



UNIVERSITÀ DI PISA
DIPARTIMENTO DI INGEGNERIA DELL'INFORMAZIONE
Dottorato di Ricerca in Ingegneria dell'Informazione

Doctoral Course

“Exploiting computer architecture issues to improve performance and fairness in commercial multicores”

Prof. Julio Sahuquillo

UPV (Universidad Politécnica de Valencia) -*Valencia, Spain*

Short Abstract:

Current multicore processors implement shared resources like the last level cache (LLC) or the main memory. Requests from multiple applications compete among them in the shared resources. As a consequence, inter-application interference rises and the performance of individual applications slowdown with respect to isolated execution. This problem aggravates when considering simultaneous multithreading (SMT) cores, which implement shared resources within the core, resources that are critical for performance like the caches or the register files.

Current Operating Systems do not consider, or poorly consider these interferences, yielding the system to underrated performance.

In this course, we will study where the interference can arise, how to measure it, and we will study some approaches to reduce the interference. With this aim, we will focus first on some microarchitectural issues of current processors, the main reasons why multicore processors emerged, and on multicore evaluation. Once the basics are studied, we will focus on current technologies, like Intel CAT, that allow us to mitigate the slowdown that individual application suffer and improve the overall performance. Main memory and GPU aspect will be also studied.

We will also perform some labs to reinforce the theoretical concepts studied during the course.

Course Contents in brief:

- 1. Superscalar and multithreaded processors
- 2. Multicore Processors, why?
- 3. Multicore Evolution
- 4. Performance Evaluation of Multicores
- 5. Accounting Architectures
- 6. Cache memories: concepts and problems and advanced topics
- 7. Main memory - GPU

Total # of hours: 20

References:

Basically, references will consist on specific papers depending on the studied topic.

- [1] Kristof du Bois, Stijn Eyerman, Lieven Eeckhout: Per-Thread Cycle Accounting in Multicore Processors, TACO 2013
- [2] Stijn Eyerman, Lieven Eeckhout, Tejas Karkhanis, and James E. Smith: A performance counter architecture for computing accurate CPI components, ASPLOS 2006
- [3] S. Eyerman et al., System-level performance metrics for multiprogram workloads, IEEE Micro, vol. 28, no. 3, pp. 42–53, May 2008
- [4] Josué Feliu, Julio Sahuquillo, Salvador Petit, José Duato: Bandwidth-Aware On-Line Scheduling in SMT Multicores. IEEE Trans. Computers 65(2): 422-434 (2016)
- [5] Josué Feliu, Stijn Eyerman, Julio Sahuquillo, Salvador Petit: Symbiotic job scheduling on the IBM POWER8. HPCA 2016: 669-680, 2015
- [6] Josué Feliu, Julio Sahuquillo, Salvador Petit, José Duato: Addressing Fairness in SMT Multicores with a Progress-Aware Scheduler. IPDPS 2015: 187-196, 2014
- [7] Josué Feliu, Salvador Petit, Julio Sahuquillo, José Duato: Cache-Hierarchy Contention-Aware Scheduling in CMPs. IEEE Trans. Parallel Distrib. Syst. 25(3): 581-590 (2014)
- [8] Josué Feliu, Julio Sahuquillo, Salvador Petit, José Duato: L1-bandwidth aware thread allocation in multicore SMT processors. PACT 2013: 123-132
- [9] Mark D. Hill, and Michael R. Marty, Amdahl's Law in the Multicore Era, Computer July 2008
- [10] P. Michaud, Demystifying multicore throughput metrics, IEEE CAL, vol. 12, no. 2, pp. 63–66, 2013.
- [11] Kunle Olukotun, et al., The Case for a Single-Chip Multiprocessor, ASPLOS 1996
- [12] Vicent Selfa, Julio Sahuquillo, Lieven Eeckhout, Salvador Petit, María Engracia Gómez, Application Clustering Policies to Address System Fairness with Intel's Cache Allocation Technology. PACT 2017: 194-205

CV of the Teacher

Julio Sahuquillo received his BS, MS, and PhD degrees in Computer Science from the Universidad Politécnica de Valencia (UPV), Spain. He is a Full Professor at the Department of Computer Engineering. He has taught multiple courses on computer organization and architecture.

His research topics have included multiprocessor systems, cache design, instruction-level parallelism, GPUs, photonics interconnects, and power dissipation. He has been the principal investigator of multiple research projects including two Spanish National projects and four

research projects funded by the Valencian Government. He is the principal investigator of the UPV in the H2020 ExaNeSt European Project, two and four

During these projects he has advised 9 PhD Thesis on computer architecture, and he is currently advising 5 PhD students. Some of the advised PhD Thesis have received important awards like the thesis advised to Alejandro Valero who received the 2013 Intel Doctoral Student Honor Programme Award (Alejandro Valero) with an amount of 28000 dollars, and the “Premio extraordinario de tesis doctoral at the UPV”. The Thesis advised to Josué Feliu in 2017 has been awarded with the “2017 Premio Sociedad Científica Informática de España y la Fundación BBVA” with 5000 euros”. According to google scholar his work has received 1500 cites (November 2017) and receives around 140 to 160 cites each year.

He has published more than 45 papers in journals listed in the JCR and more than 90 papers in refereed conferences. His work has been published in top conferences like IEEE International Symposium on High-Performance Computer Architecture (HPCA), International Symposium on Microarchitecture (Micro), Parallel Architectures and Compilation Techniques (PACT), IPDS, DATE, etc. Regarding journals his work has been published 16 times in IEEE Transactions (IEEE Transactions on Computers (6 times), IEEE Transactions on Parallel and Distributed Systems (5 times), IEEE Transactions on Very Large Scale Integration, ...) and other top journals in Q1 and Q2 of the JCR list.

He has also chaired more than 20 workshops, collocated in conjunction with IEEE supported conferences.

Room and Schedule

Room: *Aula Riunioni del Dipartimento di Ingegneria dell'Informazione, Via Diotisalvi, Pisa* on July 2018 (from 9 to 14)

Schedule:

- | | |
|-----------------------|--|
| 9/7/2018 – 5 hours. | 1. Superscalar and multithreaded processors |
| | 2. Multicore Processors, why? |
| | 3. Multicore Evolution |
| 10/07/2018 – 5 hours. | 4. Performance Evaluation of multicores |
| | 5. Accounting architecture |
| | 6. Cache memories (I): concepts and problems. |
| 11/07/2018– 5 hours | 7. Cache memories (II): advanced topics from real hardware machine perspective |
| | Cache Sharing, fairness, prefetching, and scheduling |
| | 8. Introduction to GPUs |
| | 9. Main memory controller and memory organization: concept and problem |
| 12/07/2018– 5 hours | Labs. Practical Experience |