Applied Soft Computing 24 (2014) 432–446

Contents lists available at ScienceDirect

Applied Soft Computing

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

Review Article

Multi-objective evolutionary algorithms for energy-aware scheduling on distributed computing systems

Mateusz Guzek[a], Johnatan E. Pecero[b,*], Bernabé Dorronsoro[c], Pascal Bouvry[b]

[a] Interdisciplinary Centre for Security Reliability and Trust, University of Luxembourg, Luxembourg
[b] Faculty of Science, Technology, and Communications, University of Luxembourg, Luxembourg
[c] Laboratoire d'Informatique Fondamentale de Lille, University of Lille 1, France

ARTICLE INFO

Article history:
Received 9 June 2012
Received in revised form 15 May 2014
Accepted 4 July 2014
Available online 24 July 2014

Keywords:
Evolutionary algorithms
Multi-objective optimization
Scheduling
Energy efficiency
Dynamic voltage scaling

ABSTRACT

The ongoing increase of energy consumption by IT infrastructures forces data center managers to find innovative ways to improve energy efficiency. The latter is also a focal point for different branches of computer science due to its financial, ecological, political, and technical consequences. One of the answers is given by scheduling combined with dynamic voltage scaling technique to optimize the energy consumption. The way of reasoning is based on the link between current semiconductor technologies and energy state management of processors, where sacrificing the performance can save energy.

This paper is devoted to investigate and solve the multi-objective precedence constrained application scheduling problem on a distributed computing system, and it has two main aims: the creation of general algorithms to solve the problem and the examination of the problem by means of the thorough analysis of the results returned by the algorithms.

The first aim was achieved in two steps: adaptation of state-of-the-art multi-objective evolutionary algorithms by designing new operators and their validation in terms of performance and energy. The second aim was accomplished by performing an extensive number of algorithms executions on a large and diverse benchmark and the further analysis of performance among the proposed algorithms. Finally, the study proves the validity of the proposed method, points out the best-compared multi-objective algorithm schema, and the most important factors for the algorithms performance.

Contents

Introduction 433
Problem description 433
System 433
Constraints 434
Objectives 434
Energy-aware scheduling: related work 434
Algorithms description 435
MOEAs schemas 435
Representation 435
Operators 437
Grouping recombination 437
Mutation 437
Evaluation 437

* Corresponding author. Tel.: +352 4666445733.
E-mail addresses: Mateusz.Guzek@uni.lu (M. Guzek), Johnatan.Pecero@uni.lu (J.E. Pecero), bernabe.dorronsoro_diaz@inria.fr (B. Dorronsoro), Pascal.Bouvry@uni.lu (P. Bouvry).

http://dx.doi.org/10.1016/j.asoc.2014.07.010

Experimentation 438
Problem instances 438
Algorithms testing and configuration 439
Performance metrics 439
Convergence test 440
Performance of the proposed recombination 440
Results 441
Solution number analysis 441
Quality indicators analysis 442
Comparison with other methods 442
Analysis of the results 444
Conclusions and future work 445
Acknowledgements 446
References 446

Introduction

The energy efficiency of Information Technologies (IT) is one of the biggest current issues in the field of computing. The global data center power is estimated as 38.9 GW, and increased by 19% in 2012, and consecutively by 7% in 2013 [1]. The rapid increase of energy consumption of IT infrastructures, caused by growing scale and power of computing systems, resulted in development of a new discipline of IT – called commonly as GreenIT. Researchers and engineers make advances in every aspect of the domain to ensure increasing energy efficiency, with an important distinction between two categories of energy usage optimization, called static and dynamic power management [2]. The difference between them is that static power management takes place during the design time of the IT element on which it is applied. Oppositely, dynamic power management is a technique, which is executed during running of such element. One of the main dynamic management methods to optimize performance in computing systems is to use the best schedulers.

Classically, the optimal schedule is the one that executes the workload minimizing one of the execution time functions, e.g. minimizing total execution time. However, an energy-efficient scheduling algorithm has to minimize consumed energy. Such algorithms exploit power management technologies available in hardware to achieve this aim. As the biggest influence for power consumption of a server is its processor [3], this study focuses on minimizing its energy consumption. The processors manufacturers offer two main technologies: resource hibernation and Dynamic Voltage Frequency Scaling (DVFS) [2,4].

In this work, we focus on DVFS technology that exploits the characteristics of power function of electronic circuits. The most common modern circuit technology, complementary metal-oxide-semiconductor (CMOS), has a convex power function of supplied voltage and frequency. Additionally, frequency in such circuit is linear function of voltage. Therefore, using low voltage levels leads to energy savings. This technique may also result in decreased Quality of Service (hereinafter, QoS), as decreasing processing speed may increase total execution time. However, this technique is much more adaptable to changes than resource hibernation, as DVFS transition time (30–150 μs) is much shorter than hibernating a resource (few seconds, which decreases system responsiveness and thus QoS) [5]. Such elasticity of DVFS enables its direct incorporation into the scheduling process for each task. The only requirement needed to successfully and meaningfully apply that technique is that single task execution times are significantly bigger than DVFS transition time.

The most common state-of-the-art technique using DVFS is called slack reclamation [6]. It is a post-processing algorithm, which takes an already constructed schedule and uses tasks' slack times to reduce performance of processors whenever it is possible without increasing total completion time. It has been adopted in many algorithms as it is easy to apply and gives considerable energy savings (see Section 'Energy-aware scheduling: related work'). The main assumption that leads to the next generation of algorithms using DVFS is that involving it into the scheduling step itself will lead to greater energy savings. This work follows this direction, as used methods allow to apply DVFS at each step of a schedule creation.

The problem of precedence constrained task scheduling is NP-hard in the simplified case of equal length tasks, homogeneous processors, and no communication costs [7]. This work tackles a generalization of the problem, with precedence constraints, heterogeneous processors using DVFS technology, communication costs, and another objective – the energy consumption.

The contributions of this paper are threefold. (1) We propose three scheduling algorithms to solve the heterogeneous multiprocessor multi-objective scheduling problem. They based on state-of-the-art multi-objective (MO) algorithms schemas, with new grouping crossover and mutation operators. (2) We do a thorough empirical study of the problem, evaluating the performance of different operators in the algorithms and the influence of instance parameters on the solutions obtained. Finally, (3) we identify the most important factors of the problem, as well as the best performing algorithm for the problem, with statistical confidence.

The remainder of the paper is structured as follows: in Section 'Problem description' the tackled problem is described. The state of the art on energy-aware scheduling is presented in Section 'Energy-aware Scheduling: Related Work'. Section 'Algorithms description' describes the proposed solution for the problem and Section 'Experimentation' analyses the results of simulations over the large set of instances. Section 'Conclusions and future work' includes conclusion and presents future research directions.

Problem description

The addressed scheduling problem deals with the optimal allocation of a set of tasks that compose a parallel application to the set of processing elements in a distributed system. The target is minimizing both the total execution time and energy consumed by the execution of the application.

System

The distributed computing system consists of a set R of m heterogeneous processors. For each processor r_l, a set of DVFS pairs D^l is defined. A DVFS pair k, denoted as d_k, is a tuple (v_k, s_k), where $v_k \in \mathbb{R}^+$ is the operational voltage of the processor and $s_k \in \mathbb{R}^+$ is the ratio of the operational speed to the maximal processor's speed, further called as relative speed.