Simulation Modelling Practice and Theory 34 (2013) 48–63

Contents lists available at SciVerse ScienceDirect

Simulation Modelling Practice and Theory

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

A novel black-box simulation model methodology for predicting performance and energy consumption in commodity storage devices

Laura Prada, Javier Garcia, Alejandro Calderon, J. Daniel Garcia *, Jesus Carretero

Computer Architecture Group, University Carlos III of Madrid, Av. Universidad Carlos III, 22, 28270 Colmenarejo, Madrid, Spain

A R T I C L E I N F O

Article history:
Received 27 August 2012
Received in revised form 17 January 2013
Accepted 19 January 2013
Available online 28 February 2013

Keywords:
Storage devices simulation
Black-box probabilistic models
I/O systems performance
Scalable performance evaluation
I/O energy consumption
Scalable simulation

A B S T R A C T

Traditional approaches for storage devices simulation have been based on detailed and analytic models. However, analytic models are difficult to obtain and detailed models require a high computational cost which may be not affordable for large scale simulations (e.g. detailed data center simulations). In current systems like large clusters, grids, or clouds, performance and energy studies are critical, and fast simulations take an important role on them.

A different approach is the black-box statistical modeling, where the storage device, its interface, and the interconnection mechanisms are modeled as a single stochastic process, defining the request response time as a random variable with an unknown distribution. A random variate generator can be built and integrated into a bigger simulation model. This approach allows to generate a simulation model for both real and synthetic complex workloads.

This article describes a novel methodology that aims to build fast simulation models for storage devices. Our method uses as starting point a workload and produces a random variate generator which can be easily integrated into large scale simulation models. A comparison between our variate generator and the widely known simulation tool DiskSim, shows that our variate generator is faster, and can be as accurate as DiskSim for both performance and energy consumption predictions.

1. Introduction

As demonstrated by the successful emergence of the Green500 [8] list, which provides a ranking of the most energy-efficient supercomputers in the world, energy has become as significant as performance. Consequently, the performance-per-watt has been established as a new metric to evaluate systems. Many researchers have shown interest in identifying in which cases there is room for improvement in the context of power efficiency. As a result, a lot of inefficiencies in relation to energy have been identified. Research works show that a CPU resting in an idle state reaches about 50% of peak power consumption [10]. Storage subsystems alone represent roughly 10–25% of the power consumed by the data center [14]. High power consumption in storage systems is expected, as an idle machine with one processor and two disks can easily spend as much power on disks as on processors [22]. Storage consumption can become a greater problem in storage subsystems where the average number of disks per machine is in the dozens.

* Corresponding author. Tel.: +34 918561316. Fax: +34 918561270.
E-mail address: josedaniel.garcia@uc3m.es (J.D. Garcia).

1569-190X/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.simpat.2013.01.006

In this way, simulation techniques are commonly used for both performance and energy consumption evaluation of many applications and systems. In order to obtain realistic results, data access to either a file system or database management system cannot be ignored. However, a key element in the data access studies is the storage device simulation model. A storage device simulation model accepts as input the parameters of an application workload (as a flow of device requests) and generates a performance metric prediction. The output performance measurement may be a general performance metric such as average bandwidth, throughput, and latency. Such metrics give an idea of the global performance of the device. However, if the device simulation model is integrated on top of a large system model, such as file systems or database managers, a detailed metric as request's response time is needed. Even more, disk's manufactures provide response time as an average value however, these metrics are not suitable for detailed simulations.

Scalability is one the most critical issues for current system's wide simulators. This is the case of large clusters, peer to peer or volunteer computing models, grid storage infrastructures [7], and content delivery networks. In all these cases, expensive realistic simulations spend a large amount of computational resources and computation time.

Traditionally, storage devices have been modeled by means of detailed or analytic models based on devices geometry [26], zone splitting [35,5], or the use of read-ahead caches [9] and request reordering [30], reaching the emulation level in many cases. An alternate approach is given by the black-box simulation models, where almost no knowledge of the storage device is required. The main advantage of the black-box simulation models is that access patterns of the storage devices can be modeled using a sequence of random variables, which models the time required to service disk's requests. This sequence of random variables is a stochastic process. The goal of these simulation models are the generation of values which fits this stochastic process. Experimental data must be obtained and analyzed to fit the distribution behind.

In a previous work [11], we presented initial results using black-box simulation models. The model introduced in the previous work was not based on probabilistic distributions but on histogram representations. In this way, many experimental data must be stored to get enough accuracy, requiring high memory costs. Additionally, other issues such as effects of queuing, sequentiality, and inactivity periods were not considered.

We propose a novel black-box simulation model methodology, named Back-Box Model based on Probability distributions (BBMP). Our solution is based on probability distributions, which are specially fast for generating values and providing response times predictions from disk drives. Effects of queuing, sequentiality, caching, and inactivity periods are considered as well.

The main contributions of this work are the following. First, we present BBMP that aims to predict disk response times in a very fast way, saving time in large simulation process. This solution is also easily extensible to other storage devices such as SSDs (Solid State Disks). Second, our method incorporates the usage of a new response time measurement tool, which can be used in any kind of disk, with any kind of interface. Third, unlike other works [20], our evaluations use real workloads, or synthetic with characteristics typical of real workloads. Also, we compare our built models with the ones generated for the widely known simulation tool DiskSim [5]. Finally, as we will show in our evaluation section, the proposed method can be used not only to predict performance, but also to obtain energy consumption estimations from disk drives.

The rest of this article is organized as follows. In Section 2, we briefly discuss about related work. Section 3 gives an overview of the method used to build the simulation model. Section 4 describes the procedure used to obtain disk response times. Section 5 explains how to construct models and how to implement random variate generators. Next, Section 6 includes a description of an energy consumption model. The experimental results are presented in Section 7. Finally, Section 8 presents our conclusions.

2. Related work

Preliminary work has been performed by proposing simulation models that fall into three main model categories: Analytic, Detailed, and Back-box. Table 1 shows an advantage/drawback comparison of different simulation models for storage devices.

Analytic models use mathematical equations that summarize disk behaviors. Thus, predictions are usually fast. However, analytic models are difficult to obtain as it is required to know the hard disk drive internals, and for some of its parts it is not easy to reach to an equation that describes them. Examples of such models include single disk models [30,35] or array disk models [37,17,42]. Detailed models are usually very accurate, as they emulate the hard disk drive behavior. One widely known detailed simulator is DiskSim [5]. The DiskSim simulation tool is able to model commodity hard drives by using parameters which are extracted from disks, using semi-automated algorithms [40] or by means of the DIXtrac disk charac-

Table 1
Advantages and drawbacks of different disk model categories.

	Advantage	Drawback
Analytic	Fast	Difficult to obtain
Detailed	Accurate	High cost
Back-box	Easy to obtain	Accurate?