Available online at www.sciencedirect.com

ScienceDirect

Electronic Notes in Theoretical Computer Science 337 (2018) 67–86

Electronic Notes in Theoretical Computer Science

www.elsevier.com/locate/entcs

Epistemic Uncertainty Propagation in Power Models

Marco Gribaudo[1], Riccardo Pinciroli[2]

Dipartimento di Elettronica, Informatica e Bioingegneria
Politecnico di Milano
Via Ponzio 34/5, Milano 20133, Italy

Kishor Trivedi[3]

Department of Electrical and Computer Engineering
Duke University
Durham, NC 27708, USA

Abstract

Data-centers have recently experienced a fast growth in energy demand, mainly due to cloud computing, a paradigm that lets the users access shared computing resources (e.g., servers, storage, etc.). Several techniques have been proposed in order to alleviate this problem, and numerous power models have been adopted to predict the servers' power consumption. Some of them consider many server resources, some others account for only the CPU, that has proven to be the component responsible for the largest part of a server's power consumption. All these models work with generally inaccurate input parameters. However, none of them takes into account the effects of such inaccuracy on the model outputs. This paper investigates how epistemic (parametric) uncertainty affects a power model. Studying the impact of epistemic uncertainty on power consumption models makes it possible to consider loads with a probability density while investigating the battery depletion time or the amount of energy required for a given task.

Keywords: power consumption, energy consumption, power models, epistemic uncertainty, parametric uncertainty, uncertainty propagation, M/M/c/K

1 Introduction

Servers' energy consumption is now one of the main issues to be considered while developing new applications. For example, the U.S. data-centers energy consumption is estimated to be 73 billion kWh in 2020 [31]. Another energy consumption related problem is about managing those big infrastructures in order to improve

[1] Email: marco.gribaudo@polimi.it

[2] Email: riccardo.pinciroli@polimi.it

[3] Email: ktrivedi@duke.edu

https://doi.org/10.1016/j.entcs.2018.03.034
1571-0661/

their efficiency and fault tolerance [19]. Furthermore, uninterruptible power supplies (UPS) are used for emergency power when the main power source fails. To better investigate these problems and to decrease the global amount of energy consumed by data-centers, accurate power consumption models are required. In past years, scientific and business organizations provided several power models to estimate the data-centers power consumption [11,1,14]. Depending on the granularity at which the problem is studied, the power models consider only few resources or the whole data-center.

The power consumption models are generally solved with fixed values for their input parameters. However, since those parameters are estimated from finite number of collected samples, they introduce into the model some uncertainties due to partial information. Thus, the uncertainty needs to be propagated from the inputs to the output. This kind of parametric uncertainty is called epistemic uncertainty [21].

In this paper we wish to apply epistemic uncertainty propagation to power consumption estimation, when the data-center arrivals forms a homogeneous Poisson process and service times are exponentially distributed (i.e., the rates are fixed over time). Although several works in literature take into consideration aleatory uncertainty while analyzing a general model [13,30], just a few consider the epistemic uncertainty. Indeed, while aleatory uncertainty is due to the natural variations of the physical phenomenon modeled by the system, epistemic uncertainty is introduced into the model by a lack of knowledge (i.e., finite number of observations, that in our case makes estimation of Λ and M inaccurate) and needs to be propagated to the output. Albeit in the former case the uncertainty is reduced improving the model itself, in the latter one it may be curtailed by collecting a larger amount of samples for a more accurate input parameter estimation. For this reason, we assume some of the input parameters of the power model are assessed with uncertainty, thus to be stochastic. Those input parameters become input random variables (r.v.) with a probability density function (pdf). Thus, the power model does not return an exact value, but some stochastic results with a confidence interval. Hereinafter, we refer to r.v. with capital letters, and use the small ones to refer to their observed values. In order to propagate the uncertainty from model inputs to its outputs, we adopt a multi-dimensional integration technique [21,23,22] that has been already applied to dependability models. To the best of our knowledge, epistemic uncertainty has never been applied to models that estimate the servers' power consumption.

In order to investigate how epistemic uncertainty propagation affects the power consumption of a machine, we collect several samples while performing some benchmarks, thus considering different resource utilizations. In particular, we observed an ASUS desktop computer with an Intel i7-3770 CPU@3.4GHz, with 4 cores and a Simultaneous multi-threading (SMT) level of 2 (i.e., it can concurrently run 8 parallel threads, in spite of the 4 physical cores), a 16GB memory, a dedicated GeForce GTX 560 GPU, two hard disk drives and a solid state disk, and running an Ubuntu 14.04 OS. Albeit some relationships between CPU utilization and server's power consumptions are given in [11], they are affected by SMT when its level is higher