The Journal of Systems and Software 125 (2017) 365–379

Contents lists available at ScienceDirect

The Journal of Systems and Software

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

Versatile workload-aware power management performability analysis of server virtualized systems

Mohamed Escheikh[a], Kamel Barkaoui[b,*], Hana Jouini[a]

[a] SYS'COM ENIT, University of Tunis-El Manar, Tunis, Tunisia
[b] CEDRIC, CNAM Paris, France

ARTICLE INFO

Article history:
Received 22 November 2015
Revised 12 November 2016
Accepted 19 December 2016
Available online 3 January 2017

Keywords:
Virtualization
SVSs
Performability
SRNs
Workload-aware PM
Power-performance trade-off

ABSTRACT

The widespread integration of virtualization technologies in data centers has enabled in the last few years several benefits in terms of operating costs and flexibility. These benefits maybe boosted through join optimization of power management (PM) and dependability for virtualized systems. This indeed involves developing appropriate models to better understand their performability behavior whenever they are exposed to predictable (e.g. rejuvenation) and unpredictable breakdowns. We propose in this paper a performability analysis of server virtualized systems (SVSs) using a workload-aware PM mechanism based on non-Markovian Stochastic Reward Nets (SRNs) modeling approach. This analysis investigates interactions and correlations between several modules involving workload-aware PM mechanism, dynamic speed scaling processing, virtual machine (VM) and virtual machine monitor (VMM) both subject to software aging, failure and rejuvenation. We show through numerical results, using quantitative and qualitative metrics, how performance, power usage and efficiency are impacted by workload-aware PM mechanism. We show also how judicious choice of tunable attribute (i.e. *Timeout*) of the proposed PM mechanism with respect to workload can lead to a good power-performance trade-off.

1. Introduction

The increasing today's trend in Information Technologies (IT) toward modern services such as social media, rich internet content and e-commerce has generated unprecedented on-demand traffic requiring scalable processing capabilities. This new context has encouraged datacenters' operators to invest massively in architectures based on virtualization technology. The objective is to take advantage from scalability, reliable availability and cost-effectiveness benefits enabling significant enhancements in terms of return on investment (ROI). However to achieve successful migration toward these new architectures, cost-effective PM strategies should be adopted. Such strategies need to rationalize power consumption while meeting QoS and availability requirements. Optimizing design, development and deployment processes of virtualized systems involve usually many pertinent questions to be considered regarding PM, dynamic speed scaling processing, availability and performance. In order to achieve these purposes, performance and dependability analyses should be jointly investigated. The objective is to provide adaptive and fault tolerant virtualized resource allocation fitting with scalable on-demand requirements and matching with green computing targets. This is particularly important for datacenters evolving to larger scales and requiring operating and power consumption cost rationalization.

From infrastructure as a service (IaaS) provider's perspective, profit maximization is a high priority. In this regard, energy consumption rationalization while matching time-varying workload of bursty nature such as data-intensive applications without causing tangible performance overhead plays a crucial role (Beloglazov and Buyya, 2010). This enables IaaS providers to avoid both under-provisioning and over-provisioning and allows several benefits in terms of better resource utilization, energy consumption and cost-effective computing. Nowadays the emergence of advanced configuration and power interface (ACPI) standard (Hewlett-Packard Corporation, Intel Corporation, Microsoft Corporation, 2011) provides significant benefits to adapt PM to workload's dynamic. It enables power manageable component (PMC) to implement dynamic PM (DPM) capabilities while simplifying and accelerating the development of power-managed systems. Therefore it's of paramount importance while conceiving cost-effective PM policies to judicious use PMC capabilities in order to promote efficient resource provisioning in virtualized environments.

* Corresponding author.
E-mail addresses: mohamed.escheikh@enit.rnu.tn (M. Escheikh), kamel.barkaoui@cnam.fr (K. Barkaoui), hana.juini@gmail.com (H. Jouini).

http://dx.doi.org/10.1016/j.jss.2016.12.037
0164-1212/

1.1. Research scope

This work focuses on stochastic modeling of time-based cost-effective PM policy applied to SVS. The main objective of our contribution is to squeeze the most out of opportunistic configurations of ACPI power states in order to find convenient power-performance trade-off. The key idea is to enable adaptive switching between PM states (idle power serving states (i.e. C-states) and execution power saving states (i.e. P-states)). In this regard we propose performability analysis of a versatile workload-aware PM mechanism for SVS handling data intensive applications of bursty nature and enabling dynamic speed scaling. The analysis, based on non-Markovian SRNs modeling approach, highlights interactions and correlations between several entities involving:

1. VMM subject to software aging, failure and Cold-VM rejuvenation;
2. VM subject to software aging, failure and rejuvenation;
3. workload-aware PM mechanism;
4. resource provisioning based on dynamic speed scaling.

In this paper we extend the SVS availability SRN model investigated in Machida et al. (2010) for Cold-VM rejuvenation by retaining several features leading to more versatile and holistic SVS SRN availability model. This enables to investigate further performabilty metrics related to energy efficiency and power-performance trade-off. We particularly focus on time-based workload-aware PM and the issue to find opportunistic timeout value (the tunable parameter of the proposed PM mechanism) providing optimal power-performance trade-off for a given workload. Notice that SVS SRN availability models considered by Machida et al. (2010) (for three different VMM rejuvenation technics namely Cold-VM, Warm-VM and Migrate-VM) describe only VMM and VM features without dealing neither with handled applications nor with PM. These models are used to investigate the optimal rejuvenation periods for both VMM and VM enabling optimal SVS availability.

1.2. Research challenges

The key challenges and issues addressed through our contribution are:

1. How to develop comprehensive modeling and modular approach including several SVS entities namely VMM, VM, workload with bursty nature, resource provisioning with dynamic speed scaling and PM mechanism?
2. How to determine the right decision about switching between PMCs states in order to minimize SVS energy consumption while keeping acceptable performance in accordance with service level agreement (SLA) requirements?
3. How to implement a workload-aware PM policy to optimally solve the trade-off between delivered performance and energy savings for time-varying workload of bursty nature?
4. How to achieve good SVS power-efficiency by judiciously setting PM policy attributes with respect to workload?
5. How to define reward-based measures enabling to quantify performability metrics, efficiency and power-performance trade-off?

The rest of this paper is organized as follows: In Section 2 we discuss related works. In Section 3 we address performability analysis of SVS. We first address workload and fault tolerance issues related to SVSs before discussing the role of PMC and workload-aware PM mechanisms in SVSs. In Section 4 we present our main contribution regarding a versatile SVS availability SRN model describing SVS performability, proposed in this paper, and considering several details and interactions between two sub-modules (i.e. SVS availability SRN sub-model and workload-aware PM availability SRN sub-model). Section 5 investigates numerical results related to the versatile SVS availability SRN model to highlight power-performance trade-off through a set of pertinent metrics (i.e. firing frequency of some pertinent transitions, mean waiting time, power utilization, mean consumed power, efficiency) with respect to timeout. Finally Section 6 concludes this paper.

2. Related work

In this section we discuss recent research efforts in the area of PM and energy-efficiency in virtualized and cloud environments. In literature different modeling approaches has been proposed and investigated to better describe how the dynamics of different entities are involved in virtualized systems. A set of works has been proposed to investigate PM. Chen et al. (2005) investigated trade-off between energy consumption and performance based on SLA requirements. They highlighted time overhead of switching on/off nodes and relevant energy consumption on reliability. Nathuji et al. (2009) developed PM system for virtualized distributed architectures referred to as virtual power. They proposed the notion of soft power states and related mapping with hard power states to enable VM to manage soft power with traditional techniques. Control theory is used successfully in Gao et al. (2013) and Beloglazov et al. (2012) to tackle PM issues on servers. Elnozahy et al. (2002a) applied two PM mechanisms (Vary-On Vary-Off, VOVO) to explore power efficiency problem in homogeneous cluster hosting a single web application with SLAs specified in terms of response time requirements. Authors proposed five resource management policies aiming to estimate the CPU frequency meeting response time requirements. At fault tolerant level SRNs modeling approach was used (Elnozahy et al., 2002b; Machida et al., 2010) to build availability models for virtualized systems including VM and VMM with software rejuvenation. In (Wang et al., 2007; Xie et al., 2004) time-based rejuvenation policies under varying workload had been developed to enhance performability measure of cluster systems. A comparative study of three different rejuvenation policies is achieved through analytical models using deterministic and stochastic petri net (DSPN) models. In Lee and Zomaya (2012) authors present two energy-conscious task consolidation heuristics aiming to maximize resource utilization and explicitly accounting for both active and idle energy consumption. The proposed heuristics assign each task to the corresponding resource on which the energy consumption is minimized without performance degradation of that task. Authors in Jin et al. (2012) adopt an empirical approach to investigate how server virtualization affects the energy usage in physical servers. They identify a fundamental trade-off between on one hand the energy saving from server consolidation and on the other hand energy overhead and throughput reduction from server virtualization. In Beloglazov and Buyya (2010) authors proposed heuristics for dynamic reallocation of VMs in order to minimize energy consumption, while providing reliable QoS. In Kliazovich et al. (2012) a simulation environment is proposed for energy-aware cloud computing data centers. Simulation results obtained for different scenarios demonstrate the effectiveness of the simulator in utilizing PM schema for computing and networking components, such as voltage scaling, frequency scaling, and dynamic shutdown. In Machida et al. (2010) authors present a comparative study between availability models of three VMM rejuvenation techniques referred to as Cold-VM, Warm-VM and Migrate-VM. They show that rejuvenation trigger intervals of both VM and VMM need to be carefully chosen so as to reach high level VM availability. The provided comparative study was achieved in terms of steady-state availability and frequency of transactions lost per year.