

Autor: Marcelo Grandi Mandelli

Título: Exploration Of Runtime Distributed Mapping Techniques for Emerging Large Scale MPSoCs

Ingresso: 2011/1

Defesa: 13/07/2015, ata nr. 055 **Homologação:** xx/xx/xxxx ata nr./

Linha de Pesquisa: Sistemas Embarcados e Sistemas Digitais

Orientador: Prof. Dr. Fernando Gehm Moraes

Comissão Examinadora: César Augusto Missio Marcon (PPGCC/PUCRS), Guy Gogniat (Université de Bretagne-Sud), José Luís Güntzel (UFSC), Michel Robert (Université de Montpellier – França), Leandro Indrusiak (University of York – Reino Unido), Luciano Copello Ost (Coorientador – University of York – Reino Unido), Giles Sassatelli (Université de Montpellier – França)

ABSTRACT:

MPSoCs with hundreds of cores are already available in the market. According to the ITRS roadmap, such systems will integrate thousands of cores by the end of the decade. The definition of where each task will execute in the system is a major issue in the MPSoC design. In the literature, this issue is defined as task mapping. The growth in the number of cores increases the complexity of the task mapping. The main concerns in task mapping in large systems include: (i) scalability; (ii) dynamic workload; and (iii) reliability. It is necessary to distribute the mapping decision across the system to ensure scalability. The workload of emerging large MPSoCs may be dynamic, i.e., new applications may start at any moment, leading to different mapping scenarios. Therefore, it is necessary to execute the mapping process at runtime to support a dynamic workload. Reliability is tightly connected to the system workload distribution. Load imbalance may generate hotspots zones and consequently thermal implications, which may result in unreliable system operation. In large scale MPSoCs, reliability issues get worse since the growing number of cores on the same die increases power densities and, consequently, the system temperature. The literature presents different task mapping techniques to improve system reliability. However, such approaches use a centralized mapping approach, which are not scalable. To address these three challenges, the main goal of this Thesis is to propose and

evaluate distributed mapping heuristics, executed at runtime, ensuring scalability and a fair workload distribution. Distributing the workload and the traffic inside the NoC increases the system reliability in long term, due to the minimization of hotspot regions. To enable the design space exploration of large MPSoCs the first contribution of the Thesis lies in a multi-level modeling framework, which supports different models and debugging capabilities that enrich and facilitate the design of MPSoCs. The simulation of lower level models (e.g. RTL) generates performance parameters used to calibrate abstract models (e.g. untimed models). The abstract models pave the way to explore mapping heuristics in large systems. Most mapping techniques focus on optimizing communication volume in the NoC, which may compromise reliability due to overload processors. On the other hand, a heuristic optimizing only the workload distribution may overload NoC links, compromising its reliability. The second significant contribution of the Thesis is the proposition of dynamic and distributed mapping heuristics, making a tradeoff between communication volume (NoC links) and workload distribution (CPU usage). Results related to execution time, communication volume, energy consumption, power traces and temperature distribution in large MPSoCs (144 processors) confirm the tradeoff hypothesis. Trading off workload and communication volume improves system reliability through the reduction of hotspots regions, without compromising system performance.