shelf products [19]. Implementing randomization has been surprisingly nonintrusive. We illustrate this for two cases. Bus protocols like AMBA [12] (one of the most, if not the most, used), do not define any particular arbitration policy. This situation allows adding random arbitration policies with no impact on the protocol specification. The same happens for the cache placement and replacement. While the latter is already supported in many processors, adding the former requires combining the address being accessed with a hardware- (or software-) generated random seed [10], changed across runs, to map the address to a random cache set. This change causes the timing behavior of cache conflict scenarios that are probabilistically relevant—those whose timing behavior can only be exceeded with negligible probability—to be close to average behavior, which, in turn, is very close to the typical behavior on conventional hardware designs.

At software level, randomization has been implemented as a pass in the LLVM compiler [30] or as a source-to-source translator developed in an approach called toolchain-agnostic software randomization (TASA) [32]. Both solutions leverage the fact that the way in which functions and data (locals and globals) are placed in the source code and the binary determines their address in memory. By randomly allocating them and adding padding space among them keeps the program functionality unchanged and attains similar randomized timing to that obtained with hardware-implemented random placement. As opposed to the hardware and LLVM-based software solutions, which attain randomization at program run granularity, the TASA approach applies randomization on a per-binary basis. As a result, the probability of exceedance determined by the use of the TASA is equivalent to the execution-time exceedance probability of all systems with the same randomly generated binary. For the hardware and LLVM-based software randomization cases instead, the obtained probability is per run of the program, and therefore, has to be multiplied by its rate. Time upper bounding at a software level is managed offline, by monitoring relevant events during the analysis-time measurements [through performance monitoring counters (PMC)] and by padding execution-time observations so that their impact on the program's execution time is deterministically upper bounded. For instance, reading a PMC that returns the quantity of FP operations executed by a program, allows computing a padding equivalent to each FP operation taking the highest latency. Similarly, a bus jitter can be deterministically upper bounded by monitoring the number of bus access requests with PMCs and applying a contention model that assumes worst-case overlap among them [21].

Both hardware and software randomization and upper bounding solutions have proven effective for the performance in various platforms and with negligible implementation costs:  $\approx 1\%$  additional hardware for a four-core processor for the space domain [27], and a preprocess compiler step for TASA source-to-source transformations in the automotive domain [32]. Moreover, pWCET estimates have been shown to be typically within 20% of the HWM on conventional (time deterministic) platforms used as reference for industrial applications in the space, avionics, and automotive domains [20], [25], [29]. This evidence proves that time randomization and upper bounding do not incur untenable pessimism.

## V. EXPERIMENTAL SUPPORT EVIDENCE

To sustain our contention, we discuss an exemplary application of the MBPTA in an automotive case study targeting the AURIX TC277 [44]. We show that, even on a processor whose hardware design expressly seeks maximum determinism, the execution-time behavior of applications running on it suffers jitter created by resources that may be hard, if at all possible, for the user to control. We do *not* aim to present a full WCET analysis method for the TC277: our intent is just to show that the execution-time jitter of hard-to-predict resources like the cache—a definite and massive asset of future automotive processors [1]–[4]—can be handled with the MBPTA.

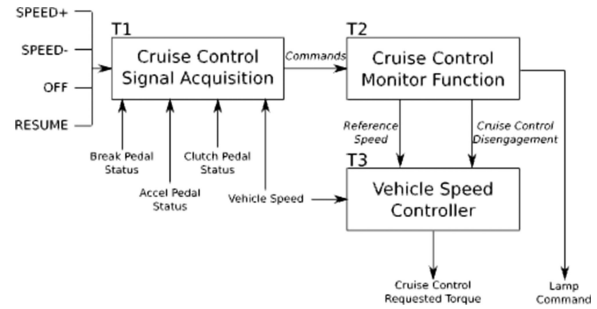


Fig. 5. Block diagram of the case-study application.

### A. Application Case

The AURIX TC277 comprises three cores (plus two additional ones that operate in lockstep mode): one energy-efficient core and two performance-efficient ones. We focus on the latter, which embed high-performance jittery resources such as caches and dynamic branch predictors. All cores are equipped with local scratchpad memories and caches, for both instruction and data, and are connected via a crossbar to a common “memory system” comprising a shared SRAM, and program/data flash memories.

The application we consider is an automotive cruise control system, whose functional code was automatically generated from a Simulink model, and a CONCERTO<sup>3</sup> model for its architectural specification. The application was run on a customized version of ERIKA Enterprise,<sup>4</sup> which implements an OSEK/VDX compliant personality. The originating Simulink model comprised in excess of 200 blocks, which corresponded to about 3000 lines of C code. After a transformation process that flattened the Simulink block hierarchy, optimally regrouped the blocks by compatible rates, and regenerated source code accordingly, the application was embedded in the following three real-time ERIKA tasks: 1) signal acquisition; 2) monitoring; and 3) speed controller, as shown in Fig. 5. A fourth task, status update, was added to the application to close the simulation loop in the experiment by stubbing and interconnecting all input and output ports.

For the purposes of this paper, we discuss the impact of the instruction cache layout on the application's execution-time behavior, to showcase its relevance as a source of jitter, and demonstrate how the MBPTA can capture its contribution in the analysis process. To this end, we deployed the application on the processor such that part of the code was stored (and cached from) the program Flash memory segment. Private stack and data were located in the local scratchpads while local shared data were mapped to the shared SRAM. The instruction cache size in the performance-efficient cores is 16 kB, with 32-B cache lines. Other sources of variability and combinations thereof are not considered here. How to jointly account for them is discussed in [21].

<sup>3</sup>CONCERTO, ARTEMIS JU, <http://www.concerto-project.org/>.

<sup>4</sup>Erika Enterprise RTOS, <http://erika.tuxfamily.org/drupal/>.

### B. Instruction Cache Jitter

To analyze the impact of the cache layout on the execution-time jitter, we used TASA, a technique that applies software randomization to off-the-shelf caches to make their response time probabilistically analyzable. As explained in Section IV, TASA performs source-to-source program code transformations where the relative location of functions, stack frames, and global variables is randomized by reordering and padding the source file so that each resultant binary incurs a randomly different cache placement. Thus, by studying the timing behavior of a statistically significant number of binaries, the impact of caches can be accounted for probabilistically. In our experiments, we generated 1000 distinct binaries with TASA, all with identical functionality, and each with different stack and global data allocations to memory positions. Simple *ad hoc* scripts were needed to automate this process: they invoke the TASA preprocess pass and compile the output of it to produce one binary; this process is repeated as many times as needed, varying the random seed so that the required number of binaries is obtained. Each such binary needs to be run once for the purposes of the MBPTA. The computational cost of this process is proportional to two characteristics of the software being considered: 1) its size and complexity and 2) its execution time. The former determines the cost of generating the binary, largely dominated by the compile time, much more complex than the TASA preprocessing step. The latter determines the cost to execute each binary. In our particular case, generating the binary and executing it took around 5 s altogether (the most part for the compilation), which serialized in less than 1.5 h for all binaries. Of course, binary generation and execution can be parallelized in multiple instances: as the process is fully independent per binary, the turn-around time would decrease roughly linearly with the degree of parallelization. The remainder of the MBPTA process (acquiring the collected execution times, and producing the pWCET distribution) takes just a few seconds.

While the cost of this process for a single program is rather low, it would increase linearly for applications that include multiple programs assigned to a criticality level that requires evidence of bounded execution time. However, each such program could be analyzed separately, thus allowing the analysis process to proceed independently (and perhaps with internal parallelism). In general, devoting around 1–2 hours of computational cost (less if parallelized) to the timing analysis of each program, with minimal user intervention, should be affordable even for complex applications.

Fig. 6 reports the execution-time variability observed for the four application tasks—for the same program path—as determined by the different randomly generated program layouts. For the system under analysis, the observed variability, which may incorporate the effects of other sources of execution-time jitter, ranged up to approximately 5%. In all cases, the HWM was quite distant from the observed average and mode.<sup>5</sup> The impact of time upper bounded jittery resources, computed offline based on PMC measurements, should then be added to

<sup>5</sup>The mode of a data sample is the most frequent value.

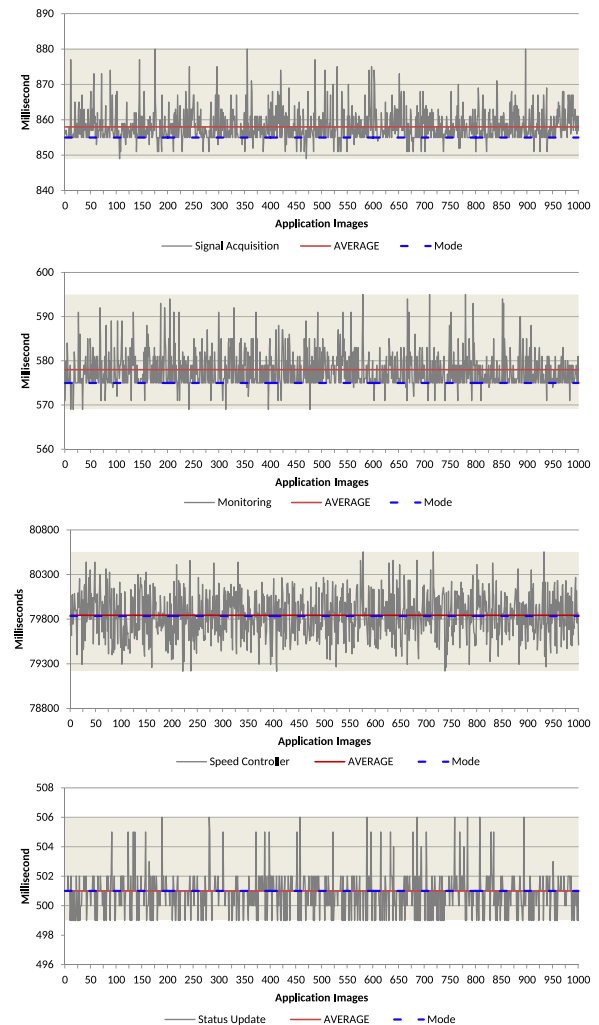


Fig. 6. Uncontrolled variability induced by the program layout.

these execution-time observations before obtaining the pWCET distribution. This kind of variability is not explored by state-of-the-art WCET analysis procedures, measurement-based and static alike. Even more critical is the fact that the traditional measurement-based techniques *do not support constructing arguments* on whether and to what extent the effect of jittery resources has been captured at analysis time. Next, we show how MBPTA can consider these effects in the determination of pWCET bounds.

### C. Application of the MBPTA

We applied MBPTA to the four tasks of the case-study application. According to the TASA prescriptions, we collected timing measurements for each such task by executing the same set of 1000 randomly generated binaries of the application. The collected observations successfully passed the statistical i.i.d. tests (a precondition to apply statistical analysis), which allowed using them as input to the subsequent probabilistic analysis process. To the latter end, we used the MBPTA method using the coefficient of variation (MBPTA-CV) [8], which applies the EVT [33] by automatically selecting the distribution



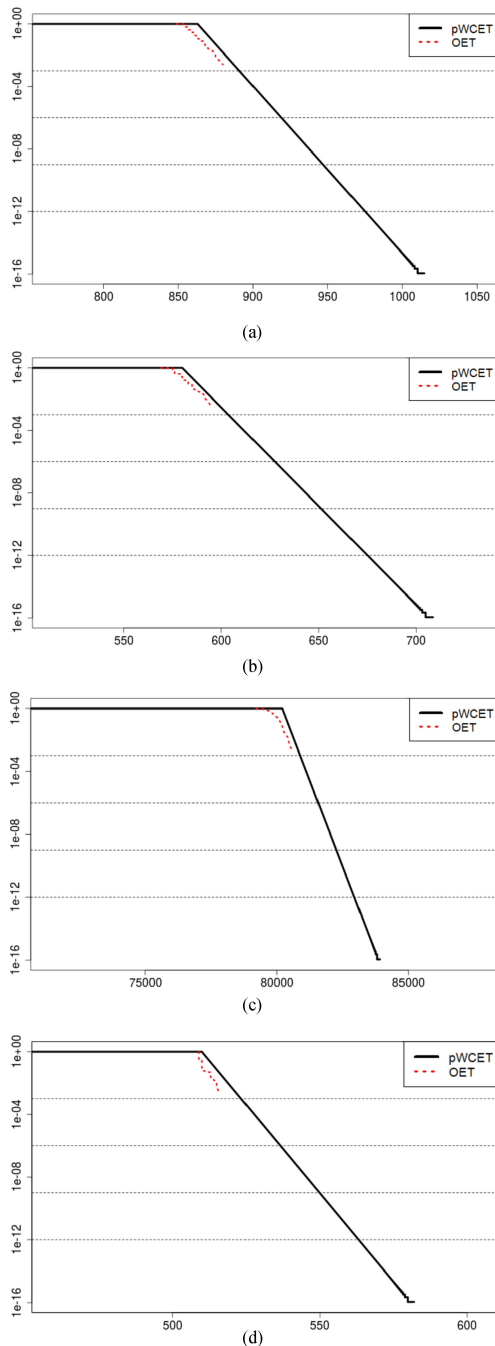


Fig. 7. pWCET distributions computed by the MBPTA for the four application tasks. (a) Signal acquisition. (b) Monitoring. (c) Speed control. (d) Update status.

parameters that best fit the maxima of the observed execution times. In all cases, 1000 measurements were sufficient for the MBPTA to converge: adding additional observations for each application would *not* change the resulting pWCET distributions shown in Fig. 7.

The red dotted line in Fig. 7 plots the observed execution times (OET), in the form of complementary cumulative distribution function (CCDF), to show that the pWCET curves (solid black lines) always tightly upper bound the observed data. The pWCET bounds for the analyzed functions at relevant exceedance thresholds are reported against the HWM in

TABLE II  
pWCET BOUNDS AT RELEVANT EXCEEDANCE THRESHOLDS  
(IN PROCESSOR CYCLES)

Test	HWM	$10^{-3}$	%	$10^{-6}$	%	$10^{-9}$	%	$10^{-12}$	%
Signal Acq.	880	891	<b>1.2</b>	919	<b>4.4</b>	947	<b>7.6</b>	975	<b>10.8</b>
Monitoring	595	603	<b>1.3</b>	627	<b>5.4</b>	651	<b>9.4</b>	674	<b>13.3</b>
Speed Control	80554	80888	<b>0.4</b>	81574	<b>1.3</b>	82260	<b>2.1</b>	82946	<b>3.0</b>
Status Update	506	511	<b>1.0</b>	521	<b>3.0</b>	531	<b>4.9</b>	541	<b>6.9</b>

Table II. The application of the MBPTA-CV to the automotive functions led to extremely tight results as, when compared to their respective HWM, the predicted pWCET bounds are always below the reference 20% margin. The low distances for higher exceedance thresholds can be partially ascribed to the overall high predictability of the execution platform.

An intuitive but wrong conclusion here might be that the 20% margin is a reliable figure in the general case. In fact, our experiments show that, limited to the processor considered in this paper, and focusing only on the instruction cache, a conservative margin at 20% would be conservatively pessimistic, and therefore, sound. Yet, for other processor architectures, with complexity similar to the technology used in the automotive domain [1]–[4], that margin would be optimistic instead, hence, unsound [25].

In actual fact, the slope of the pWCET distribution, hence, the margin above the highest observed value for the acceptable exceedance probability, depends on the particular characteristics of the program under analysis and how it uses the underlying processor hardware. Such margin (or the pWCET value itself) is guaranteed not to be exceeded with a particular probability (e.g.,  $10^{-12}$  per run). That would be the exceedance probability if *all* the sources of the jitter that have been upper bounded, always caused their highest latency. This does not happen in the general case; how often it may, cannot be told beforehand as it depends on the input-dependent behavior of the application during operation. In this situation, the MBPTA method allows the user to strictly upper bound the residual risk of execution-time exceedance. Conversely, time-deterministic approaches building on a margin set on the expert judgment for a particular platform, hardly scale to other (arbitrarily complex) platforms and also do not provide means to assess the residual risk.

## VI. CONCLUSION

ISO-26262 classifies hardware faults as either systematic or random, while it considers all software faults to be systematic. The unrelenting demand for newer value-added functionalities for computer-based systems requires the use of increasingly complex hardware and software. This trend challenges the viability of exhaustive analysis and prevention for all types of systematic faults as prescribed by the standard. This threat is especially true for the timing behavior of software applications, as the fabric of new systems denies users the ability to capture all sources of execution-time variations and to create the test scenarios needed to estimate the residual risk of failure. Recent timing analysis techniques that deliver WCET estimates with an associated probability of exceedance have the potential to overcome this limitation. However, how the quantification of the

likelihood of execution time exceedance events fits the scope and intent of safety standards such as ISO-26262 is still an open research question. In this paper, we address this question by proposing ISO-26262 adaptations to assess the residual risk associated to exceeding the timing budget assigned to a software program, in analogy to what is done for random hardware faults. This approach relies on timing-analysis techniques that deliver a probabilistic WCET bound and serve the purpose of the upper bounding residual risk. We exemplify this approach with a particular incarnation of the MBPTA, which transparently applies time randomization to selected hardware or software elements of the computing platform, in this manner, relieving the user from the burden of controlling the impact of low-level hardware elements on the software execution time. This proposal is presented in the context of the ISO-26262 software development process and the treatment of random hardware faults in the safety life cycle, with the intent of promoting the acceptance of execution-time exceedance rate quantification in the standard.

#### ACKNOWLEDGMENT

The authors would like to thank Intecs SpA, lead of CONCERTO who provided the sources of the automotive application, and the University of Padova, the build automation for the AURIX target.

#### REFERENCES

- [1] Intel GO Automated Driving Solution Product Brief. [Online]. Available: <https://www.intel.es/content/www/es/es/automotive/go-automated-accelerated-product-brief.html>, Accessed on: May 1, 2018.
- [2] NVIDIA DRIVE PX, “Scalable supercomputer for autonomous driving.” [Online]. Available: <http://www.nvidia.com/object/drive-px.html>. Accessed on: May 1, 2018.
- [3] QUALCOMM Snapdragon 820 Automotive Processor. [Online]. Available: <https://www.qualcomm.com/products/snapdragon/processors/820-automotive>. Accessed on: May 1, 2018.
- [4] RENESAS R-Car H3. [Online]. Available: <https://www.renesas.com/en-us/solutions/automotive/products/rcar-h3.html>. Accessed on: May 1, 2018.
- [5] Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment. SAE International, ARP4761, 2001.
- [6] J. Abella, D. Hardy, I. Puaut, E. Quiñones, and F. J. Cazorla, “On the comparison of deterministic and probabilistic WCET estimation techniques,” in *Proc. Euromicro Conf. Real-Time Syst.*, 2014, pp. 266–275.
- [7] J. Abella et al., “WCET analysis methods: Pitfalls and challenges on their trustworthiness,” in *Proc. IEEE Int. Symp. Ind. Embedded Syst.*, 2015, pp. 1–10.
- [8] J. Abella, M. Padilla, J. Del Castillo, and F. J. Cazorla, “Measurement-based worst-case execution time estimation using the coefficient of variation,” *ACM Trans. Des. Autom. Electron. Syst.*, vol. 22, no. 4, pp. 72–1–72–29, Jun. 2017.
- [9] AbsInt Angewandte Informatik GmbH, *aiT WCET Analyzers*. [Online]. Available: <https://www.absint.com/ait/>. Accessed on: May 1, 2018.
- [10] I. Agirre et al., “IEC-61508 SIL 3 compliant pseudo-random number generators for probabilistic timing analysis,” in *Proc. Euromicro Conf. Digit. Syst. Design*, 2015, pp. 677–684.
- [11] ARM, *ARM Expects Vehicle Compute Performance to Increase 100x in Next Decade*, Apr. 2015. [Online]. Available: <https://www.arm.com/about/newsroom/arm-expects-vehicle-compute-performance-to-increase-100x-in-next-decade.php>
- [12] ARM Ltd. “AMBA open specifications.” [Online]. Available: <http://www.arm.com/products/system-ip/amba/amba-open-specifications.php>. Accessed on: May 1, 2018.
- [13] AUTOSAR Technical Overview V2.2.1, AUTOSAR GbR, Munich, Germany, 2006. [Online]. Available: <https://www.autosar.org/>
- [14] A. Blin, C. Courtaud, J. Sopena, J. Lawall, and G. Muller, “Maximizing parallelism without exploding deadlines in a mixed criticality embedded system,” in *Proc. 28th Euromicro Conf. Real-Time Syst.*, 2016, pp. 109–119.
- [15] F. J. Cazorla, T. Vardanega, E. Quiñones, and J. Abella, “Upper-bounding program execution time with extreme value theory,” in *Proc. Int. Workshop Worst-Case Execution Time Anal.*, 2013, pp. 61–70.
- [16] Certification Authorities Software Team, “Multi-core processors—Position paper,” Certification Authorities Software, Tech. Rep. CAST-32, May 2014.
- [17] Certification Authorities Software Team, “Multi-core processors—Position paper,” Certification Authorities Software, Tech. Rep. CAST-32A, Nov. 2016.
- [18] R. N. Charette, “This car runs on code,” *IEEE Spectr.*, 2009. [Online]. Available: <https://spectrum.ieee.org/transportation/systems/this-car-runs-on-code>, Accessed on: Feb. 1, 2009.
- [19] COBHAM, “LEON3 Processor. Probabilistic platform.” [Online]. Available: <http://www.gaisler.com/index.php/products/processors/leon3>. Accessed on: May 1, 2018.
- [20] F. Cros et al., “Dynamic software randomisation: Lessons learned from an aerospace case study,” in *Proc. Design, Autom. Test Eur. Conf. Exhib.*, Mar. 2017, pp. 103–108.
- [21] E. Díaz et al., *MC2: Multicore and Cache Analysis via Deterministic and Probabilistic Jitter Bounding*. Cham, Switzerland: Springer, 2017, pp. 102–118.
- [22] E. Mezzetti and T. Vardanega, “A rapid cache-aware procedure positioning optimization to favor incremental development,” in *Proc. Real-Time Embedded Technol. Appl. Symp.*, 2013, pp. 107–116.
- [23] W. Farr, “Software reliability modeling survey,” in *Handbook of Software Reliability Engineering*. Hightstown, NJ, USA: McGraw-Hill, Inc., 1996, pp. 71–117.
- [24] M. Fernández, R. Gioiosa, E. Quiñones, L. Fossati, M. Zulianello, and F. J. Cazorla, “Assessing the suitability of the NGMP multi-core processor in the space domain,” in *Proc. ACM Int. Conf. Embedded Softw.*, 2012, pp. 175–184.
- [25] M. Fernandez et al., “Probabilistic timing analysis on time-randomized platforms for the space domain,” in *Proc. Design, Autom. Test Eur. Conf. Exhib.*, Mar. 2017, pp. 738–739.
- [26] M. H. Halstead, *Elements of Software Science (Operating and Programming Systems Series)*. New York, NY, USA: Elsevier, 1977.
- [27] C. Hernández et al., “Design and implementation of a time predictable processor: Evaluation with a space case study,” in *Proc. 29th Euromicro Conf. Real-Time Syst.*, 2017, pp. 16–1–16–23.
- [28] *Road Vehicles—Functional Safety*, International Organization for Standardization, ISO/DIS 26262, 2009.
- [29] L. Kosmidis et al., “Measurement-based timing analysis of the AURIX caches,” in *Proc. 16th Int. Workshop Worst-Case Execution Time Anal.*, 2016, pp. 9:1–9:11.
- [30] L. Kosmidis, C. Curtsinger, E. Quiñones, J. Abella, E. Berger, and F. J. Cazorla, “Probabilistic timing analysis on conventional Cache designs,” in *Proc. Design, Autom. Test Eur. Conf. Exhib.*, 2013, pp. 603–606.
- [31] L. Kosmidis et al., “Fitting processor architectures for measurement-based probabilistic timing analysis,” *Microprocess. Microsyst.*, vol. 47, pp. 287–302, 2016.
- [32] L. Kosmidis, R. Vargas, D. Morales, E. Quiñones, J. Abella, and F. J. Cazorla, “TASA: Toolchain-agnostic static software randomisation for critical real-time systems,” in *Proc. 35th Int. Conf. Comput.-Aided Design*, 2016, pp. 59–1–59–8.
- [33] S. Kotz and S. Nadarajah, *Extreme Value Distributions: Theory and Applications*. Singapore: World Scientific, 2000.
- [34] A. Kritikakou et al., “Distributed run-time WCET controller for concurrent critical tasks in mixed-critical systems,” in *Proc. 22nd Real-Time Netw. Syst.*, 2014, pp. 139–148.
- [35] S. Law and I. Bate, “Achieving appropriate test coverage for reliable measurement-based timing analysis,” in *Proc. 28th Euromicro Conf. Real-Time Syst.*, Toulouse, France, Jul. 5–8, 2016, pp. 189–199.
- [36] N. G. Leveson and P. R. Harvey, “Software fault tree analysis,” *J. Syst. Softw.*, vol. 3, no. 2, pp. 173–181, 1983.
- [37] T. J. McCabe, “A complexity measure,” *IEEE Trans. Softw. Eng.*, vol. SE-2, no. 4, pp. 308–320, Dec. 1976.
- [38] J. Nowotzsch, M. Paulitsch, D. Buhler, H. Theiling, S. Wegener, and M. Schmidt, “Multi-core interference-sensitive WCET analysis leveraging runtime resource capacity enforcement,” in *Proc. 28th Euromicro Conf. Real-Time Syst.*, 2014, pp. 109–118.

- [39] H. Pentti and H. Atte, "Failure mode and effects analysis of software-based automation systems," VTT Industrial Systems, *Espoo Finland STUK-YTO-TR 190*, 2002, p. 190.
- [40] H. Pham, *System Software Reliability* (Springer Series in Reliability Engineering). London, U.K.: Springer-Verlag, 2006.
- [41] P. Radojković, S. Girbal, A. Grasset, E. Quiñones, S. Yehia, and F. J. Cazorla, "On the evaluation of the impact of shared resources in multithreaded COTS processors in time-critical environments," *ACM Trans. Archit. Code Optim.*, vol. 8, no. 4, pp. 34-1-34-25, 2012.
- [42] Rapita Systems Ltd. Rapita Verification Suite (RVS). [Online]. Available: <https://www.rapitasystems.com/products/rvs>, Accessed on: May 1, 2018.
- [43] P. H. Seong, *Reliability and Risk Issues in Large Scale Safety-Critical Digital Control Systems*, 1st ed. London, U.K.: Springer, 2008.
- [44] *AURIX Application Kit TC277 TFT*, hitex. [Online]. Available: <http://www.ehitex.de/application-kits/infineon/2531/aurix-application-kit-tc277-tft>. Accessed on: May 1, 2018.
- [45] R. Wilhelm *et al.*, "The worst-case execution-time problem overview of methods and survey of tools," *ACM Trans. Embedded Comput. Syst.*, vol. 7, pp. 1-53, May 2008.

**Irene Agirre** received the M.Sc. degree in embedded systems from Mondragon Unibertsitatea, Mondragón, Spain. She is a Ph.D. student with the Universitat Politècnica de Catalunya, Barcelona, Spain.

She is a Researcher with the Dependable Embedded Systems Department, IK4-Ikerlan Technology Research Center, Mondragón, Spain. Her research interests include functional safety of real-time embedded systems.

**Francisco J. Cazorla** received the Ph.D. degree in computer science from the Polytechnic University of Catalonia, Barcelona, Spain, in 2005.

He is the Director of the Computer Architecture Operating System Research Group, Barcelona Supercomputing Center and a Researcher at IIIA-CSIC, Barcelona, Spain. His research interests include hardware design and performance analysis of embedded real-time and high-performance systems.

**Jaume Abella** received the Ph.D. degree in computer science from the Polytechnic University of Catalonia, Barcelona, Spain, in 2005.

He is a Senior Ph.D. Researcher with the CAOS group at the Barcelona Supercomputing Center, since 2009. His research interests include timing and functional verification of critical real-time systems and performance analysis. He is a member of HiPEAC.

**Carles Hernandez** has been a Ph.D. Researcher with the Barcelona Supercomputing Center, Barcelona, Spain, since 2012. He has participated in several European research projects and with the European Space Agency on time predictable and reliable high-performance processor designs. In 2012, he was an intern with the IP verification group at Intel Munich.

**Enrico Mezzetti** received the Ph.D. degree in computer science from the University of Bologna, Bologna, Italy.

He is currently a Researcher with the CAOS group at BSC. His research interests include industrial-level timing analysis of embedded real-time systems with a special focus on multi and many core platforms.

**Mikel Azkarate-askatsua** received the Ph.D. degree in computer science from Technische Universität Wien, Vienna, Austria.

He leads the real-time systems research group at IK4-Ikerlan, Mondragón, Spain. He coordinates the SAFEPOWER H2020 project and several other R&D activities on the implementation of embedded systems for real-time control systems, control system modeling, and functional-safety certification.

**Tullio Vardanega** received the Ph.D. degree in computer science from the Technical University of Delft, Delft, The Netherlands.

He was a staff member with the European Space Agency Research and Technology Center. He has been with the University of Padova, Padova, Italy, since 2002. His research interests include high-assurance real-time systems and software engineering methods to develop them.