Computers and Electrical Engineering 70 (2018) 83–98

Contents lists available at ScienceDirect

Computers and Electrical Engineering

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

An online temperature-aware scheduling technique to avoid thermal emergencies in multiprocessor systems☆

Qaisar Bashir[a], Muhammad Naeem Shehzad[b,*], Muhammad Naeem Awais[b], Sobia Baig[b], Muhammad Ghaffar Dogar[b], Aamir Rashid[c]

[a] University of Lahore, Lahore, Pakistan
[b] COMSATS University Islamabad, Lahore Campus, Pakistan
[c] University of Engineering and Technology Taxila, Pakistan

ARTICLE INFO

Keywords:
Embedded systems
Temperature-aware scheduling
Multiprocessing
Switching time
Simulation

ABSTRACT

Reliability, performance and power consumption of a real-time multiprocessor system are negatively affected by high temperatures, spatial gradients and thermal cycles. Software-based load balancing techniques have been used for thermal management. The transfer of workload in such techniques is based on temperature estimation. In this paper, we proposed an online temperature-aware scheduling technique that performs load balancing based on dynamic temperature measurement at a fixed ambient temperature. Contrary to the static techniques that utilize the principle of temperature prediction, the proposed technique does not require any thermal history of workload and is effective for any kind of workload without prior knowledge. Moreover, it reduces the energy consumption and avoids the workload switching delays among the cores. The simulation results show that the technique reduces overall temperature up to 5%, thermal cycles up to 3% and lowers the temporal and spatial gradients compared to the commonly used techniques.

1. Introduction

There has been an enormous growth in the population of embedded systems in the recent years. With the continuous development of complex software applications, there is a constant demand for improved processing power. Superior processing capability enables processors to handle computationally intense applications in an effective manner. Initially, the focus of the research was to improve the clock speed of a single processor. However, soon it was realized that the speed of a single processor cannot be enhanced beyond a certain limit because of excessive heat dissipation and power consumption. Thereafter, with the advancement in CMOS technology, multiprocessor cores operating at higher frequencies have become a common practice to achieve higher processing capabilities. Although such systems improve performance, however they incur multiple problems with reference reliability and power utilization. The thermal emergencies like high chip temperature, temperature spatial and temporal gradients and thermal cycles are some of the major side effects of multiprocessor systems. Another reason behind the elevated chip temperature is high power density of the chip. This rise is because of the recent progress in fabrication methods, where the size of the chip is continuously scaled down, but the circuit operating voltages are not reduced with the same proportion.

The high chip temperatures challenge the reliability of the system. Mean time to failure (*Mttf*) of a device, that reflects the

☆ Reviews processed and recommended for publication to the Editor-in-Chief by Associate Editor Dr. J. Garcia-Alfaro.

* Corresponding author.
E-mail address: naeem.shehzad@ciitlahore.edu.pk (M.N. Shehzad).

https://doi.org/10.1016/j.compeleceng.2018.06.002
Received 5 January 2018; Received in revised form 1 June 2018; Accepted 3 June 2018
0045-7906/ © 2018 Elsevier Ltd. All rights reserved.

reliability of the system, is adversely affected by high temperature and is expressed as Eq. (1).

$$Mttf = B.\ e^{E_A/k.T_0} \tag{1}$$

where B is the constant, E_A is the activation energy measured in electron volt (ev), T_0 is the operating temperature and k is the Boltzman constant. The Eq. (1) shows that when the operating temperature of the device is increased, its *Mttf* decreases exponentially. This is due to the fact that the elevated temperature significantly increases the chances of dielectric breakdown (DB), stress migration (SM) and electro-migration (EM) which initially causes the degradation of the device operation and subsequently the permanent failure of the device. Similarly, the thermal cycles which are produced due to the large difference in temperature between different parts of chip cause temporal and spatial gradients that in turn reduce the life span of the chip. Eq. (2) expresses the mathematical relation between thermal cycles and failure rate of the device.

$$\gamma \propto \Delta T^Q \tag{2}$$

where γ is the failure rate of the device, ΔT is the change in temperature magnitude, and Q is the frequency of thermal cycles. Eq. (2) shows that the failure rate of the device increases with more frequent and larger temperature swings.

Other than the above discussed thermal emergencies, the ambient temperature also directly affects the temperature of the cores in the system. A research on multiprocessor Core i7 technology reveals that 1 °C increase in ambient temperature leads to 1.05 °C increase in CPU temperature [1]. Moreover, higher chip temperatures also increase the total power consumption of the system. Since most of the embedded systems are battery operated and have very limited cooling options, the negative effect of temperature on power drainage for such systems cannot be ignored in modern processor technology. Expensive packaging solutions and cooling mechanisms are used to mitigate the thermal effects, however these solutions increase the system cost. In short, the temperature as well as the energy and power management for these systems has become a challenging problem in modern processor technology.

Several software-based techniques have been proposed in literature to address the thermal issues, including the average temperature, temperature peaks, thermal cycles and temperature spatial gradient [2]. Few of them are static in nature and switch the workload on the basis of estimated core temperature using offline thermal history of the specific workload. However, these techniques do not provide accurate results, if the run-time workload is different from the ones for which the thermal behaviors are recorded. On the contrary, the workload is switched dynamically on the basis of the current value of core temperature in online techniques. These techniques are efficient in the sense that they do not require any prior knowledge about the nature of the workload. Nevertheless, most of the online techniques do not consider the switching delays between different states of cores (idle, sleep, deepsleep), which results in wastage of time resources. Moreover, most of the existing scheduling based techniques address only few of the many thermal emergencies at a time.

In this research work, we have proposed an online thermal-aware scheduling technique. The main contributions of the work are:

- The technique addresses the thermal emergencies including the average temperature, temperature peaks, thermal cycles and temperature spatial gradient in the presence of fixed ambient temperature based on principle of thermal balancing.
- Contrary to the existing online techniques, it considers the switching delays between different energy states of cores (idle, sleep, deepsleep) and evades them to boost the performance.
- It lowers the energy/power consumption of the system by controlling the thermal issues.

The rest of the paper is organized in the following manner. Section II highlights the related work. Section III describes system model, power model and task model. Section IV describes the proposed temperature-aware online scheduling technique. Section V explains the experimental setup. Section VI consists of results and discussion followed by conclusion drawn in section VII.

2. Related work

Comprehending the seriousness of thermal issues in multiprocessor systems, numerous techniques were introduced previously to manage real-time scheduling by considering thermal constraints. Most of them involve dynamic power management and dynamic voltage and frequency scaling to control the peak temperature, energy, and performance. Some of the well known related techniques are described here.

2.1. Dynamic power management (DPM) techniques

DPM technique turns off the system components which are not in use. DPM techniques include simple predictive, timeout, model-based stochastic, online, session clustering, adaptive (based on learning), and many more [2]. All the existing techniques can be classified into two broader categories (1) predictive techniques (2) stochastic techniques.

Predictive techniques develop the relationship between past and current workload events and predict the future workload event. Based on workload, some components are turned on and others are turned off. Predictive techniques do not work efficiently when there is no prior information of the system workload or it is changing rapidly. Authors in [3] proposed a Predictive Dynamic Thermal Management (PDTM) technique, in which accuracy of the workload prediction is improved by using separate predictor [3].

Stochastic techniques are based on Markov model. In this model, workload uncertainty is considered and a trade-off between power and performance is performed. Different states (idle, sleep, deepsleep, standby) of devices are modeled as stochastic processes [4]. Authors in [5] introduced a technique that accumulates all the idle intervals before the transition into the low energy state and

Download English Version:

https://daneshyari.com/en/article/6883268

Download Persian Version:

https://daneshyari.com/article/6883268

Daneshyari.com