Semi-partitioned scheduling and task migration in dataflow networks

Cannella, E.

STELLINGEN

Propositions belonging to the Ph.D. dissertation:

Semi-partitioned Scheduling and Task Migration in Dataflow Networks

by Emanuele Cannella

1. Let P be an application modeled as a Polyhedral Process Network. Using a dedicated middleware, it is possible to guarantee correct communication among processes of P even if these processes are remapped at run-time. (This dissertation, Chapter 3)
2. Let P be an application modeled as a Polyhedral Process Network. The code used to run each process of P to a processor can be generated in an automated way. Moreover, the code of processes can also be automatically instrumented to allow process migrations at run-time. (This dissertation, Chapter 4)
3. Let C be an application modeled as a Cyclo-Static Dataflow graph. This application can be scheduled as a set of real time periodic tasks using any scheduling algorithm that guarantees bounded deadline tardiness for each task. (This dissertation, Chapter 5)
4. Let C be an application modeled as a Cyclo-Static Dataflow graph. By scheduling C using a semi-partitioned scheduling approach instead of a purely partitioned one, designers can achieve the same application throughput while requiring less processors, at the cost of higher application latency and memory requirements. (This dissertation, Chapter 5)
5. Let S be an application modeled as an Synchronous Dataflow graph. If S contains stateless actors, the Earliest Deadline First semi-partitioned stateless (EDF-ssl) scheduling algorithm can be used to achieve the same application throughput, compared to partitioned algorithms, with a lower energy consumption. The energy savings come at the cost of increased memory requirements and latency of the application. (This dissertation, Chapter 6)
6. In embedded systems research, papers with a strong mathematical foundation should not be automatically considered valuable. The applicability of the proposed techniques to real-life problems should be evaluated first.
7. In embedded systems research, a paper with a strong connection to real-life problems, but only incremental contributions, is more valuable than a revolutionary paper with loose connection to real-life problems.
8. Reading research papers is a must for any PhD student. However, sometimes, a ten-minute talk with a colleague or a supervisor can be more enlightening than several hours spent studying research papers.
9. Research proposals in the field of embedded systems should always be linked to industry. With respect to that, limitations on research funding from other institutions can even be considered a godsend.