Computers & Industrial Engineering 63 (2012) 652–662

Contents lists available at SciVerse ScienceDirect

Computers & Industrial Engineering

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

# A column generation approach for power-aware optimization of virtualized heterogeneous server clusters

Hugo H. Kramer [a,*], Vinicius Petrucci [b], Anand Subramanian [c,b], Eduardo Uchoa [a]

[a] Universidade Federal Fluminense, Departamento de Engenharia de Produção, Rua Passo da Pátria, 156 – Bloco E, 4° andar, São Domingos, Niterói-RJ 24210-240, Brazil
[b] Universidade Federal Fluminense, Instituto de Computação, Rua Passo da Pátria, 156 – Bloco E, 3° andar, São Domingos, Niterói-RJ 24210-240, Brazil
[c] Universidade Federal da Paraíba, Departamento de Engenharia de Produção, Centro de Tecnologia, Campus I – Bloco G, Cidade Univesitária, João Pessoa-PB 58051-970, Brazil

ARTICLE INFO

*Article history:*
Available online 12 August 2011

*Keywords:*
Column generation
Energy management
Server virtualization

ABSTRACT

Increasingly, clusters of servers have been deployed in large data centers to support the development and implementation of many kinds of services, having distinct workload demands that vary over time, in a scalable and efficient computing environment. Emerging trends are utility/cloud computing platforms, where many network services, implemented and supported using server virtualization techniques, are hosted on a shared cluster infrastructure of physical servers. The energy consumed to maintain these large server clusters became a very important concern, which in turn, requires major investigation of optimization techniques to improve the energy efficiency of their computing infrastructure.

In this work, we propose an efficient approach to solve a relevant cluster optimization problem which, in practice, can be used as an embedded module to implement an integrated power and performance management solution in a real server cluster. The optimization approach simultaneously deals with (i) CPU power-saving techniques combined with server switching on/off mechanisms, (ii) the case of server heterogeneity, (iii) virtualized server environments, (iv) an efficient optimization method, which is based on column generation techniques. The key aspects of our approach are the basis on rigorous and robust optimization techniques, given by high quality solutions in short amount of processing time, and experimental results on the cluster configuration problem for large-scale heterogeneous server clusters that can make use of virtualization techniques.



## 1. Introduction

An increasing number of clusters of servers have been deployed in large data centers to support the development and implementation of many kinds of services supporting different applications in a scalable and efficient computing environment, for example, focused on Web-based applications. A typical server cluster is a distributed system that consists of hundreds or thousands of machines linked by a fast network (Cardellini, Casalicchio, Colajanni, & Yu, 2002). These cluster architectures are becoming common in utility/cloud computing platforms (Church, Greenberg, & Hamilton, 2008; Hayes, 2008), such as Amazon EC2 and Google AppEngine. In these platforms, the network services are mostly hosted on several shared physical servers and can have distinct workloads that vary over time.

These cluster platforms usually have great processing and performance demands, incurring in high energy costs and indirectly contributing to increase $CO_2$ generation and then to the environmental deterioration (McKinsey & Company, 2008). Accordingly to Koomey (2007), the power consumption of servers and their associate infrastructure in 2005 was about 1.2% of the total consumed in US at an estimated cost of US$ 2.7 billion, and, in the world, this cost was estimated in US$ 7.2 billion. The energy consumed to maintain these large data centers became a very important concern, which in turn, requires major investigation of optimization techniques to improve the energy efficiency of their computing infrastructure (Bianchini & Rajamony, 2004; Fan, Weber, & Barroso, 2007; Ranganathan, 2010).

To support the execution of multiple independent network services, modern server cluster platforms (also known as cloud computing) rely on virtualization techniques to enable the usage of different virtual machines (VMs)—operating system plus software applications—on a single physical server. Server virtualization has been widely adopted in data centers around the world for improving resource usage efficiency; particularly helping to make these computing environments more energy-efficient. Several virtual machine monitors or hypervisors, which act as a layer between the virtual machine and the actual hardware, have been developed to support server virtualization, such as Xen (Dragovic et al., 2003) and VMware (Haletky, 2008).

* Corresponding author.
*E-mail addresses:* hugoharry@gmail.com (H.H. Kramer), vpetrucci@ic.uff.br (V. Petrucci), anand@ct.ufpb.br (A. Subramanian), uchoa@producao.uff.br (E. Uchoa).

0360-8352/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cie.2011.07.023

The adoption of virtualization technologies for power-aware optimization in server clusters turns out to be a challenging research topic. Specifically, the ability to dynamically distribute server workloads in a virtualized server environment allows for turning off physical machines during periods of low activity, and bringing them back up when the demand increases. Moreover, server on/off mechanisms can be combined with CPU DVFS (*Dynamic Voltage and Frequency Scaling*), which is a technique that consists of varying the frequency and voltage of the CPU at runtime according to processing needs, in order to provide even better power optimizations. Examples of DVFS capabilities implemented by current microprocessors are Intel's "Enhanced Speedstep Technology" (Intel Corp.'s, 2011), and AMD's "PowerNOW!" (Advanced Micro Devices, Inc., 2011a).

Research studies show that servers in data centers are loaded between 10% and 50% of peak, with a CPU utilization that rarely surpasses 40% (Barroso & Hölzle, 2007). Thus, the consolidation of the different workloads of services in a cluster, using server virtualization techniques, can reduce energy consumption and increase the utilization of physical servers (Kusic, Kephart, Hanson, Kandasamy, & Jiang, 2009; Srikantaiah, Kansal, & Zhao, 2008; Wang, Wang, Chen, & Zhu, 2008). In addition, techniques for processors energy-saving, typically considered major consumers of power in a server, have been used in the literature (Bertini, Leite, & Mossé, 2010; Horvath, Abdelzaher, Skadron, & Liu, 2007; Rusu, Ferreira, Scordino, Watson, & Mossé, 2006).

The case of heterogeneity in server cluster was addressed by Heath, Diniz, Carrera, Meira, and Bianchini (2005). However, their approach considers a non-virtualized cluster and addresses a different optimization problem. Related optimization approaches, based on the bin packing problem, for configuring virtualized servers are described in the literature (see Bichler, Setzer, & Speitkamp, 2006; Khanna, Beaty, Kar, & Kochut, 2006). However, their models are not designed for power-aware optimization. A heuristic-based solution for the power-aware consolidation problem of virtualized clusters, without adopting DVFS, is presented in Srikantaiah et al. (2008), but it does not show any analysis and guarantees to find solutions that are at least near to the optimal. In the same direction, a recent power-aware optimization algorithm for VM consolidation is presented in Wang and Wang (2010). However, their algorithm relies on an independent DVFS policy that performs locally in each server, which may not lead to an optimal solution of which servers should be active and their respective CPU frequencies. In contrast, our solution provides an integrated optimization approach by leveraging combined server on/off and DVFS techniques.

A dynamic resource provisioning framework is developed in Kusic et al. (2009) based on lookahead control optimization, but their approach does not consider DVFS. A power-aware migration framework for virtualized HPC (High-performance computing) applications is presented in Verma, Ahuja, and Neogi (2008). Similarly to our approach, it relies on virtualization techniques used for dynamic consolidation, although the application domains are different and DVFS technique is not employed. Other related approaches, such as presented in Elnozahy, Kistler, and Rajamony (2003), Chen et al. (2005), Rusu et al. (2006), Khargharia, Hariri, and Yousif (2008) and Bertini et al. (2010) also rely on DVFS techniques and include server on/off mechanisms, based on the seminal papers of Pinheiro et al. (2001) and Chase et al. (2001), for power-aware optimizations. However, these approaches are not designed (and not applicable) for virtualized server clusters. That is, they do not consider the optimization of multiple service workloads in a shared cluster infrastructure.

In this paper, we propose an efficient approach to solve a relevant cluster optimization problem which, in practice, can be used as an embedded module to implement a power and performance management solution in a real server cluster. The test instances used to evaluate the proposed optimization approach are derived from real measurements of power and performance data of real servers. To the best of our knowledge, this is a novel approach that simultaneously deals with (i) the CPU DVFS technique combined with server on/off mechanisms, (ii) the case of server heterogeneity, (iii) virtualized server environments, (iv) an efficient optimization approach, which is based on Column Generation (CG) techniques. The key aspects of our approach are the basis on rigorous and robust optimization techniques, given by high quality solutions in short amount of processing time, and experimental results on the cluster configuration problem for large-scale heterogeneous server clusters that can make use of virtualization techniques.

The remainder of the paper proceeds as follows. Section 2 formaly describes the optimization problem dealt in the current work. Section 3 explains the proposed CG approach. Section 4 presents and discusses the results obtained. Section 5 contains the concluding remarks of this work.

## 2. Cluster optimization problem

The cluster optimization problem is to determine the most power efficient configuration (i.e., which servers must be active and their respective CPU frequencies) that meets the required performance of multiple services on a shared cluster of servers, as well as to decide on the best mapping of these services to servers. The goal of our optimization approach is to reduce power consumption in the virtualized cluster while meeting service performance demands.

### *2.1. Server cluster scenario*

In our approach, the target cluster scenario (shown in Fig. 1) consists of a group of replicated physical servers. The server cluster presents a single view to the clients through a special entity termed *dispatcher server*, which distributes incoming requests among the actual *servers* that process the requests. The dispatcher server is also responsible for monitoring and managing the current configuration of the servers. According to Church et al. (2008), it is reasonable to have server clusters that are built on the order of thousands to tens of thousands of physical servers, whereas medium clusters are on the order of hundreds of servers.

A cluster configuration solution is given by (a) which servers must be active and their respective CPU frequencies and (b) a corresponding mapping of the services to physical servers. We assume a virtualized server environment, where multiple services can be hosted on a single physical server. Using server virtualization, each service runs as a virtual machine that includes its own operating system (e.g., Linux) and any additional software or library required to support the respective service execution, such as a Web or e-mail service.

To make decisions on which CPU frequency a server must operate, our approach relies on DVFS technique available on current microprocessors. This technique allows for dynamically adjusting the performance states (P-states) at which the server can operate when the CPU is active, which consists of a pre-defined set of frequency and voltage combinations. A DVFS-capable processor has only discrete levels of operating points.

As an example of DVFS power management capability, Table 1 shows the P-states and power consumption for the Intel Pentium M 1.6 GHz processor, whose data are available on a White Paper from Intel (Intel Corp.'s, 2011). Notice that performance state *P*0 is the highest P-state and *Pn* is the lowest one. In practical terms, for instance, the processing performance of a web server (measured in requests per second) is linearly proportional to its CPU frequency.