

PROPOSITION DE SUJET DE THESE

Intitulé : Cache interference mitigation through bandwidth regulation for RISC-V based architectures

Référence : **TIS-DTIS-2025-10**
(à rappeler dans toute correspondance)

Début de la thèse : 01/10/2025

Date limite de candidature : 15/06/2025

Mots clés

HW architecture, multi-core, embedded systems, interference mitigation, timing analysis

Profil et compétences recherchées

Profile: Bac+5/M2 in computer science, telecommunications, electronics or similar

Skills: HDL (knowledgeable), embedded systems (knowledgeable), computer architecture (knowledgeable)

Présentation du projet doctoral, contexte et objectif

The need for more performance and, at the same time, less Size, Weight and Power constraints (SWaP) have driven computer architectures to move from monocoresh designs to multicore ones [1]. As a result, some resources are shared among different cores (e.g., shared cache, DDR SDRAM), leading to high inter-core interference. This kind of interference may be responsible for high execution time variability, and consequently, be the cause for a loss in execution determinism and performance. This problem has since become a concern in the safety-critical real-time community where assuring a task finishes within its deadlines is imperative. Efforts have been made into it by trying to mitigate their effects, suppress them or make them predictable via software or hardware [2], most of the time at expense of performance.

Among the resources shared in a multicore platform, the memory hierarchy and especially the shared Last Level Cache (LLC) is a strong source of interference due to concurrent accesses, evictions, snooping for coherence, attached buffers saturation, etc. [3]. To cope with this issue, different types of techniques and mechanism have been proposed. Spatial techniques, such as cache partitioning or cache locking [4], are the most common ones. Fundamentally, their aim is to dedicate a portion of the shared LLC exclusively to each of the cores sharing it, i.e., effectively turning a shared LLC into a private one. Other techniques, however, opt to leave the cache undivided and control interference. For example, bandwidth regulators [5] allows to manage the number of requests issued by a core for a given time. This last approach requires a software control logic which may reduce performance in case the proper hardware is not present.

We propose to focus on this last type of solutions for mitigating interference at the LLC while minimizing performance degradation. We believe hardware designs can satisfy these two requirements. To carry out this task, open-hardware designs will be considered as a base (e.g., RISC-V based processors [6], system on chip [7]) and pertinent ameliorations or additions will be implemented upon them. Therefore, the main goal of this thesis is to explore hardware mechanisms for memory requests regulation at the shared LLC in order to mitigate the interference impact of this type of caches while making them dynamic, adaptable and performant. To achieve the previous, the following points will be treated:

- Elaboration of the state-of-the-art of the existing interference and mitigation mechanisms.
- Based on the previous, propose and formalize new interference mitigation mechanisms.
- Verify the proposed mechanisms via tests on an FPGA or a formal proof.

[1] R.M. Ramanathan. "White paper: Intel® Multi-Core Processors Making the Move to Quad-Core and Beyond", 2006.

[2] T. Lugo, S. Lozano, J. Fernández and J. Carretero, "A Survey of Techniques for Reducing Interference in Real-Time Applications on Multicore Platforms," in *IEEE Access*, vol. 10, pp. 21853-21882, 2022, doi: 10.1109/ACCESS.2022.3151891.

[3] Michael Garrett Bechtel and Heechul Yun, "Denial-of-Service Attacks on Shared Cache in Multicore: Analysis and Prevention", in *CoRR*, vol 1903.01314, 2019.

[4] R. Mancuso, R. Dudko, E. Betti, M. Cesati, M. Caccamo and R. Pellizzoni, "Real-time cache management framework for multi-core architectures," 2013 IEEE 19th Real-Time and Embedded Technology and Applications Symposium (RTAS), Philadelphia, PA, USA, 2013, pp. 45-54, doi: 10.1109/RTAS.2013.6531078.

[5] H. Yun, G. Yao, R. Pellizzoni, M. Caccamo, and L. Sha, "MemGuard: Memory bandwidth reservation system for efficient performance isolation in multi-core platforms," in *Proc. IEEE 19th Real-Time Embedded Technol. Appl. Symp. (RTAS)*, Apr. 2013, pp. 55-64

- [6] Zhao, Jerry and Abraham Gonzalez. "Sonic BOOM: The 3rd Generation Berkeley Out-of-Order Machine." (2020).
[7] A. Amid et al., "Chipyard: Integrated Design, Simulation, and Implementation Framework for Custom SoCs," in IEEE Micro, vol. 40, no. 4, pp. 10-21, 1 July-Aug. 2020, doi: 10.1109/MM.2020.2996616.

Collaborations envisagées

IRIT (Thomas Carle)

Laboratoire d'accueil à l'ONERA

Département : Traitement de l'information et Systèmes

Lieu (centre ONERA) : Toulouse

Contact : Alfonso Mascarenas-Gonzalez

Tél. : 05 62 25 29 09 Email : alfonso.mascarenas@onera.fr

Directeur de thèse

Nom : Frédéric Boniol

Laboratoire : ONERA

Tél. : 05 62 25 26 62

Email : frederic.boniol@onera.fr

Pour plus d'informations : <https://www.onera.fr/rejoindre-onera/la-formation-par-la-recherche>