Time-Predictable Design for Cyber-Physical Systems

Wei Zhang
Associate Professor
Department of Electrical and Computer Engineering
Southern Illinois University Carbondale
Carbondale, IL 62901
Email : zhang@engr.siu.edu

1. Introduction

Embedded computing systems are increasingly networked and conjoined with physical systems in many application domains spanning from automobiles and aircrafts to power grids and assisted living. Our society is more and more relied on these cyber-physical systems that tightly integrate the computation and physical systems. Many cyber-physical systems are under real-time constraints, and thus a critical research challenge is how to ensure time predictability of the distributed cyber-physical systems that consist of sensors, microprocessors, actuators and wireless/wired communication networks, considering the fact that predicting the worst-case execution time (WCET) of modern microprocessors is already overwhelmingly complicated due to architectural features such as cache memories, pipelines and branch prediction [1].

2. Time-Predictable Multi-Core Architecture

The computer industry is rapidly moving towards the multi-core processors. Compared to uniprocessors, multi-core chips offer a significant boost in processing capability while consuming less energy and less board space. Therefore, multi-core processors can potentially benefit future high-performance cyber-physical systems. For example, from the viewpoint of time predictability, segregating different real-time tasks onto different CPUs makes it easier to determine whether they'll meet their deadlines [2]. Also, fault tolerance is typically vital for safety-critical CPS such as fly- and drive-by-wire systems, and a multi-core processor can cost-effectively improve reliability by exploiting redundant cores that are tightly coupled on the same chip.

Multi-core processors, however, are generally designed for improving the average-case throughput, which will further complicate the WCET analysis as compared to uniprocessors, mainly due to the possible interferences from other co-running threads in using the shared resources, for instance the shared caches and memory, which can significantly increase the worst-case execution time (WCET). Actually, our recent work has discovered the timing anomaly of multi-core computing [4], which makes accurate timing analysis for multi-core chips extremely complex, if not impossible. Therefore, to ensure time-predictable multi-core computing for future cyber-physical systems, we believe it is a necessity to customize the multi-core architectural design.

More specifically, we propose to design a time-predictable and high-performance multicore/manycore architecture (TMulticore) to provide scalable and predictable performance for future high-end cyber-physical systems. In TMulticore, it is important to either prioritize or separate the shared resources to avoid the inter-core interferences; while at the same time, it should minimize the impact on the performance. For instance, in a multi-core processor with a shared L2 cache, we can give higher priority to real-time threads so that their data and instructions cannot be evicted by non-real-time threads. Alternatively, it is feasible to explore intelligent cache partitioning for balancing time predictability and performance. We believe time-predictable multi-core design by eliminating the interferences on the integrated system can bring other benefits to CPS as well, including the increase of composability and the prevention of fault propagation.

3. Integrated Timing Analysis for Distributed Systems

Unlike most traditional embedded systems, cyber-physical systems are typically designed as networks of interacting elements rather than standalone devices [3]. Predicting the worst-case execution time of networked systems is more challenging than calculating the WCET of a standalone computing system. Most of WCET work has focused on predicting the worst-case performance of individual microprocessors. While there are relative a few research efforts on timing analysis for real-time network applications, e.g. [5], these studies generally either target very simple networks or compromise the tightness of timing analysis.

We envision that future timing analysis for cyber-physical systems need to closely integrate the timing analysis modules for both computing and communication to derive tight WCET for the whole system. While separating the timing analysis for computation and communication may reduce the complexity of analysis, it may be too conservative as well. The reason is that the worst-case computation time and the worst-case communication time may not happen at the same time. Therefore, integrating time analysis for both computation and communication for real-time tasks can increase the tightness of analysis. Specifically, the timing analyzer can exploit a compiler to extract useful data flow, control flow and communication patterns, as embedded processors usually use pre-defined instructions or registers to transmit data. Then integer linear programming (ILP) can be used to integrate the timing analysis of computation tasks and communication tasks to derive precise WCET.

4. Retargetable Timing Analysis Tools for Heterogeneous CPS

Cyber-physical systems used in various applications often employ different microprocessors and different network protocols. The microprocessors may vary from simple microcontroller to multithreaded or multicore processors. And the communication protocols may vary from time-triggered to event-triggered protocols. For instance, in a car like Volvo S80, it contains more than 30 embedded processors, which communicate across several networks. These

heterogeneous computing and communication models lead to different timing; therefore, the timing analysis tool must be re-designed for each computing and communication model used in the CPS, which will result in significant design efforts and may affect time-to-market. We envision that to support future heterogeneous cyber-physical systems efficiently, the timing analysis tool must be made retargetable, just as a compiler should be retargetable to different processor architectures.

We believe it is important to provide abstract timing models of computing and communication to the timing analysis tool. The timing analysis is purely based on abstract timing models, which is independent of specifics of architectures and networking protocols. Therefore, when a CSP uses a different kind of microprocessor or networking protocol, we simply need to update the abstract timing model to derive the WCET, without having to re-write the timing analyzer.

5. Summary

In this paper, we have described the critical challenge of time-predictable design for future CPS. First, we envision that multi-core processors will be increasingly used in future high-end CPS to improve performance, energy efficiency and fault tolerance or to reduce delay, jitters, and cost. Today's multi-core architecture design generally aims at improving average-case performance, which can significantly compromise time predictability. Consequently, it is crucial to design time-predictable multi-core processors to remove the inter-core interferences and to support real-time computing. In a typical CPS, which embedded processors communicate with each other through networks, it is important to integrate the processor timing analysis with the communication timing analysis to derive tight WCET for the system. Also, we emphasize the importance of retargetable timing analyzer to support heterogeneous cyber-physical systems.

References :

[1] R. Wilhelm, J. Engblom, A. Ermedahl, N. Holsti, S. Thesing, D. Whalley, G. Bernat, C. Ferdinand, R. Heckman, T. Mitra, F. Mueller, I. Puaut, P. Puschner, J. Staschulat, P. Stenstrom. The Worst-case execution time problem - overview of methods and survey of tools. In ACM Transactions on Embedded Computing Systems, January 2007.
[2] W. Wolf. The good news and the bad news (Embedded Computing Column). IEEE Computer, Nov 2007.
[3] E. Lee. Cyber physical systems: design challenges. Technical Report No. UCB/EECS-2008-8, University of California, Berkeley.
[4] J. Yan, W. Zhang. WCET analysis for multi-core processors with shared instruction caches. In Proc. of RTAS, April 2008.
[5] Y. Liang, A. Roychoudhury and T. Mitra. Timing analysis of body area network applications. In Proc. of the 7th International Workshop on Worst-Case Execution Time.