Journal of Computer and System Sciences 82 (2016) 174–190

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Power and performance management for parallel computations in clouds and data centers

Keqin Li *

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

ARTICLE INFO

Article history:
Received 24 December 2014
Received in revised form 27 June 2015
Accepted 1 July 2015
Available online 21 July 2015

Keywords:
Cloud computing
Data center
Energy-efficient scheduling
Parallel task
Performance analysis
Precedence constraint
Simulation

ABSTRACT

We address scheduling independent and precedence constrained parallel tasks on multiple homogeneous processors in a data center with dynamically variable voltage and speed as combinatorial optimization problems. We consider the problem of minimizing schedule length with energy consumption constraint and the problem of minimizing energy consumption with schedule length constraint on multiple processors. Our approach is to use level-by-level scheduling algorithms to deal with precedence constraints. We use a simple system partitioning and processor allocation scheme, which always schedules as many parallel tasks as possible for simultaneous execution. We use two heuristic algorithms for scheduling independent parallel tasks in the same level, i.e., SIMPLE and GREEDY. We adopt a two-level energy/time/power allocation scheme, namely, optimal energy/time allocation among levels of tasks and equal power supply to tasks in the same level. Our approach results in significant performance improvement compared with previous algorithms in scheduling independent and precedence constrained parallel tasks.

1. Introduction

1.1. Motivation

Modern supercomputing centers and data centers consume astonishing amount of power and generate huge amount of pollution. For instances, the Tianhe-2 system in the National Super Computer Center in Guangzhou (China), the world's fastest supercomputing system as of November 2014, has power consumption of 17,808 kW [27]. Assuming an electricity charge of 0.15 USD per kWh, the system has electricity cost of more than 23 million USD per year. It is well known that the extra costs of cooling and electrical losses are at about similar or even higher scale. The four Facebook's data centers consumed about 678 million kWh of energy in 2012, costing more than 25 million USD per data center [25]. The mix of fuels consumed to generate the energy used by Facebook consisted of 34% coal, 22% nuclear, 19% renewable, 15% natural gas, and 10% "uncategorized". It is also well known that among 100 units of source energy, only 33 units of electrical energy can be successfully generated and delivered to data centers. Therefore, energy consumption in data centers are actually tripled in terms of natural resources consumption. Furthermore, the amount of carbon emissions associated with Facebook's Prineville, Oregon, data center in 2012 was 104,000 metric tons of greenhouse gases. Sustainable high-performance computing and cloud computing has been a grand challenge for human civilization and an extremely important research direction.

* Fax: +845 257 3996.
E-mail address: lik@newpaltz.edu.

http://dx.doi.org/10.1016/j.jcss.2015.07.001

Traditionally, parallel applications are run on high-performance computers (e.g., supercomputers and clusters), which are very expensive and not available to ordinary people. Cloud computing provides the tools and technologies such as computing resource sharing and virtualization to support data and computing intensive parallel applications with much more affordable prices compared to traditional parallel computing infrastructures [22]. Cloud computing is also the most effective approach to green computing due to resource virtualization, automation software, pay-per-use and self-service, and multitenancy [26]. While traditional parallel computing pursues the goal of high performance, cloud computing provides the opportunity to effective and efficient management of power and performance for parallel computations in green clouds and data centers simultaneously. Such balanced consideration of system performance and energy consumption for large-scale parallel applications is very important, since such applications typically consume significant computing resources, times, and energy. Furthermore, cloud computing has the capability to provide energy-efficient methods and algorithms to deal with the power-performance tradeoff at the data center level, when numerous users submit their service requests at the same time. Energy-efficient cloud computing [10] is the issue to be considered in this paper.

1.2. Related research

Increased energy consumption causes severe economic, ecological, and technical problems. Power conservation is critical in many computation and communication environments and has attracted extensive research activities. Reducing processor energy consumption has been an important and pressing research issue in recent years. There has been increasing interest and importance in developing high-performance and energy-efficient computing systems [17–19]. There exists an explosive body of literature on power-aware computing and communication. The reader is referred to [1,8,9,63–65,77] for comprehensive surveys.

Software techniques for power reduction are supported by a mechanism called *dynamic voltage scaling* [24]. Dynamic power management at the operating system level refers to supply voltage and clock frequency adjustment schemes implemented while tasks are running. These energy conservation techniques explore the opportunities for tuning the energy-delay tradeoff [62]. In a pioneering paper [66], the authors first proposed the approach to energy saving by using fine grain control of CPU speed by an operating system scheduler. In a subsequent work [68], the authors analyzed offline and online algorithms for scheduling tasks with arrival times and deadlines on a uniprocessor computer with minimum energy consumption. These research have been extended in [4–6,12,32,35,47–49,69] and inspired substantial further investigation, much of which focus on real-time applications. In [3,21,23,28,30,34,38,50,52,54,57,58,60,61,67,73–76] and many other related work, the authors addressed the problem of scheduling independent or precedence constrained tasks on uniprocessor or multiprocessor computers where the actual execution time of a task may be less than the estimated worst-case execution time. The main issue is energy reduction by slack time reclamation.

There are two considerations in dealing with the energy-delay tradeoff. On the one hand, in high-performance computing systems, power-aware design techniques and algorithms attempt to maximize performance under certain energy consumption constraints. On the other hand, low-power and energy-efficient design techniques and algorithms aim to minimize energy consumption while still meeting certain performance goals. In [7], the author studied the problems of minimizing the expected execution time given a hard energy budget and minimizing the expected energy expenditure given a hard execution deadline for a single task with randomized execution requirement. In [11], the author considered scheduling jobs with equal requirements on multiprocessors. In [14], the authors investigated the relationship among parallelization, performance, and energy consumption, and the problem of minimizing energy-delay product. In [20], the authors addressed joint minimization of carbon emission and maximization of profit. In [33,37], the authors attempted joint minimization of energy consumption and task execution time. In [46], the authors studied the problem of scheduling a bag-of-tasks application, i.e., a collection of independent stochastic tasks on a heterogeneous platform with deadline and energy consumption budget constraints. In [59], the authors simultaneously addressed three constraints, i.e., energy, deadline, and reward, for both homogeneous and heterogeneous applications. In [71], the authors devised a novel reliability maximization with energy constraint algorithm, which can effectively balance the tradeoff between high reliability and energy consumption. In [78], the authors considered task scheduling on clusters with significant communication costs. System level power management has been investigated in [16,29,36,51,55]. Other studies were reported in [2,53,72].

In [40–45], we addressed energy and time constrained power allocation and task scheduling on multiprocessors with dynamically variable voltage and frequency and speed and power as combinatorial optimization problems. In [40,43], we studied the problems of scheduling independent sequential tasks. In [41,44], we studied the problems of scheduling independent parallel tasks. In [42], we studied the problems of scheduling precedence constrained sequential tasks. In [45], we studied the problems of scheduling precedence constrained parallel tasks.

1.3. Our contributions

In this paper, we address scheduling independent and precedence constrained parallel tasks on multiple homogeneous processors in a data center with dynamically variable voltage and speed as combinatorial optimization problems. In particular, we consider the problem of minimizing schedule length with energy consumption constraint and the problem of minimizing energy consumption with schedule length constraint on multiple processors. The first problem has applications