PhD position

Continental Automotive – Université de Toulouse (IRIT lab)

The advent of multi-core processors in real-time embedded systems raises the problem of dealing with the interference that occurs between tasks running in parallel on several cores, when they simultaneously request access to shared components such as a sequential interconnect or the memory controllers. Simultaneous requests are handled sequentially, which hampers parallelism and negatively impacts the actual performance of the system by slowing down the tasks execution. It also has an impact on the worst-case execution/response time analysis of the task system, as the static analysis of the amount of interference adds pessimism. The problem of dealing with interference in multi-core processors is a hot topic in the real-time community, and has not been solved satisfactorily so far.

Heterogeneous cluster for real-time applications

The objective of this thesis is to propose a multi-core cluster architecture that enables the timing predictability of real-time applications by controlling the level of interference from the interconnect down to the execution core: this will allow bounding the worst-case execution or response time of the tasks run on this cluster.

The main existing approaches to control interference are: (a) enforcing temporal isolation, e.g. using a TDMA bus [2] or NoC that prevents access conflicts at the cost of performance, since the static allocation of the bus to cores does not permit its full usage; (b) software solutions based on variants of the multi-phase model [1] in which the tasks are organized in such a way that they access shared resources only in specific execution phases, which requires a costly adaptation of existing software with no guarantee that the task scheduler can avoid all the conflicts.

In this thesis, we will explore hardware solutions. The main idea is to dedicate one core per cluster (subsequently called the control core) to perform all the accesses to the main memory and to handle task synchronizations. Computation cores, which are equipped with private local memories, then execute tasks that do not access shared resources, apart from their local memory which they only share with the control core. The control core loads the necessary instructions and data from/to the global memory to/from the local memories of the computation cores.

This approach covers several aspects:

- hardware design of a cluster: structure of the interconnect, organization of local memories, architecture of the control core[1]. The control core must be equipped with facilities to control the transfers between the main memory and the local memories, as well as the synchronizations between tasks that execute on different cores. We envision that it could be based on an open- source RISC-V core.

[1] [1]The other cores in the cluster can be standard cores.

- scheduling strategies for the memory accesses performed by the control core: these strategies must allow a precise analysis of the task response times (including their execution time on their core and interactions with the control core)

- hardware design of a timing-predictable network-on-chip that allows communication and synchronization between clusters.

- analysis of intra- and inter-cluster communication to verify timing predictability

References

[1] Rodolfo Pellizzoni, Emiliano Betti, Stanley Bak, Gang Yao, John Criswell, Marco Caccamo, and Russell Kegley. A predictable execution model for cots-based embedded systems. In 17th IEEE Real-Time and Embedded Technology and Applications Symposium, 2011.

[2] Paul Pop, Petru Eles, and Zebo Peng. Bus access optimization for distributed embedded systems based on schedulability analysis. In Conference on Design, Automation and Test in Europe (DATE), 2000.

Qualifications

The candidate should have a master's degree in the field of computer science or electronics and information technology.
Good skills in digital circuit design with VHDL/SystemVerilog and in computer architecture are required. Experience with C programming and compilation is desirable. Knowledge in the field of real-time scheduling would be a plus.
Good speaking and writing skills in English are required.

Application

The position is to be filled in Sept.-Oct. 2023. The employer is Continental Automotive. The work will be done partly in the premises of Continental Automotive in Toulouse and partly in the research laboratory IRIT of the University of Toulouse III (www.irit.fr).

Interested candidates are asked to send a motivation letter and a curriculum vitae by e-mail to:
philippe.cuenot@continental-corporation.com and christine.rochange@irit.fr