Mini Review
Volume 15:12, 2022

Journal of Computer Science & Systems Biology

ISSN: 0974-7230

Open Access

A Review on Designing the Energy Usage of Digital Entities Used in the Telco Cloud

Lisandro Stepanov*
Department of Information Science, University of Oslo, Problemveien 7, 0315 Oslo, Norway

Abstract

Several surveys of the findings of research into the power consumption of virtual entities (VEs, also known as virtual machines (VMs) or containers) have been published. Our contribution to this work is a thorough examination of the dynamics of research itself: the challenges, approaches, pitfalls, fallacies, and research gaps, without ignoring the research results. The prospective researcher who wants to understand the dynamics of research into predictive modelling and supporting measurements of power consumption by individual VEs relevant to the telco cloud is our target audience. A thorough frequency analysis is used to characterise dynamics, which we do using a novel method we developed that is unique in its ability to parse research literature. A prospective researcher can obtain a thorough characterization of the problems, approaches, developments, formal methods, pitfalls, fallacies, and research gaps that characterise this research space by using the visual aids we provide and our observations through cross-cutting themes. Among the themes identified by our survey, we noted that all of the problem categories we identified touch on one or more of a set of seven major variables that may affect power consumption by virtual entities and the resulting model representations: workload type, virtualization agent (VM or container) characteristics, host machine resources and architecture, temperature, operating frequency, attribution of a fraction of consumed power.

Keywords: Telco cloud • Telecommunications • Power consumption

Introduction

We highlight here the misconception of the Data Plane Development Kit's (DPDK) power efficiency (commonly misrepresented as a power hog), the often-unacknowledged limitations of widely used linear models, the problematic use of benchmarks in model validation, the failure to precisely identify the physical contexts of some experimental research, and the influence of synthetic workload generators on measurement. We have also highlighted the unavoidable need to precisely define the scope and limitations of models, as well as the fallacy of seeking a "universal" power model. We identify four research gaps, container power consumption modeling, the effect of overcommitment on power efficiency, investigation and classification of DPDK applications; and modelling of power consumption by virtualized I/O. (a challenge which is starting to receive some attention).

It is difficult to precisely measure a VE's power consumption because measurements of its host's power consumption cannot be directly related to it. Individual VEs co-hosted on a physical machine's power consumption cannot be measured using hardware power metres. Furthermore, the power consumption of a VE varies depending on its hosting machine. As a result, precise modelling of energy- and/or power consumption is required for accurate measurement of VEs. Accuracy in power and energy measurements is a general requirement for designing energy- and/or power-efficient operations. It is also necessary for billing in multi-tenant environments, so that the Infrastructure Provider (IPr) can charge customers a reasonable amount for the resources (including energy) they consume. Before delving into the scope of our survey, we pose three key questions that will guide our work.

Literature Review

Energy efficiency in the Internet (and in computing and telecommunication networks in general) has become a significant problem, receiving increasing attention since the early years around 2000 (see, for example, and references therein), beginning with cloud computing infrastructures and then extending to mobile and fixed networks. Indeed, it has been demonstrated that smaller data centres, into which telecommunications points of presence (PoPs) can be classified, account for approximately 95% of data centre energy use in the United States. Furthermore, when compared to hyperscale server farms (the remaining 5%), this use is relatively inefficient. Indeed, the effect of virtualization technologies on power consumption in public telecommunications networks (PTNs) remains unknown. There is a widespread belief that Network Functions Virtualization (NFV) will result in lower energy consumption as a result of resource consolidation and increased flexibility in turning unused hardware (HW) on and off as needed. However, "massive introduction of general-purpose HW enabled by NFV would tend to increase power requests in comparison to specialised HW solutions". As a result, power-aware management and control mechanisms are required in these environments.

Simultaneously, the complexity of these mechanisms, as well as the level of human intervention, must be limited in order to keep operational expenditures (OPEX) within reasonable limits. Comparative analyses of infrastructure implementations with and without virtualization are one approach to understanding this impact. This approach is used in the study of the evolved packet core (EPC). This research demonstrates that the virtualized implementation is in fact less energy efficient. Unfortunately, the scope of virtualization and containerization within the converged wireless and wireline infrastructure is extremely broad, and a single "use-case" cannot be generalised to an overall statement. As a result, we note that the scope of our survey requires an operational context, which we propose in the following.

We propose that our network-operations context is the telco cloud. In attempts to describe it, the term "Telco cloud" evokes a number of common terms. Four such terms are virtualization, software-defined networking

---

***Address for Correspondence:** Lisandro Stepanov, Department of Information Science, University of Oslo, Problemveien 7, 0315 Oslo, Norway, E-mail: LisandroStepanov30@gmail.com*

Received: 07 December, 2022, Manuscript No. jcsb-23-85066; **Editor assigned:** 08 December, 2022, Pre QC No. P-85066; **Reviewed:** 21 December, 2022, QC No. Q-85066; **Revised:** 26 December, 2022, Manuscript No. R-85066; **Published:** 03 January, 2023, DOI: 10.37421/0974-7230.2022.15.447

(SDN), automation, and orchestration. Edge computing, containerization, microservices, and resilient infrastructure are also popular terms. We propose three key observations that organise these terms into a unified picture of the telecom cloud. The complementary collaboration of the PTN operator's (PTNO) network, computing, and storage infrastructure with that of global infrastructure and application providers is well demonstrated in. A distributed cloud infrastructure operates at network junctions (transport and interconnect). It consists of (cloud) infrastructure owned and operated by the PTNO, public cloud providers, and enterprises that use their joint service.

Discussion

We survey predictive energy and power models, as well as measurements that aid in the qualitative and/or quantitative prediction of telco cloud consumption by individual VEs [1-3]. The rationale that drove our selection can be summarised as follows: we sought works that measure real-time power consumption by VEs and/or model real-time power consumption by VEs. The VE is the sole focus of measurements and modelling. However, the devil is in the details, so the specifics of this simple rationale must be worked out. One more specific point concerns the VEs themselves. There are software technologies that are functionally critical to VEs, which we will discuss further in later sections. Works that measure and/or model the power consumption of such technologies are eligible. It is not difficult to justify this claim of scope.

Because power consumption is a scalar quantity, lowering the power consumption of a VE component results in lowering the VE's power consumption. To be fair, the translation is not direct. A generalisation of Amdahl's law occurs to me: improvement in a component, as measured by some metric, is dampened by the ratio of that component's use (as measured by that metric) to the system's (the VE's) use (measurable by the same metric). However, the finer point alluded to at the start of this paragraph can be safely summarised as follows. In scope is research into the measurement and modelling of power consumption by a VE component. For example, we would include the Data Plane Development Kit (DPDK) because it performs a critical networking function (VEs that serve as virtual network functions (VNFs)); Include a study that compares the power consumption of a VE when using two different implementations of input/output virtualization technology, such as SR-IOV (single-root IO virtualization) and paravirtualization [4,5].

Exclude a study that compares the power consumption of different network adapter (or network interface card (NIC)) architectures unless it reveals the impact of these architectures on the power consumption of VEs.

Our final point is about workload. The universal power model is a fallacy, and the main reason is that the interaction between workload and architecture cannot be fixed indefinitely. This is not to say that modelling is a waste of time. It simply means that the model must be subjected to validity constraints in terms of workload and architecture. As a result, we do not rule out modelling and measurement because of their workload or architectural scope. We do, however, observe that such models are traps for those who use them without being aware of their limitations.We'll now go over some recent, related surveys, emphasising their methodology and analytical approach [6,7].

Against this backdrop, we summarise the survey's novelty. Much research has been conducted in recent years to model the power consumption of servers in data centres and cloud environments in general. Several surveys cover various aspects of modelling and energy efficiency approaches for servers and virtual entities.Presents an analysis of power models from the micro- to macro-level. This survey examines various aspects and levels of hardware and software-centric modelling techniques. Researchers investigated models based on computing resources (CPU, memory, storage, and I/O), system architecture (such as single or multiple cores), GPU availability, system/network components, operating systems, and virtualization environments. They categorise existing models at various layers, ranging from architecture level modelling to power models for entire data centres.

Conclusion

According to the classification in, power models are affected by various organisational contexts, such as the power consumed by system components, running applications, and/or process execution strategy. These metrics, on the other hand, can be used to generate additive component-based, regression-based, or machine-learning-based power models. Additive models typically present an aggregated view of server power consumption, which could be based on different resources (such as CPU, RAM, I/O, and disc) or the server's static and dynamic power consumption. Power is mostly related to the dynamic evolution of some measured system parameters in regression-based models. Power modelling with machine learning techniques is a sophisticated research area that can be divided into supervised, unsupervised, reinforcement, and evolutionary learning.

Acknowledgement

None.

Conflict of Interest

Authors declare no conflict of interest.

References

1. Lawrence, Neal and Stefan Elbel. "Analysis of two-phase ejector performance metrics and comparison of R134a and CO_2 ejector performance." *Sci Technol Built Environ* 21 (2015): 515-525.
2. Chunnanond, Kanjanapon and Satha Aphornratana. "An experimental investigation of a steam ejector refrigerator: The analysis of the pressure profile along the ejector." *Appl Therm Eng* 24 (2004): 311-322.
3. Martínez, William Ruíz, Yesid Díaz-Gutiérrez and Roberto Ferro-Escobar, et al. "Application of the internet of things through a network of wireless sensors in a coffee crop for monitoring and control its environmental variables." *TecnoLógicas* 22 (2019): 155-170.
4. Ku, Meng-Lin, Wei Li and Yan Chen, et al. "Advances in energy harvesting communications: Past, present, and future challenges." *IEEE Commun Surv Tutor* 18 (2015): 1384-1412.
5. Sharma, Neelam, B.M. Singh, and Karan Singh. "QoS-based energy-efficient protocols for wireless sensor network." *Sustain Comput Inform Syst* 30 (2021): 100425.
6. Roscoe A.W and G. M. Reed. "A timed model for communicating sequential processes." *Theor Comput Sci* 58 (1988).
7. Kim, Hye Yeon and Frederick T. Sheldon. "Testing software requirements with z and statecharts applied to an embedded control systemt0t1." *Softw Qual J* 12 (2004): 231-264.

How to cite this article: Stepanov, Lisandro. “A Review on Designing the Energy Usage of Digital Entities Used in the Telco Cloud.” J Comput Sci Syst Biol 15 (2022): 447.