



# Energy efficient scheduling of real-time tasks on multi-core processors with voltage islands



Jun Liu, Jinhua Guo\*

Department of Computer and Information Science, University of Michigan at Dearborn, USA

## HIGHLIGHTS

- A Voltage Island Largest Capacity First (VILCF) algorithm is proposed.
- VILCF fully utilizes the capacity of voltage islands to achieve energy efficiency.
- The energy consumption is optimal when the workloads of all cores are balanced.
- The approximation ratio is bounded by a value depending on the number of islands.

## ARTICLE INFO

### Article history:

Received 15 January 2015

Received in revised form

8 May 2015

Accepted 5 June 2015

Available online 26 June 2015

### Keywords:

Voltage island

Dynamic voltage frequency scaling

Real-time

Task scheduling

Energy-efficiency

## ABSTRACT

This paper studies energy efficient scheduling of periodic real-time tasks on multi-core processors with voltage islands, in which cores are partitioned into multiple blocks (termed voltage islands) and each block has its own power source to supply voltage. Cores in the same block always operate at the same voltage level, but can be adjusted by using Dynamic Voltage and Frequency Scaling (DVFS). We propose a Voltage Island Largest Capacity First (VILCF) algorithm for energy efficient scheduling of periodic real-time tasks on multi-core processors. It achieves better energy efficiency by fully utilizing the remaining capacity of an island before turning on more islands or increasing the voltage level of the current active islands. We provide detailed theoretical analysis of the approximation ratio of the proposed VILCF algorithm in terms of energy efficiency. In addition, our experimental results show that VILCF significantly outperforms the existing algorithms when there are multiple cores in a voltage island.

© 2015 Elsevier B.V. All rights reserved.

## 1. Introduction

As the rate of the clock speed improvements of a single processor slowed, manufacturers shifted to add more independent processors (known as cores) into a single chip, which is called Multi-core Processor. Multi-core processors can achieve higher throughput with the same clock frequency. In addition, the proximity of multiple cores on the same chip allows the cache coherency circuitry to operate at a much higher frequency. However, power management becomes more challenging for real-time systems using multi-core processors.

To reduce overall power consumption in multi-core processors, the use of multiple supply voltages has been widely adopted. Often, this approach is realized through the use of core based voltage-islands [1]. A voltage island is a group of on-chip cores powered by the same voltage source, independently from the chip-level

voltage supply. More voltage islands and fine-grained control of shutdown mechanism lead to lower power consumption. One core per voltage island is the most energy efficient design. However, the more voltage islands a chip has, the greater difficulty in routing the power lines. A trade-off is to have multiple cores in a voltage island. For example, Intel SCC (Single-chip Cloud Computer) 48-core x86 processor has 6 different voltage islands, and each island of 8 cores can run at a different voltage.

Power consumption of chips generally breaks down into two sources [2], dynamic power and leakage power. The dynamic power of a core can be minimized by executing at low frequency using Dynamic Voltage and Frequency Scaling (DVFS) technique [3,4]. In general, if the workloads across all cores are balanced, it would allow each core to work at relatively low frequency and thus consume less dynamic power. Leakage power mainly comes from leakage current. When all cores in an island are idle, the island can be powered off using dynamic power management (DPM) [5] to reduce the amount of leakage.

As the demand for concurrent processing and increased energy efficiency grows, it is expected that multi-core processors will

\* Corresponding author.

E-mail addresses: [juliu@umich.edu](mailto:juliu@umich.edu) (J. Liu), [jinhua@umich.edu](mailto:jinhua@umich.edu) (J. Guo).

become widely used in real-time or embedded systems. In this paper, we study the periodic real-time task scheduling problem on multi-core processors with voltage islands. Cores in a voltage island always operate at the same voltage level. However, different voltage islands can have different voltages. This model is more realistic and reflects the reality of the current multi-core architecture. Our *main contributions* are as follows:

- We propose a Voltage Island Largest Capacity First (VILCF) algorithm for energy efficient scheduling of periodic real-time tasks on multi-core processors. It achieves better energy efficiency by fully utilizing the remaining capacity of an island before turning on more islands or increasing the voltage/frequency of the current islands.
- We derive the lower bound of the energy consumption of the optimal solutions for Voltage Island Energy Efficient Scheduling (VIEES) problem. The energy consumption is optimal when the workloads of all cores in a voltage island are balanced. However, the optimal solution is NP-hard to compute.
- We provide detailed theoretical analysis of the approximation ratio of the proposed VILCF algorithm, in terms of energy efficiency, against the lower bound of the optimal solution. The analysis shows that the approximation ratio is bounded to a value depending on the number of voltage islands.
- We perform experiments to evaluate VILCF. Our simulation results show that VILCF significantly outperforms the existing algorithms when there are multiple cores in a voltage island.

The rest of this paper is organized as follows. The closely related work is reviewed in Section 2. Section 3 describes the power, task, and Voltage Islands models. The problem definition and Voltage Island Largest Capacity First (VILCF) algorithm is presented in Section 4. The lower bound of the energy consumption of the optimal solution is discussed in Section 5. Section 6 presents a thorough analysis of the approximation ratio of the proposed VILCF algorithm in term of energy efficiency. The performance evaluation results are discussed in Section 7. Finally, we conclude the paper and describe future work in Section 8.

## 2. Related work

Power-aware computing and energy efficient scheduling of real-time tasks have been well studied in the past decade. Due to the convexity of power consumption function, dynamic voltage and frequency scaling (DVFS) that slows down the processing speed of processors is a commonly used technique for energy savings [3,4,6]. Based on the DVFS technique, many real-time task scheduling schemes have been developed for uniprocessor and multiprocessors, e.g., [7,8]. In our previous works [9,10], we studied power saving for servers and server farms under the task response time constraint. We applied both DVFS and DPM [5] techniques to optimize power consumption of servers while ensuring the response time.

There have been relatively limited studies on scheduling real-time tasks on multi-core processors. Most of them assume per-core DVFS, where each core has its own power supply and can change its voltage and frequency level independently from other cores, e.g., [11–14]. Chen et al. [11] investigated energy-efficient scheduling of periodic real-time tasks over multiple DVFS processors with the leakage power consideration and proposed a polynomial time algorithm, Largest Task First (LTF), to derive task mappings that try to execute at a critical frequency.

Some recent works have explored the voltage islands techniques. However, they focused on either configurations of voltage islands or different applications. Qi and Zhu [15] studied symmetric and asymmetric block configurations for multi-core processors,

which contain the same and different number of cores on each block, respectively. They investigated the block-partitioned core configurations for multi-core processors and evaluated the energy efficiency for different block configurations. Hu et al. [16] minimized power consumption through architecting voltage islands in core-based designs and presenting an algorithm for simultaneous voltage island partitioning, voltage level assignment, and physical-level floor planning. Ozturk et al. [17] observed that at the compile time, the loads across of cores in a chip are imbalanced. They developed algorithms to map the application codes to voltage islands and assign different voltages to different processors to provide energy saving. Santiago and Chen in [18] studied the periodic real-time task scheduling and proposed a Single Frequency Approximation (SFA) scheme that uses a single voltage and frequency for executing, particularly, the lowest voltage and frequency that satisfies the timing constraints in voltage island based multi-core processors. However, they used the original Largest Task First (LTF) [11] to assign real-time tasks onto the cores in the island, which is less efficient when there are multiple cores in an island.

In this paper, we propose a Voltage Island Largest Capacity First (VILCF) algorithm to map periodic real-time tasks on multi-core processors. Not only does it balance the workloads throughout all cores, but also it reduces the number of active islands.

## 3. System models

We study the power model, task model, and voltage island model in this section.

### 3.1. Power model

In this study, we adopt a practical and widely accepted power model [15,16,19]. Power consumption of a core is composed of two parts, leakage power and dynamic power. We denote dynamic power  $P_d$ , which can be adjusted by DVFS.  $P_d$  increases as working speed of core increases. We denote the working speed of cores  $s$ . The function of dynamic power to  $s$  is:

$$P_d(s) = C_{ef} V_{dd}^2 s \quad (1)$$

where  $V_{dd}$  denotes working voltage of cores,  $C_{ef}$  denotes effective switch capacitance,  $s = \kappa \frac{(V_{dd} - V_t)^2}{V_{dd}}$ ,  $V_t$  denotes threshold voltage, and  $\kappa$  denotes hardware-design-specific constant, ( $V_{dd} \geq V_t \geq 0$ ,  $\kappa > 0$ , and  $C_{ef} > 0$ ).

We assume leakage power is a constant. The total power consumption per core is the sum of leakage power and dynamic power:  $P_d(s) + \beta_2$ . We denote leakage power  $\beta_2$ . We normalize the power consumption function as:

$$P(s) = s^3 + \beta \quad (2)$$

where  $s^3$  represents dynamic power, and  $\beta$  represents leakage power, see [11].

From Eq. (2), we know that  $P(s)$  is a convex and non-decreasing function of the core speed  $s$ . We draw the curve of this function in Fig. 1,  $s_{\min}$  and  $s_{\max}$  denote the lower bound and upper bound of  $s$  respectively. From the power to speed curve, we find that there exists a critical speed  $s_0$  at which power consumption is optimal, (in this study, we assume  $s_{\min} \leq s_0 \leq s_{\max}$ ). When cores run at critical speed,  $\frac{P_d(s)}{s}$  is minimum. By solving  $\frac{dP(s)}{ds} = 0$ , we derive  $s_0 = \sqrt[3]{\frac{\beta}{2}}$ , see Fig. 1.

### 3.2. Periodic real-time task model

We focus on a set of periodic independent real-time tasks denoted by  $T = \{\tau_i(\phi_i, p_i, c_i), i = 1 \dots n\}$ . Each task  $\tau_i$  contains a

Download English Version:

<https://daneshyari.com/en/article/424883>

Download Persian Version:

<https://daneshyari.com/article/424883>

[Daneshyari.com](https://daneshyari.com)