J. Parallel Distrib. Comput. 123 (2019) 13–25

Contents lists available at ScienceDirect

J. Parallel Distrib. Comput.

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

Optimal task execution speed setting and lower bound for delay and energy minimization

Keqin Li

Department of Computer Science, State University of New York, New Paltz, NY 12561, USA

HIGHLIGHTS

- Investigate scheduling a set of independent sequential tasks on identical processors.
- Find the optimal task execution speed setting analytically for delay and energy minimization.
- Establish lower bound for the minimum schedule length with a given energy consumption constraint.
- Establish lower bound for the minimum energy consumption with a given schedule length constraint.
- Perform experimental study on the performance of list scheduling algorithms.

ARTICLE INFO

Article history:
Received 22 November 2017
Received in revised form 7 July 2018
Accepted 3 September 2018
Available online xxxx

Keywords:
Delay minimization
Energy minimization
Energy-efficient task scheduling
List scheduling
Lower bound
Optimal speed setting
Static power consumption

ABSTRACT

The current technology trend reveals that static power consumption is growing at a faster rate than dynamic power consumption. In this paper, energy-efficient task scheduling is studied when static power consumption is a significant part of energy consumption which cannot be ignored. The problems of scheduling a set of independent sequential tasks on identical processors so that the schedule length is minimized for a given energy consumption constraint or the energy consumption is minimized for a given schedule length constraint are investigated. For a given schedule, the optimal task execution speed setting for delay and energy minimization is found analytically. Lower bounds for the minimum schedule length of a set of tasks with a given energy consumption constraint and the minimum energy consumption of a set of tasks with a given schedule length constraint are established. Our lower bounds are applicable to sequential or parallel, and independent or precedence constrained tasks, on processors with discrete or continuous speed levels, and bounded or unbounded speed ranges. The significance of these lower bounds is that they can be used to evaluate the performance of any heuristic algorithms when compared with optimal algorithms. Experimental study on the performance of list scheduling algorithms is performed and it is shown that their performance is very close to the optimal. To the best of the author's knowledge, this is the first paper that provides such analytical results for energy-efficient task scheduling with both dynamic and static power consumptions.

1. Introduction

Reducing processor energy consumption has been a significant research issue in the last two decades, and a huge body of literature has been published [1,29,40]. Processor power consumption includes two components, i.e., dynamic power consumption and static power consumption. It was believed that dynamic power consumption is the dominant part of processor energy consumption, and some research ignored static power consumption [15,39].

However, as transistors become smaller and faster, static power dissipation (i.e., the power due to leakage current in the absence of any switching activity) has become increasingly significant. Because leakage current flows from every transistor that is powered on, with increasing die size and integration, static power will become a significant part of processor power consumption. Static power dissipation is equal to the product of the supply voltage and the leakage current. While the rate of reduction of supply voltage is decreasing, leakage current is increasing exponentially [5]. The current technology trend reveals that static power consumption is growing at a faster rate than dynamic power consumption. Leakage current increases about 7.5 times and leakage power increases about 5.0 times every generation, while active power remains roughly constant [4]. In just a few processor generations, the curves will intersect. Technology scaling is increasing both the absolute and relative contributions of static power dissipation [28]. Static power consumption has noticeable influence on energy consumption and energy-delay product (EDP) [24]. It was demonstrated that if static power consumption is tuned during designing and

E-mail address: lik@newpaltz.edu.

https://doi.org/10.1016/j.jpdc.2018.09.003
0743-7315/

manufacturing, it is possible to save up to 35% reduction in energy consumption and achieve up to 20% improvement in the EDP [23].

One major challenge in the study of energy-efficient task scheduling algorithms is lack of performance analysis and comparison between a heuristic solution and an optimal solution, as traditionally conducted in scheduling theory [8] and other areas of approximation algorithms for NP-hard problems [12]. The main weakness of most existing researches is that they only compare the performance of heuristic algorithms among the algorithms themselves, not with an optimal algorithm [3,10,14,25,30–35,38,41]. Furthermore, there is little analytical result on the worst-case or average-case performance ratio, although some attempt has been made without consideration of static power dissipation [15,16,19,20,26]. This is essentially due to the sophistication of energy-efficient task scheduling algorithms and the apparent lack of the understanding of optimal solutions.

To tackle the above challenge and weakness, one effective approach has been developed in [15,16], i.e., establishing a lower bound for the optimal solution and comparing a heuristic solution with the lower bound. The advantages of a lower bound are two fold. First, it is easy to obtain based on just a few parameters, and thus can be easily incorporated into any scheduler in a real system. Second, we can still assess the performance of a heuristic algorithm when compared with an optimal algorithm even we do not know the optimal solution. If the ratio of a heuristic solution to the lower bound is close to one, the performance of a heuristic algorithm is close to the optimal. Even though a performance ratio cannot be derived, it can be obtained experimentally by simulations or numerically by calculations. This method has been successful in studying the performance of various energy-efficient algorithms in scheduling sequential or parallel tasks, and independent or precedence constrained tasks [15–21]. However, such an effort has been effective only when the static power consumption is ignored.

The motivation of this paper is to make further progress towards this direction when static power consumption is a significant part of energy consumption which cannot be ignored. The main contributions of the paper are as follows.

- We investigate the problems of scheduling a set of independent sequential tasks on identical processors so that the schedule length is minimized for a given energy consumption constraint or the energy consumption is minimized for a given schedule length constraint. In particular, for a given schedule, we are able to find the optimal task execution speed setting analytically for delay and energy minimization.
- We establish lower bounds for the minimum schedule length of a set of tasks with a given energy consumption constraint and the minimum energy consumption of a set of tasks with a given schedule length constraint. Our lower bounds are applicable to sequential or parallel, and independent or precedence constrained tasks, on processors with discrete or continuous speed levels, and bounded or unbounded speed ranges. The significance of these lower bounds is that they can be used to evaluate the performance of any heuristic algorithms when compared with optimal algorithms.
- We perform experimental study on the performance of list scheduling algorithms and show that their performance is very close to the optimal.

To the best of the author's knowledge, this is the first paper that provides such analytical results for energy-efficient task scheduling with both dynamic and static power consumptions. All researchers in this area can benefit from our work in the sense that they can compare the performance of their heuristic algorithms with the lower bounds derived in this paper.

The paper is organized as follows. In Section 2, we present preliminary information, including the power consumption model and problem definitions. In Section 3, we develop lower bound for delay minimization, i.e., minimizing schedule length with energy consumption constraint. In Section 4, we develop lower bound for energy minimization, i.e., minimizing energy consumption with schedule length constraint. In Section 5, we demonstrate experimental data for some heuristic algorithms. In Section 6, we extend our lower bounds to parallel tasks and precedence constrained tasks and other power consumption models. In Section 7, we conclude the paper.

2. Background information

Assume that we are given n independent sequential tasks to be executed on m identical processors. Each task can be executed on any of the m processors. There is no precedence constraint (i.e., dependency) nor communication cost among the tasks. (Note: Extensions of our results to parallel tasks and precedence constrained tasks and other situations are discussed in Section 6.) Let r_i represent the execution requirement (measured in the number of processor cycles or the number of instructions) of task i, where $1 \le i \le n$. The processors can be either computing cores in the same node, or computing cores across different nodes, as long as the cores are homogeneous.

We use the following power consumption model adopted by many researchers [6,9,22,27,36,37]. It is well known that dynamic power consumption p (i.e., the switching component of power) can be accurately modeled by a simple equation, i.e., $p = aCV^2f$, where a is an activity factor, C is the loading capacitance, V is the supply voltage, and f is the clock frequency. In the ideal case, the supply voltage and the clock frequency are related in such a way that $V \propto f^{\phi}$ for some constant $\phi > 0$. The processor execution speed s is usually linearly proportional to the clock frequency, namely, $s \propto f$. For ease of discussion, we will assume that $V = bf^{\phi}$ and $s = cf$, where b and c are some constants. Hence, we know that the dynamic power consumption is $p = aCV^2f = ab^2Cf^{2\phi+1} = (ab^2C/c^{2\phi+1})s^{2\phi+1} = \xi s^{\alpha}$, where $\xi = ab^2C/c^{2\phi+1}$ and $\alpha = 2\phi + 1$. Let p_i represent the dynamic power (measured in watts) consumed to execute task i, which is $p_i = \xi s_i^{\alpha}$, where s_i is the execution speed of task i (measured in GHz or the number of billion instructions per second). Let ψ be the static power consumption (measured in watts). Therefore, the total power consumption is $\xi s_i^{\alpha} + \psi = \xi(s_i^{\alpha} + \psi/\xi)$. Since ξ is a constant which only creates scaling effect, for ease of discussion, we will assume that $\xi = 1$ and simply say that $p = \psi/\xi$ is static power consumption. Hence, the power required to execute task i is $p_i + p = s_i^{\alpha} + p$.

The execution time (measured in seconds) of task i is $t_i = r_i/s_i$. The energy (measured in joule) consumed to execute task i is $e_i = (p_i+p)t_i = (p_i+p)r_i/s_i = r_i(s_i^{\alpha}+p)/s_i = r_i(s_i^{\alpha-1}+p/s_i)$. We observe that

$$\frac{\partial e_i}{\partial s_i} = r_i\left((\alpha-1)s_i^{\alpha-2} - \frac{p}{s_i^2}\right).$$

Hence, when $\partial e_i/\partial s_i = 0$, that is,

$$(\alpha-1)s_i^{\alpha-2} = \frac{p}{s_i^2},$$

which implies that when

$$s_i = s^* = \left(\frac{p}{\alpha-1}\right)^{1/\alpha},$$

e_i has its minimum value of

$$e_i^* = r_i\left((s^*)^{\alpha-1} + \frac{p}{s^*}\right),$$

which is actually

$$e_i^* = r_i p^{1-1/\alpha}\frac{\alpha}{(\alpha-1)^{1-1/\alpha}}.$$