Probabilistically Time-Analyzable Complex Processors in Hard Real-Time systems

Mladen Slijepcevic[1,2], Jaume Abella[2], Francisco J. Cazorla[2,3]

[1]Universitat Politècnica de Catalunya (UPC),
[2]Barcelona Supercomputing Center (BSC-CNS),
[3]Spanish National Research Council (IIIA-CSIC)

mladen.slijepcevic@bsc.es

Abstract-*Critical Real-Time Embedded Systems (CRTES) feature performance-demanding functionality. High-performance hardware and complex software can provide such functionality, but the use of aggressive technology challenges time-predictability. Our work focuses on the investigation and development of (1) hardware mechanisms to control inter-task interferences in shared time-randomized caches and (2) manycore network-on-chip designs meeting the requirements of Probabilistic Timing Analysis (PTA).*

I. Introduction

Industry developing CRTES, such as Aerospace, Space, Automotive, and Railways, face relentless demands for increased processor performance to support advanced new functionalities. Multiple indicators suggest that these demands will continue to grow across almost all sectors of the CRTES industry. Multicores are well accepted as one of the main design paradigms to increase performance. However, current generation CRTES, based on relatively simple single-core processors, are already extremely difficult to analyse in terms of their temporal behaviour. The advent of multicore and manycore platforms exacerbates this problem, rendering traditional temporal analysis techniques unable to scale and ineffectual, with potentially dire consequences for the quality and reliability of future products. In this context multicores for CRTES must balance the achievement of trustworthy and low Worst-Case Execution Time (WCET) estimates, high performance and low design complexity while meeting the needs of mixed-criticality workloads.

II. Background

A. PTA

Timing analysis techniques for time-deterministic hardware [1] deliver a single WCET estimate. However, the pessimism of the WCET estimate grows if not enough information about hardware internal behaviour is available or hardware is complex and not amenable to WCET analysis.

Conversely, PTA [2,3,4,5] provides a distribution of WCET estimates so that the particular value at a given exceedance probability - a so called Probabilistic WCET (pWCET) estimate - can be theoretically exceeded with a probability upper-bounded by the exceedance threshold chosen, which can be arbitrarily low (e.g., 10^{-12} per hour), thus largely below the probability of hardware failures. In this work we focus on the Measurement Based version of PTA (MBPTA) as it is closer to industrial practice.

MBPTA uses Extreme Value Theory (EVT)[6] on the execution time measurements obtained at analysis time. EVT is a well-known statistical method to approximate the tail of distributions, and so to derive the pWCET distribution. Thus, the execution time value at the desired exceedance threshold can be used as the pWCET estimate for the program under analysis. Figure 1 shows a hypothetical result of applying EVT to a collection of 1,000 observed execution times.

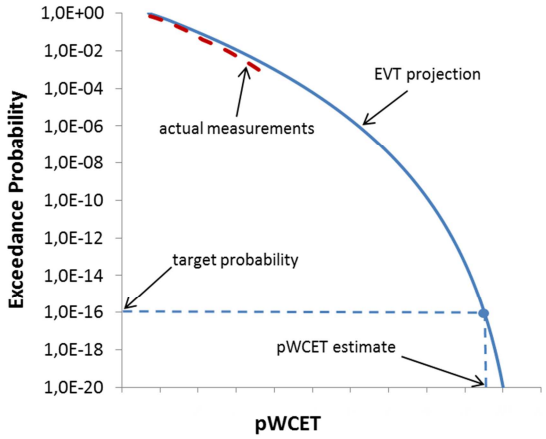

Fig. 1. Example of the 1-CDF and tail projection.

III. Probabilistically Upped-Bounding Inter-Task Interference Features

Shared caches in multicores challenge WCET estimation due to inter-task interferences. Hardware and software cache partitioning [7,8,9,10] address this issue although they complicate data sharing among tasks and the Operating System task scheduling and migration. We propose a technique [11] that overcomes the limitations of cache partitioning by enabling the estimation of trustworthy and tight WCET estimates for systems equipped with fully-shared (non-partitioned) time-randomised last level caches (LLCs). The main principle behind our proposal is that, while in a time-deterministic LLC interferences depend on when (time) and where (the particular cache in which) misses occur, a time-randomised LLC removes any dependence on the particular addresses accessed and its assigned cache set. This makes that the LLC interferences that a task suffers only depend on how often (frequency) its co-runner tasks miss in cache – thus evicting data – and not the particular address generating the miss. Based on this analysis we propose a simple hardware mechanism that limits the miss frequency of tasks in each core at analysis and deployment time in a manner that probabilistic upper-bounds can be obtained for the effect in the LLC of one

task on the other co-running tasks. Our approach removes cache partitioning constraints while making WCET estimates tighter.

III. NoC

Among manycore shared resources the network-on-chip (NoC) has prominent impact on programs' execution time and WCET estimates, as it connects cores to memory and/or shared cache levels. Among existing NoC designs, only buses have been proven MBPTA-compliant [12] for different arbitration policies. However, bus scalability is limited since bus latency increases rapidly with the number of cores. We propose a new tree-based NoC design that is compatible with MBPTA requirements and that delivers scalability towards medium/large core counts.

IV. Conclusions

Guaranteed performance needs of CRTES require using high-performance hardware designs but those jeopardise time predictability needed by conventional timing analyses. MBPTA has emerged recently as a powerful method to derive WCET estimates for critical tasks in safety-related systems on the top of complex hardware. We present new techniques to obtain time-composable WCET estimates on the top of shared non-partitioned LLCs, thus removing partitioning constraints, and MBPTA-compliant tree NoCs that outperform buses in multicores with 8/16 cores.

Acknowledgment

The research leading to these results received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under the PROXIMA Project (www.proxima-project.eu) grant agreement no. 611085. This work has also been partially supported by the Spanish Ministry of Science and Innovation under grant TIN2012-34557 and the HiPEAC Network of Excellence. Mladen Slijepcevic is funded by the Obra Social Fundación la Caixa under grant Doctorado "la Caixa"—Severo Ochoa. Jaume Abella is partially supported by the Ministry of Economy and Competitiveness under Ramon y Cajal postdoctoral fellowship number RYC-2013-14717.

References

[1] 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.
[2] F.J. Cazorla et al. PROARTIS: Probabilistically analysable real-time systems. ACM TECS, 2013.
[3] L. Cucu-Grosjean et al. Measurement-based probabilistic timing analysis for multi-path programs. In ECRTS, 2012.
[4] S. Edgar and A. Burns. Statistical analysis of WCET for scheduling. In RTSS, 2001.
[5] J. Hansen, S Hissam, and G. A. Moreno. Statistical-based wcet estimation and validation. In WCET Workshop, 2009.
[6] S. Kotz and S. Nadarajah. Extreme value distributions: theory and applications. World Scientific, 2000.
[7] H. Kim, A. Kandhalu, and R. Rajkumar. A coordinated approach for practical os-level cache management in multi-core real-time systems. In ECRTS, 2013.
[8] J. Liedtke, H. Hartig, and M. Hohmuth. OS-controlled cache predictability for real-time systems. In RTAS, 1997
[9] B. Ward, J. Herman, C. Kenna, and J. Anderson. Making shared caches more predictable on multicore platforms. In ECRTS, 2013.
[10] M. Paolieri et al. Hardware support for WCET analysis of hard real-time multicore systems. In ISCA, 2009.
[11] M. Slijepcevic et al. Time-analysable non-partitioned shared caches for real-time multicore systems. In DAC, 2014.
[12] J. Jalle et al. Bus designs for time-probabilistic multicore processors. In DATE, 2014.