Microelectronics Journal 54 (2016) 23–31

Contents lists available at ScienceDirect

Microelectronics Journal

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

Statistical energy optimization on voltage–frequency island based MPSoCs in the presence of process variations

Song Jin [a], Yinhe Han [b], Songwei Pei [c,*]

[a] Department of Electronic and Communication Engineering, School of Electrical and Electronic Engineering, North China Electric Power University, PR China
[b] Key Laboratory of Computer System and Architecture, Institute of Computing Technology, Chinese Academy of Sciences, PR China
[c] Department of Computer Science and Technology, Beijing University of Chemical Technology, PR China

ARTICLE INFO

Article history:
Received 6 November 2014
Received in revised form
5 April 2016
Accepted 3 May 2016

Keywords:
System energy
Process variation
Voltage–frequency island
Multiprocessor system-on-chip

ABSTRACT

Energy efficiency has become a primary design concern for embedded multiprocessor system-on-chips (MPSoCs). Recently, voltage–frequency island (VFI)-based design paradigm was proposed to optimize system energy by combining with task scheduling. However, the ever-increasing variations cause large uncertainty for delay and power of the processors, resulting in performance parameters of the VFIs also manifesting as statistical distribution. As a result, it is more difficult for the deterministic energy optimization to achieve desirable performance yield, defined as the probability of the design meeting timing constraints of the system. In this paper, we propose a variation-aware statistical energy optimization framework which takes account of performance yield constraints in the overall optimization flow. We observed that statistical features of parameter distributions in homogeneous and heterogeneous platforms are different. We hence define two effective metrics, namely energy optimization sensitivity (*EOS*) and lowest operating voltage (*LOV*), to make our framework adapt to both of the platforms. Experimental results demonstrate that on average our framework achieves performance yield improvement of 45% than the deterministic scheme. In respect of energy optimization, on average our framework achieves energy reduction of 33% than the existing statistical task scheduling algorithm.

1. Introduction

Technology scaling dramatically increases integration capacity of the transistors, enabling design of embedded systems to transform into multiprocessor system-on-chip (MPSoC) [1]. Typically, MPSoC can provide an entire system function for high-performance applications by integrating various processing elements (PEs) into a single chip, including microprocessor, DSP and memory, etc. For the embedded MPSoCs powered by the battery, energy efficiency has become a primary design concern [2]. It is desirable to execute applications with minimum energy, meanwhile still meeting real-time constraints of the system.

Recently, the concept of voltage–frequency island (VFI) was introduced into the MPSoC design for effective power management [3,4]. VFI-based design partitions the PEs into multiple voltage islands. Each island can operate at its own voltage and frequency. Combining with task scheduling, such design paradigm manifests great potential in system energy optimization. Consequently, various energy optimization schemes targeting VFI-based MPSoCs were proposed [5–10]. Generally speaking, the existing schemes first perform an energy-aware deterministic task scheduling algorithm. Then the static voltage assigning process determines the allowable lowest supply voltages of the PEs to minimize system energy. Lastly, the VFI partitioning conducts to merge adjacent PEs and unite their supply voltages. The overall optimization flow aims to execute tasks with minimum energy while still meeting the deadline constraints.

However, process variation (PV) exacerbated with the relentless scaling of transistor's feature size [11], bringing significant challenges to the deterministic energy optimization solutions. With PV effects, frequency and power of the PE always deviate from the designated values after chip fabrication. Regarding such uncertainty, performance parameters of the partitioned VFIs should be treated as random variables, making execution time and power of the scheduled task manifest as statistical distribution. As a result, any energy optimization schemes should regard meeting deadline constraint of the executed tasks as the prerequisite. Unfortunately, under serious variations, deterministic energy optimization relied on nominal value cannot guarantee to meet the predefined deadline constraints. While worst case design is costly

* Corresponding author.
E-mail addresses: jinsong@ncepu.edu.cn (S. Jin), yinhes@ict.ac.cn (Y. Han), peisw@mail.buct.edu.cn (S. Pei).

http://dx.doi.org/10.1016/j.mejo.2016.05.004
0026-2692/

and may not be a viable choice. Obviously, statistical nature in performance parameters under PV effects should be taken into account in the VFI-based energy optimizations.

In this paper, we propose a variation-aware statistical energy optimization framework targeting VFI-based MPSoCs. Unlike prior work, we conduct energy-aware task scheduling, voltage assigning and VFI partitioning in a statistical manner by considering probability nature in execution time and power of the task. To guarantee meeting deadline constraints in different process corners, performance yield [12] is used as a primary constraint throughout the optimization flow, defined as the probability of the design meeting timing constraints of the system. Moreover, we found that parameter distribution (frequency or power) of the PEs in homogeneous platform are similar, which is different from the case in heterogeneous one. We hence define two metrics to guide the optimization flow, called energy optimization sensitivity (*EOS*) and lowest operating voltage (*LOV*). These two metrics not only make our framework adapt to both of homogeneous and heterogeneous platforms, but also help us to explore more aggressive optimization space. Compared with the deterministic solution, on average our framework can improve performance yield of 45% over the nominal case or reduce energy of 39% over the worst case. Under the same performance yield constraints, on average energy reduction of 33% is achieved by our framework, compared with the statistical task scheduling algorithm.

The rest of the paper is as follows. Section 2 introduces related work. Section 3 presents the preliminaries. Section 4 details our framework. Section 5 presents the experimental results. We conclude in Section 6.

2. Related work

The large amount of literature has been conducted on energy optimization targeting VFI-based hardware platform. Majzoub et al. [5] proposed voltage island formation scheme to minimize system energy in the presence of PVT variations. Orgas et al. [6] proposed an energy-aware VFI partitioning scheme which performs deterministic task scheduling (EAS) [7] and VFI partitioning in a unified way. Zhang et al. [8] formulated task scheduling and voltage selection as an integer linear programming problem (ILP) to optimize system energy for MPSoCs. Jang et al. [9] proposed an energy optimization framework by combining VFI partitioning with NoC mapping. However, under serious variations, these deterministic schemes cannot guarantee to achieve desirable performance yield. Marculescu et al. [10] analyzed deterministic bounds of application completion times for single and multiple VFI-based platforms. Nevertheless, their work concentrated on timing yield analysis rather than energy optimization.

Considering PV-induced uncertainty, statistical task scheduling was widely adopted to meet timing constraints of the system. Wang et al. [12] proposed performance yield to quantify the scheduling effectiveness, defined as the probability that a task schedule meets the predefined deadline constraints. Huang et al. [13] proposed a variation-aware quasi-static task scheduling algorithm which relies on Monte Carlo simulation to calculate the performance yield. Singhal et al. [14] proposed a non-probabilistic approach to reduce computation complexity of performance yield. Chon et al. [15] proposed a task allocation and scheduling method which takes the impact of resource sharing into consideration. Above mentioned literature, however, only focused on improving performance yield whereas ignoring energy optimization. In [16], Ghorbani proposed the concept of energy yield and formulated yield-aware task scheduling as an ILP problem. The work in [16] does not fit to VFI-based platform. Moreover, rationality of energy yield remains as an open problem since it is desirable for an embedded system to consume energy as low as possible. Our prior work [17] proposed a variation-aware statistical energy optimization flow targeting heterogeneous MPSoCs. However, it cannot be applied to homogeneous platform directly. The reason mainly lies on the different features in parameter distributions of the PEs between such two kinds of platforms, which will be discussed below.

3. Preliminaries

3.1. Target hardware platform and applications

Without loss of generality, the target MPSoC platform is the tile-based structure, as shown in Fig. 1(a). Each tile contains a PE and a router. Inter-tile communication is supported by networks-on-chip (NoC). The MPSoC platform can be homogeneous or heterogeneous. An homogeneous MPSoC is the one in which the PEs contained are all of the same kind, an example production being TILE64 [25]. On the contrary, an heterogeneous MPSoC utilizes a mix of CPU, DSP, VLIW, and function specific HW cores, an example being the Nexperia platform [18]. As many embedded SoCs [3,4], PEs in the target platforms are assumed to have voltage scaling capability which possesses five discrete voltage/frequency levels. The mixed-voltage/mixed-frequency FIFOs located at the VFI boundary are used for data synchronization.

As for the target applications, we focus on high deterministic applications described as the communication task graphs. As shown in Fig. 1(b), the task graph is a directed acyclic graph (DAG). Each vertex denotes a task. Arcs denote control and data dependencies between the task and its preceding tasks. For all the leaf nodes, in general, there are deadline constraints which denote latest finish time of the tasks (e.g., $T_d(3)$ and $T_d(5)$ in Fig. 1(b)).

3.2. Energy model

Given the target platform, the total system energy consumed on executing task graph can be divided into computation energy E_{comp} and communication energy E_{comm}. For executing a task graph containing n tasks, the total computation energy consumed can be expressed as:

$$E_{comp} = \sum_{i=1}^{n} NC_j \times C_i \times V_i^2 \quad (1)$$

where NC_j denotes execution cycles of task j, which is independent of voltage and frequency of the PE. C_i denotes the average switching capacitance per cycle of PE i that task j executed on. V_i denotes the voltage that PE i operated on.

For calculating the total communication energy, as in [6,27], bit energy consumed on transmitting one bit between two PEs is first defined as:

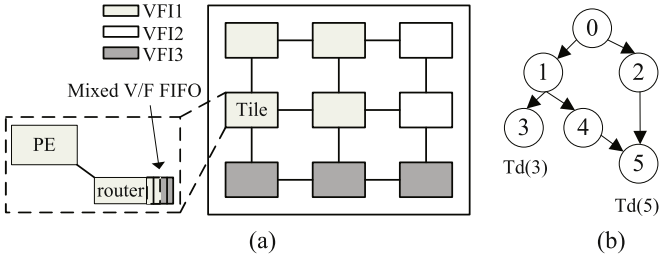

Fig. 1. (a) Target MPSoC platform, and (b) task graph.