Microprocessors and Microsystems 000 (2016) 1–12

Contents lists available at ScienceDirect

Microprocessors and Microsystems

journal homepage: www.elsevier.com/locate/micpro

Temperature-aware multi-application mapping on network-on-chip based many-core systems

Shan Cao[a,*], Zoran Salcic[b], Zhaolin Li[c], Shaojun Wei[d], Yingtao Ding[a,*]

[a] Institute of Microelectronic Technology, School of Information and Electronics, Beijing Institute of Technology, China
[b] Department of Electrical and Computer Engineering, University of Auckland, New Zealand
[c] Tsinghua National Laboratory for Information Science and Technology, Research Institute of Information Technology, Tsinghua University, China
[d] Tsinghua National Laboratory for Information Science and Technology, Institute of Microelectronics, Tsinghua University, China

ARTICLE INFO

Article history:
Received 16 October 2015
Revised 11 March 2016
Accepted 12 March 2016
Available online xxx

Keywords:
Temperature-aware
Many-core
Thermal coupling
Task mapping
Network-on-Chip

ABSTRACT

As the number of processor cores increases and the core size shrinks, chip temperature has become a critical design issue for Network-on-Chip (NoC) based many-core systems. However, few task mapping methods consider temperature optimization, and even fewer in multi-application scenarios. In this paper, we propose a temperature-aware multi-application mapping scheme for NoC-based many-core systems. Chip temperature is reduced by balancing workloads among cores while few communication overheads are introduced. The applications are firstly partitioned into subgraphs, and the inter-subgraph communication is minimized. A submesh region is then allocated to each subgraph so as to avoid long distance communication. The processing speed of each core is estimated to decide the submesh size according to current temperature distribution, so the submesh allocation is efficient for balanced core temperatures. Finally, task mapping within each pair of subgraph and submesh is performed. The core temperatures, the heating of adjacent cores, chip layout and communication overheads are all considered. Besides, the voltage and frequency level of each task is also scaled down for temperature reduction while timing constraints are guaranteed. Experimental results show that the peak temperature is reduced by 4.8°C compared to current approaches and the temperature variance is reduced to 42% of the initial temperature variance. The energy consumption and communication overheads are also reduced.

1. Introduction

1.1. Background

NoC-based many-core systems are becoming widely popular and used, driven by the demands of the applications, lower cost and power conservation. The ever increasing power consumption and transistor density create severe thermal problem for NoC-based many-core systems. However, limited chip area and power restrict the ability of cooling techniques for addressing thermal concerns.

As feature size shrinks, both core size and distances between cores are decreased. Thermal coupling between neighbouring cores dramatically increases and it is predicted it may increase by up to 65% in the 15nm feature size [1] in the steady state. The situation becomes even worse in many-core systems, where the heating by other cores appears much faster and grows quickly due to the decreased core spacing. High temperatures have several significant impacts on chips. For example, the processing speed is slowed down due to the degraded carrier mobility and increased interconnect resistivity; the device life time is shortened and package reliability is reduced because multiple failure mechanisms are accelerated as temperature increase [2]. Moreover, leakage current increases exponentially with threshold voltage, which depends strongly on temperature. Thus the elevated temperature increases the energy consumption, and this again results in the rise of chip temperature.

In recent years, thermal management techniques have got wide attention, focusing on temperature monitoring [3,4], floorplanning [5,6], cooling techniques [7,8], microarchitectural techniques [9,10], and compiler-assisted techniques [11,12]. Compiler-assisted techniques, such as task mapping, can be more advantageous than others because they can reduce associated hardware overheads [13]. However, most current task mapping schemes for NoC systems aim at performance improvement and power reduction and result in the increase of chip temperature. For example, a large group of mapping schemes map tasks to a concentrated area to reduce delay and energy consumption induced by communication

* Corresponding author. Tel.: +86-10-68918625; fax: +86-10-68918625.
E-mail addresses: caoshan@bit.edu.cn (S. Cao), ytd@bit.edu.cn (Y. Ding).

http://dx.doi.org/10.1016/j.micpro.2016.03.010
0141-9331/

[14–16]. As a result, thermal coupling is exacerbated in the concentrated area, and local temperature is elevated. Other mapping schemes reduce power consumption by low power techniques [17,18], such as using Dynamic Voltage and Frequency Scaling (DVFS) to reduce the processing speed. However, chip peak temperature does not decrease necessarily with the reduction of power consumption because temperature is related to power density, not only power consumption [19]. Besides, these methods usually seek to achieve global power reduction instead of reducing local power density.

The temperature-aware task mapping schemes are studied in recent years, and two issues remain to be addressed. Firstly, tasks are generally mapped evenly to many cores [20,21] in temperature-aware methods to alleviate thermal problems, whereas communication overheads may be significantly elevated, resulting in increased latency, power consumption and even deadlock. Therefore, it is necessary to make trade-offs between hop distance reduction and workloads distribution. Secondly, multiple levels of parallelism are exploited on NoC-based many-core systems, not only at the instruction and task levels, but also at the higher application level. To realize the application level parallelism, design focus is shift to the multi-application scenarios [22,23] where different applications are executed on different submesh regions of NoC simultaneously. Howevere, few submesh allocation approaches in multi-application scenarios optimize the core temperatures. When temperature is taken as a critical concern, the submesh allocation of each application may vary. For example, cores with higher initial temperature have to work at lower speed so as to reduce chip peak temperature, and vice verse.

1.2. Proposed approach overview

In this paper, we propose a multi-application task mapping scheme of NoC-based many-core systems, aiming to reduce the peak temperature while inducing as small communication overheads as possible. To this end, temperature is balanced among cores by estimating processing speed and evaluating the final temperature of each core; a submesh region is allocated to each task set so as to restrict the excessive increase of communication cost. It is implemented in the following three steps. Firstly, applications are represented by task graphs, and the task graphs with longer execution time and more tasks are partitioned into subgraphs. The inter-subgraph communication is minimized so that long-distance communication is not induced. Then, each subgraph is allocated to a submesh region by estimating the speed of each core according to the current temperature distribution. Communicating subgraphs may have shared cores at the border to reduce communication overheads further. Finally, task mapping is performed within each pair of subgraph and submesh, and a lowest possible voltage and frequency (V/F) pair is selected for each task as long as the timing constraints are satisfied. The current temperature distribution, chip layout, thermal coupling, and communication overheads are all considered for the mapping.

To evaluate the efficiency of the proposed mapping scheme, both random task graphs and real applications are simulated. A NoC-based many-core architecture with various NoC sizes are based on mesh topology. Temperature, energy consumption and performance are all simulated. The peak temperature is reduced by 4.8°C on average compared to the method without thermal consideration. The temperature variance is reduced to 42% of the original temperature variance. Besides, energy consumption and the communication cost are also reduced compared to current approaches.

The rest of the paper is organized as follows. Section 2 reviews the related work. The architecture, application and temperature models are built in Section 3, and the target problem to be addressed in this paper is described as well. Section 4 describes the proposed temperature-aware multi-application mapping flow and the proposed algorithms. Experimental results are presented and analyzed in Section 5. Finally, Section 6 summarizes the paper.

2. Related work

Based on the NoC-based many-core systems, a large number of task mapping methods have been proposed in the last few years. Most of them optimize the delay and power consumption induced by communication. Zhou et al. [14] proposed a genetic algorithm based delay model for NoC mapping to minimize average delay. Singh et al. [15] presented a run-time heuristic on NoC-based MPSoC to map tasks in close proximity to minimize communication overheads and alleviate NoC congestion. Hu et al. [16] proposed an online scheduling method to reduce the communication energy consumption on NoC by runtime communication analysis and remapping. Many other methods of task mapping on NoCs focus on power and energy reduction by adjusting processing voltage and frequency. Yang et al. [17] presented a contention-aware energy management scheme to map tasks to Voltage-Frequency Island (VFI)-based NoC systems by both static and dynamic approaches. Das et al. [18] proposed a joint method to optimize the throughput, computation energy and communication energy of heterogeneous Multi-Processor SoC (MPSoC). However, the concentration of tasks and ignorance of thermal profiles may actually increase the localized workloads and elevate core temperatures.

Several task mapping approaches for temperature reduction are proposed recently as the thermal problem becomes severe in multi-core and many-core systems. Murali et al. [24] considered the performance and thermal trade-off on MPSoC by frequency assignment. Performance is maximized by convex optimization while temperature and power constraints are met. For 3-Dimensional MPSoCs, Cox et al. [25] extracted a 3D IC thermal model and presented a thermal-aware resource allocation approach to map throughput-constrained streaming applications. Cheng et al. [26] also solved the task allocation issue on 3D MPSoCs. A tradeoff is made between temperature and interconnect energy by proposing a three-step heuristic. Das et al. [27] proposed a task scheduling method for multiprocessor system to minimize the computation and communication energy, processor aging and chip temperature by task mapping and voltage-frequency scaling. However, initial core temperatures are not considered and long distance communication still exists.

For the mapping of applications, the submesh allocation method is typically used to explore application level parallelism and to avoid long distance communication on NoCs with mesh topology. Yang et al. [22] proposed a mapping method by finding a region for each application and then mapping tasks to the corresponding region. However, thermal concerns are not addressed in this work. Liao et al. [28] presented a virtual submesh allocation scheme for NoC-based many-core systems by dynamically building virtual submeshes considering core temperatures to cool down overheated cores. However, virtual submesh is built which increases communication overheads, and the task mapping is not solved. Also, the size of resulting submesh in this work actually may be affected by current core temperature distribution.

Thus in this paper, we propose a task mapping scheme aiming at peak temperature reduction on NoC-based many-core systems. In multi-application scenarios, the submesh allocation is performed by taking fully consideration of current core temperatures. Task mapping for temperature reduction is augmented with scaling-down V/F level on individual cores. While balancing core temperatures, communication overheads are reduced as much as possible, and the timing constraints of the application are guaranteed.

Please cite this article as: S. Cao et al., Temperature-aware multi-application mapping on network-on-chip based many-core systems, Microprocessors and Microsystems (2016), http://dx.doi.org/10.1016/j.micpro.2016.03.010