Energy-Aware Scheduling for Real-Time Multiprocessor Systems with Uncertain Task Execution Time

Changjiu Xian
Department of Computer Science
Purdue University
West Lafayette, Indiana
cjx@cs.purdue.edu

Yung-Hsiang Lu
School of Electrical and Computer Engineering
Purdue University
West Lafayette, Indiana
yunglu@purdue.edu

Zhiyuan Li
Department of Computer Science
Purdue University
West Lafayette, Indiana
li@cs.purdue.edu

ABSTRACT

This paper presents an energy-aware method to schedule multiple real-time tasks in multiprocessor systems that support dynamic voltage scaling (DVS). The key difference from existing approaches is that we consider the probabilistic distributions of the tasks' execution time to partition the workload for better energy reduction. We analyze the problem of energy-aware scheduling for multiprocessor with probabilistic workload information and derive its mathematical formulation. As the problem is NP-hard, we present a polynomial-time heuristic method to transform the problem into a probability-based load balancing problem that is then solved with worst-fit decreasing bin-packing heuristic. Simulation results with synthetic, multimedia, and stereovision tasks show that our method saves significantly more energy than existing methods.

Categories and Subject Descriptors

C.4 [**Performance of Systems**]: design studies

Keywords

Dynamic Voltage Scaling, Multiprocessor, probability

1. INTRODUCTION

Energy consumption is an important design issue for battery-operated embedded systems. In these systems, the processor is a major energy consumer. Embedded systems often run tasks with real-time constraints. Since dynamic voltage scaling (DVS) can achieve quadratic energy savings with only linear decrease of the processor's speed, combining DVS with real-time scheduling has been extensively studied [9] [16] [18] [20] [21]. For real-time tasks with uncertain execution time, the worst-case execution time (WCET) must be considered for meeting the tasks' deadlines. Since this paper considers frequency and voltage scaling, we use "execution cycles" instead of execution time to express the workload of a task. The task's execution time is the execution cycles divided by the processor's frequency.

While most studies on DVS and real-time scheduling are for a single processor, today's embedded systems are increasingly based on multiprocessors for higher performance and lower power consumption. Energy-efficient task scheduling on multiprocessors therefore becomes an important issue. Existing studies on multiprocessor real-time scheduling take two major approaches: (a) partitioning— each task is assigned to a particular processor permanently, and (b) dynamic scheduling— a global scheduler selects tasks from a single ready queue to execute and tasks can migrate among processors. Either approach has its advantage and disadvantage [15]. In current multiprocessor systems, the partitioning approach is more common because of its simplicity and ease of implementation [3]. In this paper, we improve the energy efficiency under the partitioning approach.

Previous studies have shown that, with DVS, the total energy consumption of the multiprocessor is minimized when the workload is balanced among the processors [4] [7]. This is because of the convex relationship between the processor speed and the power consumption. These previous studies take the worst-case execution cycles (WCEC) of tasks as the workloads. However, if some tasks demand much fewer execution cycles in most cases than in their worst cases, the workloads may be poorly balanced. Even though the WCEC must be considered to guarantee meeting deadlines, the statistical distribution of the execution cycles should also be considered for better balancing. Such statistical information can be obtained through offline or online profiling [21].

In this paper, we address the problem of energy-aware scheduling for multiprocessor with the information of probabilistic distributions. We consider multiple periodic real-time tasks that execute on a set of identical processors. We derive the mathematical formulation for minimizing the expected total energy consumption while meeting the deadlines of all tasks with earliest deadline first (EDF) scheduling. As this problem is NP-hard, we present a polynomial-time heuristic method. First, we transform the problem to the load-balancing problem based on the cycle distributions of workloads, assuming that the processors have unbounded and continuous range of frequencies to choose from. The worst-fit decreasing bin packing heuristic is used to balance the load. Second, we modify the solution by including the maximum frequency as a new constraint and by assuming bounded discrete frequencies. Our simulation results with synthetic, multimedia, and stereovision task sets show that