Computer Science Review 30 (2018) 31–54

Contents lists available at ScienceDirect

Computer Science Review

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

Review article

Scheduling in distributed systems: A cloud computing perspective☆

Luiz F. Bittencourt [a,*], Alfredo Goldman [b], Edmundo R.M. Madeira [a], Nelson L.S. da Fonseca [a], Rizos Sakellariou [c]

[a] Institute of Computing, University of Campinas, Brazil
[b] Instituto de Matemática e Estatística, Universidade de São Paulo, Brazil
[c] School of Computer Science, University of Manchester, UK

ARTICLE INFO

Article history:
Received 25 April 2018
Received in revised form 15 August 2018
Accepted 17 August 2018

Keywords:
Distributed systems
Scheduling
Cloud computing

ABSTRACT

Scheduling is essentially a decision-making process that enables resource sharing among a number of activities by determining their execution order on the set of available resources. The emergence of distributed systems brought new challenges on scheduling in computer systems, including clusters, grids, and more recently clouds. On the other hand, the plethora of research makes it hard for both newcomers researchers to understand the relationship among different scheduling problems and strategies proposed in the literature, which hampers the identification of new and relevant research avenues. In this paper we introduce a classification of the scheduling problem in distributed systems by presenting a taxonomy that incorporates recent developments, especially those in cloud computing. We review the scheduling literature to corroborate the taxonomy and analyze the interest in different branches of the proposed taxonomy. Finally, we identify relevant future directions in scheduling for distributed systems.

Contents

☆ No author associated with this paper has disclosed any potential or pertinent conflicts which may be perceived to have impending conflict with this work. For full disclosure statements refer to https://doi.org/10.1016/j.cosrev.2018.08.002.

* Corresponding author.
E-mail addresses: bit@ic.unicamp.br (L.F. Bittencourt), gold@ime.usp.br (A. Goldman), edmundo@ic.unicamp.br (E.R.M. Madeira), nfonseca@ic.unicamp.br (N.L.S. da Fonseca), rizos@cs.man.ac.uk (R. Sakellariou).

https://doi.org/10.1016/j.cosrev.2018.08.002
1574-0137/

1. Introduction

The scheduling problem arises in countless areas, and it evolves over time along with industry and technology [1]. With the development of computers, scheduling in computer processors received great attention [2,3], being the most common objective the minimization of task completion times, also known as makespan. Besides some peculiarities, the basic principles remained the same as in scheduling activities among machines in production. In supercomputers, multiprocessor scheduling considers several parallel processors with the same capacity. In addition, the data source is considered to be centralized and connected by a high speed channel between processors, in a way that activities (or jobs) can exchange messages quickly.

More recently, computer networks allowed clusters of homogeneous computers to act as a multiprocessor computer with distributed data sources. However, when compared to supercomputers, clusters initially had a slow communication channel between processors, which made data exchange among processors more expensive. The scheduling of jobs in computing clusters led to another branch of research: the scheduling in distributed computer systems. With improvements in computer networks, the connection among computing nodes in clusters became faster. On the other hand, new applications demanded more and more bandwidth, storing and exchanging massive volumes of data. Multimedia and e-Science are examples of applications that handle large data sets nowadays, putting in evidence the importance of communications to improve performance and support quality of service offering in distributed systems.

Grid computing emerged in the late 90's as a heterogeneous collaborative distributed system [4] evolved from homogeneous distributed computing platforms. Grids are shared systems that enclose potentially any computing device connected to a network, from workstations to clusters. Computing grids are infrastructures that enable resource sharing by establishing use policies as well as security rules, which compose the so called Virtual Organizations (VOs) [4].

Cloud computing offers computing resources, often virtualized, as services to the users, hiding technical aspects regarding resource management [5]. Therefore, clusters and grids can be part of datacenters in the cloud computing infrastructure, demanding new optimization objectives and variables common in green computing [6] and utility computing [7]. Kwok and Ahmad stated in 1999 [3] that considering heterogeneous platforms was a challenging direction to extend scheduling algorithms. As a consequence of the popularization of these platforms in grid and cloud computing, novel scheduling concepts appeared in the literature [8–12]. While on the one hand fundamental scheduling aspects remain unchanged, on the other hand different optimization objectives ballooned the scheduling literature in the past decade. Such swell in the field brought so much information that it became challenging the recognition of the exact contribution of new results. Since challenges in scheduling still exist, Smith argues in [13] that scheduling is not a fully solved problem, he stated that *"Scheduling techniques that properly account for uncertainty, enable controlled solution change, and support efficient negotiation and refinement of constraints are crucial prerequisites, and the need to operate in the context of multiple self-interested agents is a given."*. This statement matches certain characteristics in virtualization and service in cloud computing, as we shall describe in the upcoming sections.

In a nutshell, this paper has three main contributions:

1. Propose a taxonomy for scheduling in distributed systems and introduce a taxonomy extension to cover cloud computing schedulers.
2. Classify the literature in the proposed taxonomy;
3. Identify relevant future directions for scheduling in distributed systems.

Due to its wide application, there exist a variety of approaches to the scheduling problem. This paper presents directives for scheduling researchers to identify and classify their work, as well as to provide them with an overview of existing approaches associated to their research by presenting a broad view of the scheduling problem in distributed systems as well as introducing existing works. We first introduce the problem of scheduling in distributed systems, covering advances in scheduling in cluster and grid computing. Then, an overview of the proposed taxonomy is presented, followed by the state-of-the-art in each branch of the taxonomy tree. After that, we focus on cloud computing and detail similarities and differences of scheduling in clouds with scheduling in previously existing distributed systems. We introduce a taxonomy of scheduling in cloud computing, extending the pre-cloud taxonomy. Finally, we discuss research challenges that were inherited from grid and cluster computing by the cloud computing paradigm, as well as new problems to be tackled.

This paper is organized as follows. Section 2 discusses previous work that have addressed reviews and surveys of scheduling in distributed computing. Section 3 introduces basic concepts to define scheduling problem. Section 4 briefly introduces the whole taxonomy discussed in this paper, highlighting branches that were introduced to cover scheduling in cloud computing. Section 5 presents a taxonomy of schedulers previously to the advent of cloud computing (*pre-cloud* taxonomy). Section 6 discusses and classifies the scheduling taxonomy in cloud computing (*cloud taxonomy*). Section 7 reviews the cloud computing scheduling literature according to the proposed taxonomy. Future directions in scheduling for distributed systems are discussed in Section 8, and Section 9 presents the concluding remarks.

2. Related work

In computer science, with the constant networking and middleware development, scheduling in distributed processing systems is one of the topics which has gained attention in the last two decades. Casavant and Kuhl [14] present a taxonomy of scheduling in general purpose distributed systems. The classification presented by the authors include local and global, static and dynamic, distributed and non-distributed, cooperative and non-cooperative scheduling, as well as some approaches to solve the problem, such as optimal and sub-optimal, heuristic, and approximate. This presented classification is complete in some sense, and it is still valid nowadays. However the current state of distributed systems indeed demands the addition of new branches in this taxonomy.