UNIVERSITAT POLITÈCNICA DE VALÈNCIA

MÁSTER DE AUTOMÁTICA E INFORMÁTICA INDUSTRIAL

Trabajo de Fin de Máster

ENERGY EFFICIENCY IN DATA CENTERS

Autor: Dr. Jordi Arjona Aroca

Director: Dr. Jose Enrique Simó Ten

Valencia, 17 de Diciembre de 2015

Trabajo fin de Máster

ENERGY EFFICIENCY IN DATA CENTERS

Autor:	Dr. Jordi Arjona Aroca
Director:	Dr. Jose Enrique Simó Ten

Firma del tribunal calificador:

Presidente:

Vocal:

Secretario:

Calificación:

Valencia, de de

Alea iacta est

Julius Caesar,
before crossing the Rubicon

Abstract

With the rise of cloud computing, data centers have been called to play a main role in the Internet scenario nowadays. Despite this relevance, they are probably far from their zenith yet due to the ever increasing demand of contents to be stored in and distributed by the cloud, the need of computing power or the larger and larger amounts of data being analyzed by top companies such as Google, Microsoft or Amazon.

However, everything is not always a bed of roses. Having a data center entails two major issues: they are terribly expensive to build, and they consume huge amounts of power being, therefore, terribly expensive to maintain. For this reason, increasing the energy efficiency (and hence reducing the carbon footprint) of data centers has been one of the hottest research topics during the last years. In this master thesis we propose different techniques that can have an impact in the maintenance costs of data centers of any size, from small scale to large flagship data centers.

We target energy efficiency in data centers as the main target of this master thesis. We first make a characterizing the power requirements of a data center server given that, in order to properly increase the energy efficiency of a data center, we first need to understand how energy is being consumed. We present an exhaustive empirical characterization of the power requirements of multiple components of data center servers, namely, the CPU, the disks, and the network card. To do so, we devise different experiments to stress these components, taking into account the multiple available frequencies as well as the fact that we are working with multicore servers. In these experiments, we measure their energy consumption and identify their optimal operational points. Our study proves that the curve that defines the minimal power consumption of the CPU, as a function of the load in Active Cycles Per Second (ACPS), is neither concave nor purely convex. Moreover, it definitively has a superlinear dependence on the load. We also validate the accuracy of the model derived from our characterization by running different Hadoop applications in diverse scenarios obtaining an error below 4.1% on average.

The second topic we study is the Virtual Machine Assignment problem (VMA), i.e., optimizing how virtual machines (VMs) are assigned to physical machines (PMs) in data centers. Our optimization target is to minimize the power consumed by all the PMs when considering that power consumption depends superlinearly on the load. We study four different VMA problems, depending on whether the number of PMs and their capacity are bounded or not. We study their complexity and perform an offline and online analysis of these problems. The online analysis is

complemented with simulations that show that the online algorithms we propose consume substantially less power than other state of the art assignment algorithms.

Table of Contents

List of Tables

List of Figures

Part I

Background

Chapter 1

Introduction

Internet has revolutionized our world. In 15 years it has passed from being hardly found in any home to be hardly not found in any pocket. We have become Internet-addicts and got used to continuously check our emails, videos, photos,... Through Google, Facebook, Twitter, we learn about what is happening with our friends or in any remote corner of the world, we are able to contrast news, we have access to any kind of information we are curious about. At the same time, Internet has also overturned the business world, simplifying and reducing the costs of sharing information in and between companies, and allowing any company, no matter how small it is, to have customers all over the world.

This has become real thanks to new concepts like the Internet of things, social networks, or cloud computing. However, at the end of the day, what Internet has done is putting huge amounts of data available to everyone. One of the keys of this availability of data has been the proliferation of data centers. Although some data centers are not necessarily large, like the ones usually deployed at many universities, companies or government institutions, large companies such as Google, Facebook, Amazon or Microsoft, among others, are building large scale data centers all over the world.

A large scale data center can be defined, in a nutshell, as an integrated facility housing a large amount of high end servers, up to the order of tens of thousands, interconnected by a dense network, hosting petabytes of data and consuming up to various tens of mega Watts. According to Belady [22], the building costs of a data center is between \$8M-\$30M/MW, being the average around \$20M/MW, depending on the kind of facility. These numbers lead to costs of $\$100-150$M for small/mid size facilities, while huge data centers, like Facebook's or Google's, can be in the order of \$600M.

However, although building a data center is expensive, they can be even more expensive to maintain. As we noted above, their average power consumption can be around 20MW per year, which unveils a second problem, their energy consumption. In a recent study, Van Heddeghem *et al.* [55] estimate that, between 2005 and 2012, worldwide aggregated data center energy consumption increased almost a 50%, reaching 270TWh from the previous 200TWh. This is roughly a

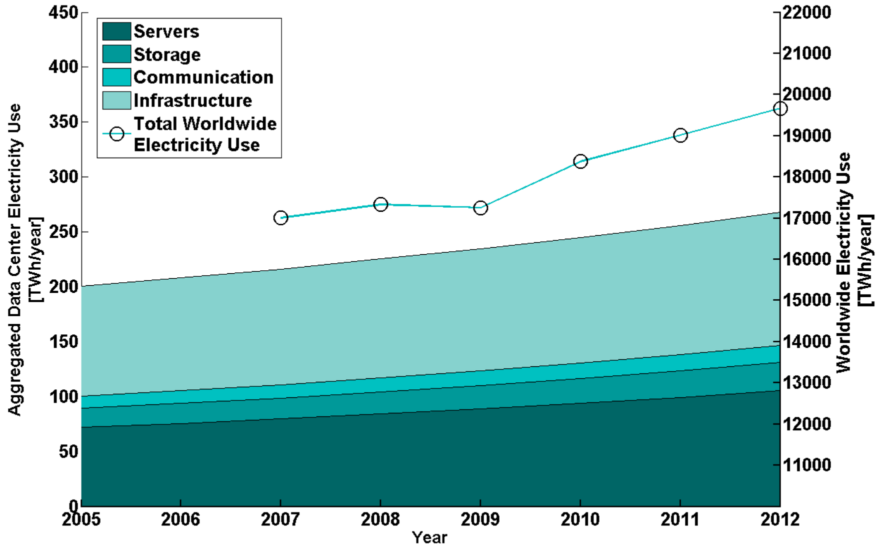

Figure 1.1: Worldwide use phase electricity consumption of data centers from 2005 to 2012 and the total worldwide electricity use from 2008 to 2012. Data center consumption is shown as the aggregation of its main consumers, servers, storage, communication and infrastructure.

1.5% of the total worldwide electricity consumption, and with a compound annual growth rate of a 4.4%.

Therefore, it is easy to see why *greening* data centers has emerged as one of the main targets of the research community during the last years. *Greening* data centers has, in fact, a two-fold objective, reducing operation costs, i.e., saving money; and improving data center sustainability, i.e., reducing the amount of energy consumption attributed to data centers. In the same way, efficiency of several data center components, for instance cooling systems, may lead to a reduction on the building costs.

Data center research has become a broad field of research. Even when we restrain ourselves to reducing building cost or increasing energy efficiency, the complexity and variety of subsystems that can be found in a data center result in a huge amount of particular problems. Even if we restrain ourselves to the increasing the energy efficiency of the different data center subsystems topic, the body of related work is overwhelming. In fact, if we look at some data provided by Barroso *et al.* [18, 19] we can see how the research in the field has contributed to change the energy consumption breakdown of data centers. In Figure 1.2 we can find the energy consumption breakdowns of a legacy data center with a PUE[1] value of around 2.0 in 2009 and 2013, in subfigures 1.2(a) and 1.2(b) respectively. Similarly, in subfigures 1.2(c) and 1.2(d), we can also see the evolution on the distribution of peak power usage in a hardware subsystem in a Google's data

[1]PUE responds to *Power Usage Efficiency* and it is one of the most commons and broadly accepted efficiency metrics. It measures the amount of cooling power needed versus the amount of electricity to run the IT infrastructure. An ideal ratio is 1.0.

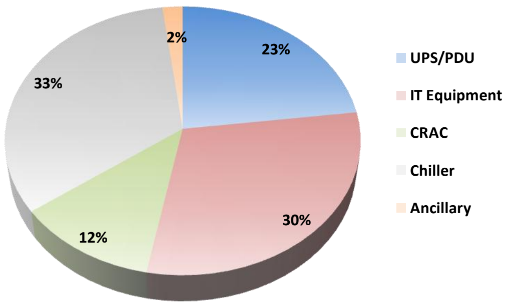

(a) Typical distribution of energy usage in a conventional datacenter with a PUE of 2.0 - 2009

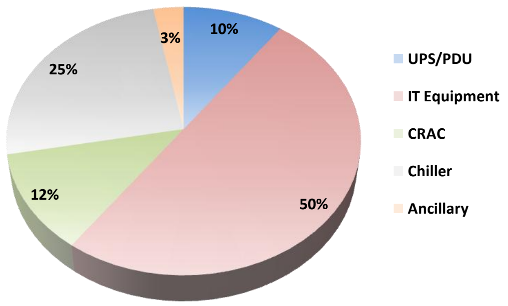

(b) Typical distribution of energy usage in a conventional datacenter with a PUE of 2.0 - 2013

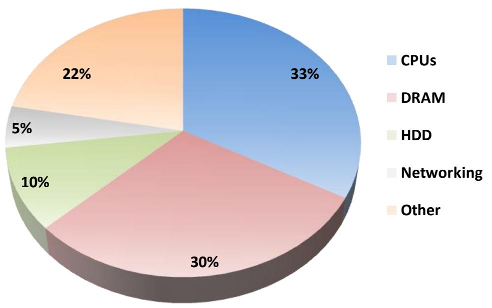

(c) Approximate distribution of peak power usage in a hardware subsystem in a Google's data center in 2007.

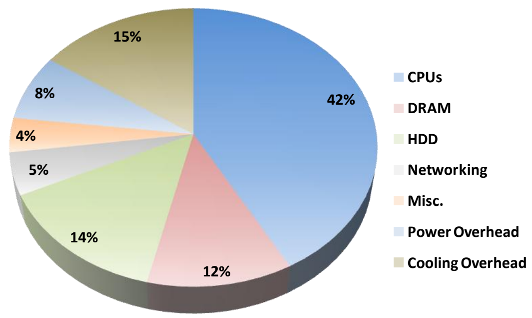

(d) Approximate distribution of peak power usage in a hardware subsystem in a Google's data center in 2012.

Figure 1.2: Evolution of the breakdown of a data center energy consumption and a hardware system between 2007-2013.

center between 2007 and 2012. This evolution is the result of intense research in multiple fields concerning each one of the pieces of hardware, usage policies or interaction between them.

However, although this evolution on the power requirements can be extended to many data centers, as for instance the (each time more) energy proportional servers, they can not be applied to all of them. One of the variables that conditions the application of these latest techniques is, for instance, the size of the data center. Big companies proudly exhibit the very low PUEs of their flagship data centers, like Facebook's PrineVille and Luleå, with $1.06 - 1.08$ and 1.07 PUE, or Google's Hamina, with 1.14 PUE. However, the resources which can be devoted to the design and construction of these data centers are not the same devoted to smaller ones or by smaller companies. Similarly, these low PUEs are usually achieved because of some particular aspects of the location of the data center, like the use of Finland's gulf water in Hamina. According to the Uptime Institute [88], the average PUE is around $1.8 - 1.89$, which gives a better idea of how much energy efficiency can still be improved.

Moreover, although a bit outdated, Bayley *et al.* [13], in 2007, presented a study quantify-

Table 1.1: Classification of Data Center types according to their size.

	Server Closet	Server Room	Localized Data Center	Mid-Tier Data Center	Enterprise Class Data Center
Size [sq ft]	< 200	< 500	< 1.000	< 5.000	> 5.000
# Servers (2005)	1.657.947	1.942.214	1.674.648	1.511.999	3.074.424
Estimated Energy consumption	11%	24%	21%	19%	24%
# Servers (est. 2009)	2.135.538	3.057.834	2.107.592	1.869.595	3.604.678

ing the amount of servers in different facilities according to their size. Some of the results of that study are presented in Table 1.1, showing the different categories, the estimated number of servers per category in 2005 as well as the distribution of energy consumption among them, and a prediction of the evolution of these numbers for 2009. These numbers show how most of the energy consumed is not necessarily in large enterprise data centers, but in smaller environments in which, in most cases, the PUE does not match the ones achieved by Facebook's or Google's flagship data centers. Nevertheless, it is important to remark that industry has become aware of this problem and more and more solutions are being provided each day to companies that only need small sized data centers, like Modular or Containerized Data Centers [5, 94] or integrated box solutions like IBM's Integrated Server Room [59]. In addition to this, Containerized Data Centers could even help to reduce the building costs of large data centers avoiding over built capacity and helping with over time scalability [89].

This master thesis is divided in two parts, a Background part, which includes this introduction and the state of the art of the two different, but related, problems we addressed in this document; and a second part, Understanding and Reducing Energy Consumption in Data Centers, that describes our work in both problems. This study intends to help to a better understanding of how energy is consumed in data centers as well as providing solutions which can be applied to data centers in any size range, from server closets to large enterprise data centers. We know provide some insights about what will be covered in this master thesis.

1.1 Understanding and Reducing Energy Consumption in Data Centers

As we have already mentioned, improving the energy efficiency of data centers has become an issue of capital importance both for economic and of environmental reasons. Due to this importance, the amount of techniques that have been proposed to help in this area is so huge that it is impossible to present them in just one document. Hence, we now only introduce some of the techniques which are relevant for this master thesis.

1.1.1 How Servers Use Power

Servers are like puzzles where each one of its pieces has its own share of power consumption. At the same time, the global power consumption of a each one of these pieces is not constant, it depends on the stress we introduce in each one of them. Depending on the amount of accesses we do to disk, to memory, on the amount of data we send to or receive from the network and on the amount of processing we do or the heat we generate, the power required by Hard Drives, Memory, Network, CPU or cooling units will vary. There are also more components, as we saw in Figure 1.2, but those are usually assumed to be the major contributors to the power consumption of data center servers.

However, servers are not power proportional [17], i.e., the total power consumed is not proportional to the amount of load being processed, and usually, the main contributor to power consumption is the fact of having the machine switched on. In the recent past, the amount of power consumed by an idle machine compared to the power consumed when it worked at full speed could easily add up to a 70% of its total power consumption. If we focus in the last 7 years, we can consult the available public data from the SPEC power benchmark web [38]. Comparing results from the last quarter of 2007 against the last available ones (second quarter of 2014), we can see a reduction on the power consumption of servers when idle compared to its peak consumption from barely a 60% in most servers to roughly a $20 - 25\%$ of power consumption.

If we consider only the active range, i.e., the power variation between the idle state and the peak consumption, it has been traditionally assumed that most of the power is consumed by the CPU. Similarly to what happens with servers, processors do not consume power linearly, in proportion to the load. Although processor power consumption has usually been modeled in a linear fashion, everything changed with the arrival of multicore processors able to work at multiple frequencies. Multicore processors introduced multiple changes. First, cpus in the same processor are able to share on-chip and on-die resources, increasing,hence, the synergies and reducing power requirements [21]. Also, there are new parameters to be considered as variable voltages and frequencies that determine CPU speed and, therefore, power requirements. Due to this new complexity, being able to understand how servers consume power has become a must if we want to devise any technique or strategy that intends to reduce power or energy consumption.

In this master thesis, we will present an empirical study were some of these aspects were analyzed and we were able to shed some light about the behavior of multicore and multifrequency machines based on real data. This knowledge can be applied in multiple techniques devoted to reduce the aggregated power and energy consumption of data centers. Not understanding the effect that placing a task in a server is going to have on its power consumption will result in non-optimal, in the best case, or in completely non-efficient, in the worst case, implementations of techniques such as speed scaling policies or virtualization strategies. We will now discuss about the latter two practices, speed scaling and virtualization, which are well known examples of techniques used to reduce the aggregated power and energy consumption in data centers.

1.1.2 Speed Scaling Based Techniques

Speed scaling is based on the ability of a processor to change its operating voltage and speed (frequency), and hence the speed and power consumption of the server. It is important to note that the values of voltage and frequency are not independent from one another. There is an intimate relation between them as, usually, the voltage conditions the range of available frequencies. In Table 1.2 we can find the different combinations of frequencies and voltages available for different processors.

Table 1.2: Relation between frequency and voltage for different processors

Processors					
AMD Opteron 6276		*Intel Xeon W3530*		*Intel Xeon E5606*	
Voltages	Frequencies	Voltages	Frequencies	Voltages	Frequencies
0.9375V to 1.3125V	1.4 *GHz*, 1.6 *GHz*, 1.8 *GHz*, 2.1 *GHz*, 2.3 *GHz*, 2.3 *GHz*	0.750V to 1.350V	1.596 *GHz*, 1.729 *GHz*, 1.862 *GHz*, 1.995 *GHz*, 2.128 *GHz*, 2.261 *GHz*, 2.394 *GHz*, 2.527 *GHz*, 2.666 *GHz*, 2.793 *GHz*, 2.794 *GHz*	0.800V to 1.375V	1.2 *GHz*, 1.333 *GHz*, 1.467 *GHz*, 1.6 *GHz*, 1.733 *GHz*, 1.867 *GHz*, 2 *GHz*, 2.133 *GHz*

One well known and extended implementation of speed scaling is *Dynamic Voltage and Frequency Scaling (DVFS)* which can be usually be found in the ample majority of servers which can be found nowadays in the market. DVFS can be configured with different governors or operating policies which will condition the way frequency adapts to the load in the system.

However, it is important to note that we can not just reduce the frequency as much as we want as it will affect the performance of the tasks being run in the machine. For this reason, usually, commercial implementations of DVFS have conservative policies whose main target is reduce operating frequency, and hence the consumption, of the machine when idle.

Most of the research in this field is devoted to find efficient policies which allow to reduce power consumption or energy consumption. It is important to remark the difference between both targets, let us give an easy example. Assume that we have a task that needs a time T to complete. If we reduce the operating frequency of the system, and hence the power consumption is reduced from C to C', it might happen that the task being run in our machine now needs a time T' to be run. If the total energy consumed $T \cdot C$ is larger than $T' \cdot C'$ we will have reduced the power consumption during a period of time, but spent more energy. This is neither good nor bad, both policies have their applications in different scenarios. However, we must remember that it is not trivial to optimize the power required and it is needed to carefully design the policies to be used.

1.1.3 Virtualization Based Techniques

We have already mentioned two important aspects of modern servers: that, with almost no exception, they are multicore and multifrequency servers, and that we pay a high cost in terms of power, just for having them idle, (i.e., powered on but not doing any task). However, think now, just for a second, that, years ago, running multiple tasks in a machine was through multithreading, i.e., running them in parallel with no isolation. When a task required some isolation, for security or other reasons, it had to be run alone in a server, what implies that the server resources not being used by that task were wasted. Keep in mind also that, in old servers the percentage of power used just for having them powered on was larger than nowadays. Additionally, there were some problematic situations with multithreading, as the existence of resource-greedy users or tasks. In order to tackle this situation we use *Virtualization*. Although virtualization was originally developed in the 1960s by IBM, it was forgotten and then recovered again in the 1990s. We can define virtualization as *"a technology that combines or divides computing resources to present one or many operating environments using methodologies like hardware and software partitioning or aggregation, partial or complete machine simulation, emulation, time-sharing, and many others*" [35]. We call each one of these operating environments a *Virtual Machine (VM)*. Hence, instead of running tasks in a per server basis or use multithreading sharing the resources pool, virtualization allows us to run tasks in a per VM basis, therefore running multiple tasks (with limited resources) independently in the same server.

Nevertheless, virtualization only opened the door for future improvements in how to reduce the power consumption of data centers from a server perspective. Two of these consequences were consolidation and virtual machine allocation techniques.

Consolidation is probably the most straightforward consequence of virtualization. Since we gained the ability of putting multiple virtual machines in one server it is logic trying to maximize the efficiency of servers. Consolidation aims to either maximize the aggregated number of tasks being run keeping a constant number of active servers, or minimize the number of servers needed to run a set of tasks. In both cases, the contribution of virtualization to increase the productivity and the energy efficiency of data centers is clear.

Similarly, and intimately related with consolidation, we have virtual machine allocation techniques. Given that the assignment of virtual machines to servers is an NP-hard problem (it can be easily reduced to problems such as bin packing or 3-partition for instance, as we will see in Chapter 4), multiple heuristics and algorithms have been devised to tackle the online version of the problem. Algorithms like First Fit, Packing, Most (Least) Loaded First ... try to obtain the best assignment of tasks to servers according to a certain magnitude, like cost, energy consumption... although in general try to minimize the number of active servers.

However, most of these algorithms are based on linear models for the power consumption of cores. Based on the insights we got with the characterization of a data center server we will study the effect of assuming a non-linear power consumption model for data centers servers. Based on this model, we will perform a competitive analysis of different VM to physical machine power aware assignment strategies under different hypothesis, like the capacity or the number of available physical machines. We will propose our own strategies and compare their power and energy consumption with the ones of some of the aforementioned algorithms (first fit, packing. . .).

1.2 Overview and Summary of Contributions

In the previous section we have enumerated several problems and challenges related to data centers. However, although each one of them are broad research topics, they are only a small subset of the challenging problems that are to be solved in such an ecosystem. We know present the problems in which we worked and the contributions we made for each one of them.

Focusing on the field of energy consumption in data center networks, we studied how energy is spent in physical machines. As we mentioned in Subsection 1.1.2, in order to devise wise consolidation algorithms or speed scaling policies for our systems, the first step is to properly understand how power is consumed in our physical machines. With the irruption of multicore servers we find a certain inconsistency in the literature as it is usual to find works that assume linear models of consumption for servers, while some others assume models different from linear (such as superlinear models). To shed some light in this problem, we performed an empirical study with 3 different servers of different architectures (Intel Xeon and AMD Opteron). Our first contribution was showing that the metric used to express load matters, and, hence, using relative magnitudes, like load percentage, might lead to deceiving results. We introduced the Active Cycles Per Second (ACPS) metric, which is an absolute magnitude, related to the frequency of operation, and that denotes the amount of computer cycles used to process load by a server per second. Our study analyzed the contribution to the total power consumption of three different components of data center servers, namely, CPU, disks, and network, and their dependencies in certain parameters, like the frequency. The most important contribution of this work is showing that the curve that defines the minimal CPU power as a function of the load is neither linear nor purely convex as has been previously assumed. Similarly, we also study the effects of the operating frequency and other parameters in the power consumption of disks and network. We validate our model by means of computing the PageRank metric of a graph and a WordCount application in a Hadoop platform, first without network activity, next with bulky network activity, and finally with a two-server cluster. We find that the energy can be estimated with an error that is below 4.1% on average and never worse than a 10%.

One of the conclusions of the previous work is that there exists an optimal point of operation distinct from the 100%. This means, in simple words, that running a server at full load and full speed, is not always optimal. This also contradicts some statements that have been traditionally assumed to be true regarding how much load has to be processed by a server.

Based on the latter study and the conclusions which can be extracted from it, a question that immediately follows is how to assign virtual machines to physical machines in a power efficient way. Our last contribution consists of an analytical study of what we call the Virtual Machine

Assignment (VMA) problem. We study this problem from different points of view, imposing, or not, restrictions on the number of physical machines or on their CPU capacity. Namely, we study 4 models, $(\cdot,\cdot)$-VMA (no restrictions in CPU capacity or number of servers), $(C,\cdot)$-VMA (CPU capacity is finite, number of servers is infinite), $(\cdot,m)$-VMA (CPU capacity is finite, number of servers is bounded) , and (C,m)-VMA (both CPU capacity and number of servers are finite). We show that the decision version of the (C,m)-VMA problem is strongly NP-complete. We show as well that the $(C,\cdot)$-VMA, $(\cdot,m)$-VMA and $(\cdot,\cdot)$-VMA problems are strongly NP-hard. Hence, there is no FPTAS for these optimization problems. We also show the existence of a PTAS that solves the $(\cdot,\cdot)$-VMA and $(\cdot,m)$-VMA offline problems. On the other hand, we prove lower bounds on the approximation ratio of the $(C,\cdot)$-VMA and (C,m)-VMA problems. With respect to the online version of these problems, we prove upper and lower bounds on the competitive ratio of the $(\cdot,\cdot)$-VMA, $(C,\cdot)$-VMA, $(\cdot,m)$-VMA, and (C,m)-VMA problems. Finally, we compare our algorithm to a modified version to other real approaches and show its advantages by simulation.

1.3 Roadmap

The rest of the master thesis is structured as follows. We start with a *Background* part, including this Introduction Chapter and Chapter 2, which presents an overview of the state of the art for each one of the problems we have mentioned. Part II, *Understanding and Reducing Energy Consumption in Data Centers* is divided in two chapters. In Chapter 3, we introduce a characterization of how energy is consumed in a physical machine by 3 of its main components, namely, CPU, disk and network. The second chapter, Chapter 4, presents an analytic study about how virtual machines should be assigned to physical machines in order to reduce the energy consumption in data centers and propose different algorithms tackling this problem. Finally, Chapter 5 concludes this master thesis with a summary of our main results, a discussion on the implications of these results and future research directions.

This master thesis covers contributions from the following literature:

- **Jordi Arjona Aroca**, Angelos Chatzipapas, Antonio Fernández Anta and Vincenzo Mancuso. "A Measurement-based Analysis of the Energy Consumption of Data Center Servers". *International Conference on Future Energy Systems (ACM e-Energy) 2014*, 63-74. June 2014.

- **Jordi Arjona Aroca**, Antonio Fernández Anta, Miguel A. Mosteiro and Christopher Thraves. "Power-efficient Assignment of Virtual Machines to Physical Machines". Workshop on Adaptive Resource Management and Scheduling on Cloud Computing (ARMS-CC).

- **Jordi Arjona Aroca**, Antonio Fernández Anta. "Empirical Comparison of Power-efficient Virtual Machine Assignment Algorithms". 4th IFIP Conference on Sustainable

Internet and ICT for Sustainability (SustainIt 2015).

- **Jordi Arjona Aroca**, Antonio Fernández Anta, Miguel A. Mosteiro and Christopher Thraves. "Power-efficient Assignment of Virtual Machines to Physical Machines". *Accepted for Publication in Future Generation Computer Systems Journal (FGCS), Special issue on "Advanced Topics in Resource Management for Ubiquitous Cloud Computing: an Adaptive Approach".*

- **Jordi Arjona Aroca**, Angelos Chatzipapas, Antonio Fernández Anta and Vincenzo Mancuso. "A Measurement-based Characterization of the Energy Consumption in Data Center Servers". *Accepted for publication in IEEE Journal on Selected Areas of Communications - Series on Green Communications and Networking (JSAC-SGCN).*

Chapter 2

Related Work

2.1 Characterizing the Energy Consumption of Data Center Servers

2.1.1 Background

The literature in the field of energy efficiency for data centers is, as we mentioned in the introduction in Chapter 1, simply overwhelming. For this reason, we focus our efforts in our particular problem, concentrating our attention on the characterization of data center servers and the energy they consume.

Indeed, although many energy saving techniques have been proposed during the recent years, such as virtualization plus consolidation or scheduling optimization [65, 78], in order to obtain full benefit of them it is crucial to have a good characterization of the servers in the data center, as a function of the utilization of the server's components. That is, it is necessary to know and understand the energy and power consumption of servers and how this changes under the different configurations. There is a large body of literature on characterizing servers' energy and power consumption. However, the existing literature does not jointly consider phenomena like the irruption of multicore servers and dynamic voltage and frequency scaling (DVFS) [96], which are key to achieve scalability and flexibility in the architecture of a server. With these new parameters, more variables come into play in a server configuration. Learning how to deal with these new parameters and how they interact with other variables is important since this may lead to larger savings.

It has been traditionally considered that the CPU is responsible for most of the power being consumed in a server, as we saw in Figure 1.2, and that this power increases linearly with the load. As we could see in that same figure, the power consumed by the CPU is significant, but the power incurred by other elements of the server, like disks and NICs (Network Interface Cards) is not negligible, and have to be taken into account. Moreover, we believe that the assumption that CPU power consumption depends linearly from the load in a server may be too simplistic, especially when the server has multiple cores and may operate at multiple frequencies. In fact, even the way load is expressed has to be carefully defined (e.g., it cannot be defined as a proportion of the maximal computational capacity of the CPU, since this value changes with the operational frequency). Therefore, more complex/complete models for the power consumed by a server are necessary. In order to be consistent, these models have to be based on empirical values. However, we found that there is a lack of empirical work studying servers energy behavior.

The study we present in Chapter 3 tries to partially fill this void by proposing a measurement-based characterization of the energy consumption of a server components with DVFS and multiple cores.

2.1.2 Related Work

There is a large body of work in the field of modeling server energy consumption and its components, both theoretically and empirically. The consumption of servers has been assumed as linear, e.g., by Wang *et al.* [95], Mishra *et al.* [76] or Beloglazov *et al.* [24], who assumed models in which energy consumption mainly depends linearly on CPU utilization. Based on the models, they proposed bin-packing-like algorithms to reduce energy consumption. Other works like the ones from Andrews *et al.* [10] or Irani *et al.* [62] proposed non-linear models, claiming that energy could be saved by running processes at the lowest possible speed.

Moving to the empirical field, we first classify works in two different groups, depending on whether they consider the effect of frequency in their analysis. We start with works not considering frequency. In this category we find articles proposing models where server components follow a linear behavior, like in [64,70,91] or more complex ones, like in [20,40,69]. In [70], Liu *et al.* proposed a simple linear model and evaluate different hardware configurations and types of workloads by varying the number of available cores, the available memory, and considering also the contribution of other components such as disks. Vasan *et al.* [91] monitored multiple servers on a datacenter as well as the energy consumption of several of the internal elements of a server. However, they considered that the behavior of this server could be approximated by a model based only on CPU utilization. Similarly, Krishnan *et al.* [64] explored the feasibility of lightweight virtual machine power metering methods and examined the contribution of some of the elements that consume energy in a server like CPU, memory and disks. Their model depends linearly on each of these components. In [40], Economou *et al.* proposed a non-intrusive method for modeling full-system energy consumption by stressing its components with different workloads. Their resulting model is also linear on the utilization of server components. Finally, Lewis *et al.* [69] and Basmasjian *et al.* [20] presented much more complex models which, apart from the contribution of different components of the server, consider extra parameters like temperature and cache misses as well as multiple cores. In particular, Lewis *et al.* [69] reported also an extensive study on the behavior of reading and writing operations in hard disk and solid state drives.

Next we move to the works which also consider frequency in their analysis. Miyoshi *et al.* [77] analyzed the runtime effects of frequency scaling on power and energy. Brihi *et al.* [28] presented an exhaustive study of DVFS using a `cpufrequtils` as we do. Main differences with our work were that they studied four different power management policies under DVFS and centered their study on the relationship between CPU and power utilization. However, they also presented interesting results about disk consumption that match partially our results, showing a flat consumption in reading operations and variations in the writing ones that they attribute to the size of the files being written. Although it was not the main objective of their work, Raghavendra *et al.* [85] performed a per-frequency and core CPU power characterization of two different blade servers. However, they claimed that CPU power depends linearly on its utilization. The main difference with our analysis is that we consider that the load supported by a server increases with the number of active cores and, hence, this load should not be represented in percentage. Gandhi *et al.* [46] published an analysis of global energy consumption versus frequency, based on DVFS and DFS and gave some intuition about the non-linearity of this relation.

Moreover, there are studies that model the energy consumption behavior for clouds and try to balance the load in order to operate the cluster in its most efficient load-power combination. MUSE [33] is one of the first works that consider a resource management architecture for data centers. Its energy efficient approach dynamically assigns jobs to the servers based on the workload (for CPU and disk) and the potential energy consumption. The authors measure the energy consumption of servers and switches involved in the cluster and conclude that at least 29% of the energy can be saved by MUSE for typical web workloads. In [87] the authors proposed a consolidation algorithm that considers the workloads of the servers in the cloud in order to find the least possible energy consumption point. Their study shows that the energy consumption of a server using variable loads for CPU and disks has an optimal operating point. Given the data from the various servers the algorithm can estimate the ideal load distribution among the servers. The authors in [14] modeled the energy consumption of data centers equipment (i.e., servers, storage, switches) for cloud computing based on existing energy consumption measurements or publicly available data sheets for each of the components (CPU, disk, network, switches). The model estimates the energy consumption per bit from the data center to the user and further analyzes the energy consumption for different types of services, i.e., storage, software, processing. However, existing works on clouds lack experimental inputs on energy consumption.

We conclude with some works that also consider frequency but do not model the energy consumption of a server. First of them, the work from Le Sueur *et al.* [67] presented an analysis of the evolution of the effectiveness of DVFS and how it is reduced in the newest and most optimized servers. They show that DVSF might be soon obsoleted by the adoption of ultra low power sleep modes. Ge *et al.* proposed PowerPack [49], a framework that includes a set of toolkits to perform an exhaustive profile of the power utilization of servers and its components. Their analysis is centered in showing the contribution of multicore system to the efficiency of several applications and, hence, no power characterization is presented. Finally, Basmadjian *et al.* [21] published an in deep analysis of the components of a processor and its contribution to the energy consumption of the CPU, shedding some light on the behavior of multicore servers. Some of their conclusions are very relevant to our work, as they show, for instance, that the energy consumption of multiple cores performing parallel computations is not equal to the sum of the power of each of those active cores. Our experiments and model support their findings and shed light on the nature of such effect.

2.2 Power Aware Assignment of Virtual Machines to Physical Machines

2.2.1 Background

The current pace of technology developments, and the continuous change in business requirements, may rapidly yield a given proprietary computational platform obsolete, oversized, or insufficient. Thus, outsourcing has recently become a popular approach to obtain computational services without incurring in amortization costs. Furthermore, in order to attain flexibility, such service is usually virtualized, so that the user may tune the computational platform to its particular needs. Users of such service need not to be aware of the particular implementation, they only need to specify the virtual machine they want to use. This conceptual approach to outsourced computing has been termed cloud computing, in reference to the cloud symbol used as an abstraction of a complex infrastructure in system diagrams. Current examples of cloud computing providers include Amazon Web Services [1], Rackspace [4], and Citrix [2].

Depending on what the specific service provided is, the cloud computing model comes in different flavors, such as infrastructure as a service, platform as a service, storage as a service, etc. In each of these models, the user may choose specific parameters of the computational resources provided. For instance, processing power, memory size, communication bandwidth, etc. Thus, in a cloud-computing service platform, various virtual machines (VM) with user-defined specifications must be implemented by, or assigned to[1], various physical machines (PM)[2]. Furthermore, such a platform must be scalable, allowing to add more PMs, should the business growth require such expansion. In this work, we call this problem the Virtual Machine Assignment (VMA) problem.

The optimization criteria for VMA depends on what the particular objective function sought is. From the previous discussion, it can be seen that, underlying VMA, there is some form of bin-packing problem. However, in VMA the number of PMs (i.e., bins for bin packing) may be increased if needed. Since CPU is generally the dominant power consumer in a server, as shown in Figure 1.2 and as we will show in Chapter 3, VMA is usually carried out according to CPU workloads. With only the static power consumption of servers considered, previous work related to VMA has focused on minimizing the number of active PMs (cf. [23] and the references therein) in order to minimize the total static energy consumption. This is commonly known as VM consolidation [66, 79]. However, despite the static power, the dynamic power consumption of a

[1]The cloud-computing literature use instead the term placement. We choose here the term assignment for consistency with the literature on general assignment problems.

[2]We choose the notation VM and PM for simplicity and consistency, but notice that our study applies to any computational resource assignment problem, as long as the minimization function is the one modeled here.

server, which has been shown to be superlinear on the load of a given computational resource [15, 53], is also significant and cannot be ignored. Since the definition of load is not precise, we use the definition we provide in Chapter 3 and define the load of a server as the amount of active cycles per second a task requires, an absolute metric independent of the operating frequency or the number of cores of a PM. The superlinearity property of the dynamic power consumption is also confirmed by the results that we show in Chapter 3. As a result, when taking into account both parts of power consumption, the use of extra PMs may be more efficient energy-wise than a minimum number of heavily-loaded PMs. This inconsistency with the literature in VM consolidation is supported by the results that will be presented in Chapter 3 and, hence, we claim that the way consolidation has been traditionally performed has to be reconsidered. In this work, we combine both power-consumption factors and explore the most energy-efficient way for VMA. That is, for some parameters $\alpha > 1$ and $b > 0$, we seek to minimize the sum of the α powers of the PMs loads *plus* the fixed cost b of using each PM.

Physical resources are physically constrained. A PMs infrastructure may be strictly constrained in the number of PMs or in the PMs CPU capacity. However, if usage patterns indicate that the PMs will always be loaded well below their capacity, it may be assumed that the capacity is unlimited. Likewise, if the power budget is very big, the number of PMs may be assumed unconstrained for all practical purposes. These cases yield 4 VMA subproblems, depending on whether the capacity and the number of PMs is limited or not. We introduce these parameters denoting the problem as **(*C*,*m*)-VMA**, where C is the PM CPU capacity, m is the maximum number of PMs, and each of these parameters is replaced by a dot if unbounded.

2.2.1.1 Problem Definition

We describe the $(\cdot,\cdot)$-VMA problem now for a better understanding of some of the works presented in the following related work section. Given a set $S = \{s_1, \ldots, s_m\}$ of $m > 1$ identical physical machines (PMs) of capacity C; rational numbers μ, α and b, where $\mu > 0$, $\alpha > 1$ and $b > 0$; a set $D = \{d_1, \ldots, d_n\}$ of n virtual machines and a function $\ell : D \to \mathbb{R}$ that gives the CPU load each virtual machine incurs[3], we aim to obtain a partition $\pi = \{A_1, \ldots, A_m\}$ of D, such that $\ell(A_i) \leq C$, for all i. Our objective will be then minimizing the power consumption given by the function

$$P(\pi) = \sum_{i \in [1,m]: A_i \neq \emptyset} \left(\mu \Big(\sum_{d_j \in A_i} \ell(d_j) \Big)^{\alpha} + b \right). \tag{2.1}$$

Let us define the function $f(\cdot)$, such that $f(x) = 0$ if $x = 0$ and $f(x) = \mu x^{\alpha} + b$ otherwise. Then, the objective function is to minimize $P(\pi) = \sum_{i=1}^{m} f(\ell(A_i))$. The parameter μ is used for consistency with the literature.

[3]For convenience, we overload the function $\ell(\cdot)$ to be applied over sets of virtual machines, so that for any set $A \subseteq D, \ell(A) = \sum_{d_j \in A} \ell(d_j)$.

We also define several special cases of the VMA problem, namely (C, m)-VMA, $(C, \cdot)$-VMA, $(\cdot, m)$-VMA and $(\cdot, \cdot)$-VMA. (C, m)-VMA refers to the case where both the number of available PMs and its capacity are fixed. $(\cdot, \cdot)$-VMA, where $(\cdot)$ denotes unboundedness, refers to the case where both the number of available PMs and its capacity are unbounded (i.e., C is larger than the total load of the VMs that can ever be in the system at any time, or m is larger than the number of VMs that can ever be in the system at any time). $(C, \cdot)$-VMA and $(\cdot, m)$-VMA are the cases where the number of available PMs and their capacity is unbounded, respectively.

2.2.2 Related Work

To the best of our knowledge, previous work on VMA has been only experimental [34,71,74, 90] or has focused on different cost functions [7,23,32,37]. First, we provide an overview of previous theoretical work for related assignment problems (storage allocation, scheduling, network design, etc.). The cost functions considered in that work resemble or generalize the power cost function under consideration here. Secondly, we overview related experimental work.

Chandra and Wong [32], and Cody and Coffman [37] study a problem for storage allocation that is a variant of $(\cdot, m)$-VMA with $b = 0$ and $\alpha = 2$. Hence, this problem tries to minimize the sum of the squares of the machine-load vector for a fixed number of machines. They study the offline version of the problem and provide algorithms with constant approximation ratio. A significant leap was taken by Alon *et al.* [7], since they present a PTAS for the problem of minimizing the L_p norm of the load vector, for any $p \geq 1$. This problem has the previous one as special case, and is also a variant of the $(\cdot, m)$-VMA problem when $p = \alpha$ and $b = 0$. Similarly, Alon *et al.* [8] extended this work for a more general set of functions, that include $f(\cdot)$ as defined above. Hence, their results can be directly applied in the $(\cdot, m)$-VMA problem. Later, Epstein *et al.* [42] extended [8] further for the uniformly related machines case. We will use these results in Chapter 4 in the analysis of the offline case of $(\cdot, m)$-VMA and $(\cdot, \cdot)$-VMA.

Bansal, Chan, and Pruhs minimize arbitrary power functions for speed scaling in job scheduling [15]. The problem is to schedule the execution of n computational jobs on a *single* processor, whose speed may vary within a countable collection of intervals. Each job has a release time, a processing work to be done, a weight characterizing its importance, and its execution can be suspended and restarted later without penalty. A scheduler algorithm must specify, for each time, a job to execute and a speed for the processor. The goal is to minimize the weighted sum of the flow times over all jobs plus the energy consumption, where the flow time of a job is the time elapsed from release to completion and the energy consumption is given by s^α where s is the processor speed and $\alpha > 1$ is some constant. For the online algorithm *shortest remaining processing time first*, the authors prove a $(3 + \epsilon)$ competitive ratio for the objective of total weighted flow plus energy. Whereas for the online algorithm *highest density first (HDF)*, where the density of a job is its weight-to-work ratio, they prove a $(2 + \epsilon)$ competitive ratio for the objective of fractional weighted flow plus energy.

Recently, Im, Moseley, and Pruhs studied online scheduling for general cost functions of the flow time, with the only restriction that such function is non-decreasing [61]. In their model, a collection of jobs, each characterized by a release time, a processing work, and a weight, must be processed by a *single* server whose speed is variable. A job can be suspended and restarted later without penalty. The authors show that HDF is $(2+\epsilon)$-speed $O(1)$-competitive against the optimal algorithm on a unit speed-processor, for all non-decreasing cost functions of the flow time. Furthermore, they also show that this ratio cannot be improved significantly proving impossibility results if the cost function is not uniform among jobs or the speed cannot be significantly increased.

A generalization of the above problem is studied by Gupta, Krishnaswamy, and Pruhs in [53]. The question addressed is how to assign jobs, *possibly fractionally*, to unrelated parallel machines in an online fashion in order to minimize the sum of the α-powers of the machine loads plus the assignment costs. Upon arrival of a job, the algorithm learns the increase on the load and the cost of assigning a unit of such job to a machine. Jobs cannot be suspended and/or reassigned. The authors model a greedy algorithm that assigns a job so that the cost is minimized as solving a mathematical program with constraints arriving online. They show a competitive ratio of α^α with respect to the solution of the dual program which is a lower bound for the optimal. They also show how to adapt the algorithm to integral assignments with a $O(\alpha)^\alpha$ competitive ratio, which applies directly to our $(\cdot, m)$-VMA problem. References to previous work on the particular case of minimizing energy with deadlines can be found in this paper.

Similar cost functions have been considered for the minimum cost network-design problem. In this problem, packets have to be routed through a (possibly multihop) network of speed scalable routers. There is a cost associated to assigning a packet to a link and to the speed or load of the router. The goal is to route all packets minimizing the aggregated cost. In [9] and [10] the authors show offline algorithms for this problem with undirected graph and homogeneous link cost functions that achieve polynomial and poly-logarithmic approximation, respectively. The cost function is the α-th power of the link load plus a link assignment cost, for any constant $\alpha > 1$. The same problem and cost function is studied in [53]. Bansal *et al.* [16] study a minimum-cost virtual circuit multicast routing problem with speed scalable links. They give a polynomial-time $O(\alpha)$-approximation offline algorithm and a polylog-competitive online algorithm, both for the case with homogeneous power functions. They also show that the problem is APX-hard in the case with heterogeneous power functions and there is no polylog-approximation when the graph is directed. Recently, Antoniadis *et al.* [11] improved the results by providing a simple combinatorial algorithm that is $O(\log^\alpha n)$-approximate, from which we can construct an $\widetilde{O}(\log^{3\alpha+1} n)$-competitive online algorithm. The $(\cdot, m)$-VMA problem can be seen as a especial case of the problem considered in these papers in which there are only two nodes, source and destination, and m parallel links connecting them.

To the best of our knowledge, the problem of minimizing the power consumption (given in Eq.2.1) with capacity constraints (i.e., the (C, m)-VMA and $(C, \cdot)$-VMA problems) has received very limited attention, in the realm of both VMA and network design, although the approaches in [10] and [16] are related to or based on the solutions for the capacitated network-design problem [31].

The experimental work related to VMA is vast and its detailed overview is out of the scope of this paper. Some of this work does not minimize energy [29, 72, 75] or it applies to a model different than ours (VM migration [80, 87], knowledge of future load [73, 87], feasibility of allocation [23], multilevel architecture [63, 74, 80], interconnected VMs [26], etc.). On the other hand, some of the experimental work where minimization of energy is evaluated focus on a more restrictive cost function [63, 93, 97].

In [63], for an energy cost model that is linear, the authors evaluate experimentally the allocation of VMs to clusters following 9 placement policies, some of them included in popular cloud platforms [44, 84]. Namely, Round Robin, Striping, Packing, Load Balancing (free CPU count), Load Balancing (free CPU ratio), Watts per Core, Cost per Core. We adapt 5 of these policies (defined later in Chapter 4) to our model and cost function for the purpose of simulations.

In [87], the authors focus on an energy-efficient VM placement problem with two requirements: CPU and disk. These requirements are assumed to change dynamically and the goal is to consolidate loads among servers, possibly using migration at no cost. In our model VMs assignment is based on a CPU requirement that does not change and migration is not allowed. Should any other resource be the dominating energy cost, the same results apply for that requirement. Also, if loads change and migration is free, an offline algorithm can be used each time that a load changes or a new VM arrives. In [87] it is shown experimentally that energy-efficient VMA does not merely reduce to a packing problem. That is, to minimize the number of PMs used even if their load is close to their maximum capacity. For our model, we show here that the optimal load of a given server is a function only of the fixed cost of being active (b) and the exponential rate of power increase on the load (α). That is, the optimal load is not related to the maximum capacity of a PM.

Part II

Understanding and Reducing Energy Consumption in Data Centers

Chapter 3

Analysis of the Energy Consumption of Data Center Servers

3.1 Overview

The work presented in this chapter is motivated by our disagreement with some of the models that have been previously proposed in the literature which state that power consumption of data center servers depends linearly on the load. Our belief is that more complex/complete models for the power consumed by a server are necessary. In order to be consistent, these models have to be based on empirical values. However, we found that, despite the large body of work in the field, there is a lack of empirical work studying servers energy behavior.

Our work tries to partially fill this void by proposing a measurement-based characterization — which is the first of its kind— of the energy consumption of a server components with DVFS and multiple cores. We evaluate here different server machines and evaluate what is the contribution to their power consumption of the CPU, hard drive disk, and network card (NIC). Our approach captures the influence of the processing frequency and the multiple cores, not only to the CPU power consumption, but also to that of disk input/output (I/O) and NIC activity.

Our contribution is threefold: (i) we propose a methodology to empirically characterize the energy consumption of a server, (ii) we provide novel, experimental-based, insights on the energy consumption behavior of the most relevant server's components, and (iii) we propose an accurate technique to estimate the energy consumption of cloud applications.

As concerns the methodology, we propose *active CPU cycles per second* (ACPS) as a new and more convenient metric for CPU load in multi-core/multi-frequency architectures. We show how to isolate the contribution of energy consumption due to CPU, disk I/O operations, and network activity by just measuring server's total energy consumption and a few activity indicators reported by the operating system. We also show that the *baseline* energy consumption of a server — i.e., the energy consumed just because the server is turned on — has a strong impact on server's total consumption.

As concerns the components' energy characterization, we show that, besides the *baseline* consumption, the CPU has the largest impact among all components, and its energy consumption is not linear with the load. Disk I/O operations are the second highest cause of consumption, and their efficiency is strongly affected by the I/O block size used by the application. Eventually, network activity plays a minor yet not negligible role in the energy consumption, and the network impact scales almost linearly with the network transmission rate. All other components (e.g., memory, fans, GPU, etc.) can be accounted for the *baseline* energy consumption, which is subject to minor variations under different operational conditions. Specifically, the main results of our measurement campaign are listed below:

- The CPU power utilization depends on the number of working cores, the CPU frequency, and the CPU load (in ACPS units). Our measurements confirm that the energy consumption with a single working core at constant frequency can be closely approximated by a linear function of the CPU load. However, given a CPU frequency, the energy consumption in multicore architectures is a concave function of the CPU load and can be approximated by a low-order polynomial. The energy consumption for a fixed CPU load is, in general, minimized by using the highest number of cores and the lowest frequency at which the load can be served. However, the minimum achievable energy consumption is a piecewise concave function of the CPU load.

- The energy consumed by hard disks for reading and writing depends on the CPU frequency and the I/O block sizes. Both reading and writing energy costs increase slightly with the CPU frequency. While the energy consumption due to reading is not affected by block size, the energy consumption due to writing increases with the block size. The reading efficiency (expressed in MB/J) is barely affected by the CPU frequency, while writing efficiency is a concave function of the block size since it boosts the throughput of writing until a saturation value is reached.

- The energy consumption and the efficiency of the NIC, both in transmission and reception, depends on the CPU frequency, the packet size, and the transmission rate. The efficiency of data transmission increases almost linearly with the transmission rate, with steeper slopes corresponding to lower CPU frequencies. Although a linear relation between transmission rate and efficiency holds for data reception as well, small packet sizes yield higher efficiency in reception.

Overall, supported by our measurements, we provide a holistic energy consumption model that only requires a few calibration parameters for every different server architecture which we want to evaluate (a universal energy model will be too simplistic and inaccurate). We validate our model by means of a server computing the *PageRank* metric of a graph and a *WordCount* application in a *Hadoop* platform, first without network activity, next with bulky network activity, and finally in the cloud. We will find that the error of our energy estimates is below 4.1% on average and never worse than a 10%.

RoadMap The rest of the chapter is organized as follows. Section 3.2 describes the methodology we used for our experiments. Section 3.3 presents our measurement campaign, for every single component which we tested. In Section 3.4 we model the energy consumption of the servers based on a few calibration parameters which we find during our measurement campaign. In Section 3.5 we discuss our findings and their implications. Finally, Section 3.6 concludes the chapter.

3.2 Methodology

In this section we introduce the measurement techniques we used to characterize the power requirements of CPU activity, disk access (read and write operations), and network activity. Our measurements start characterizing the CPU power consumption, from where we obtain information about the baseline power consumption of the system. After CPU and baseline characterization, we follow with experiments for the other two components, namely, disk and network. Note that CPU and baseline measurements are of capital importance in order to evaluate the other components, because any operation run in a machine is like a puzzle with multiple pieces and we must know what is the contribution of each one of these pieces. Consider that, we are paying a cost just for having a server switched on and the operating system running on it. Similarly, every time we run a task in the system, some CPU cycles are needed in order to execute it as well as to use the component that has to perform the task. Hence, in order to understand the contribution of any component, we first need to identify the contribution of the CPU and compute the difference with respect to the aforementioned baseline.

To explore the possible parameters which determine the power consumption of a server and to obtain statistical consistency, we run our experiments multiple times. Similarly, we run these experiments in different servers and architectures in order to validate our results and give consistency to our conclusions.

3.2.1 Collecting System Data and Fixing Frequency Parameters

One prerequisite for our experiments is to have Linux machines due to the kind of commands and benchmarks we wanted to use and, mainly, because of the possibility of adding some kernel modules and utilities,[1] which allow us to change CPU frequencies at will. In a Linux system, CPU activity stats are constantly logged, so we can periodically record the core frequency and the number of *active* and *passive* CPU ticks at each core.[2] Once we have the number of ticks and

[1] For instance cpufrequtils, acpi-cpufreq.

[2] File `/proc/stat` reports the number of ticks since the computer started devoted to *user*, *niced* and *system* processes, waiting (*iowait*), processing interrupts (i.e., *irq* and *softirq*), and *idle*. In our experiments we count both

the core frequency, since a tick represents a hundredth of second, cycles can be calculated as 100 *ticks/frequency*.

We use active cycles per second (ACPS) instead of CPU load percentage to characterize CPU load. ACPS depend on the CPU frequency used, as the higher the frequency the more the work that can be processed. In contrast, CPU load percentages cannot be compared when different frequencies are used, while the amount of ACPS that can be processed can be considered as an absolute magnitude. In order to get (set) information about the operative frequency of the system we used the `cpufrequtils` package.[3] With those tools, we can monitor the CPU frequency at which the system works and assign different frequencies to the cores. However, to limit the number of possible combinations to characterize, we assign the same frequency to all cores.

3.2.2 CPU

In order to evaluate the CPU power utilization we prepared a script based on a benchmark application, namely `lookbusy`.[4] Note that `lookbusy` allows us to load one or more CPU cores with the same load.Our `lookbusy`-based experiment follows the next steps: we first fix the CPU frequency to the lowest possible frequency in the system; then we run `lookbusy` with fixed amount of load for one core during timeslots of 30 seconds, starting with the maximum load and then decreasing the load gradually. After the last `lookbusy` run we measure the power used during an additional timeslot with *no* `lookbusy` load offered. We register the active cycles and the power used during each timeslot.

After taking these different samples for one frequency we move to the immediately higher frequency (we can list and change frequencies thanks to `cpufrequtils`) and repeat the previous steps. After going through all the available frequencies, we restart the whole process but increasing by one the number of active cores. We repeat this whole process until all the cores of the server are active. Note that when we change the frequency of the cores we change it in all of them, active or not, for consistency. Similarly, when more than one core is active, the load for all the active cores is the same.

Once explained the scheme of our experiments, we must clarify the meaning of running a timeslot with *no* load. Note that zero-load is clearly not possible as there is always going to be load in the system due to, e.g., the operating system. However, during the timeslot in which we do not run `lookbusy`, we measure the power corresponding to the operational conditions which are as close as possible to the ones of an idle system. Moreover, the decision of using timeslots of 30 seconds is to guarantee enough, yet not excessive, time for the measurements. In fact, as we start and stop `lookbusy` at the beginning and end of the timeslots, we need to ignore the first and the last few seconds of measurements in each timeslot to avoid measurement noise due to power ramps and operational transitions.

waiting and idle ticks as *passive* ticks, while we denote the aggregated value of the rest of ticks as *active*.

[3] `https://wiki.archlinux.org/index.php/CPU_Frequency_Scaling`

[4] `http://www.devin.com/lookbusy`.

The measured values of load (in ACPS) and power in each timeslot are used to obtain a least squares polynomial fittings curve. These fittings characterize the CPU power utilization for each combination of frequency and number of active cores. We will use as *baseline power utilization* of each one of these configurations the zero-order coefficient of the polynomial of these fittings curves.

3.2.3 Disks

The power consumption of the hard drive was evaluated using 2 different scripts (for reading and writing) based on the `dd` linux command.[5] We chose `dd` as it allows us to read files, write files from scratch, control the size of the blocks we write (read), control the amount of blocks written (read) and force the commit of writing operations after each block in order to reduce the effect of operating system caches and memory. We combine this tool with flushing the RAM and caches after each reading experiment.

In both our scripts we perform write (read) operations for a set of different I/O block sizes and for different data volumes to be written (read). In each case we record the CPU active cycles, the total power and time used in each one of these operations for each combination of block size and available frequency.

Finally, we identify the contribution of the hard drive to the total power utilization by subtracting the contribution of both the baseline and the CPU requirements from the measured total power.

Disk I/O experiments shed light on the relevance of the block sizes when reading or writing as well as whether there is an influence of the frequency on these operations.

3.2.4 Network

In order to evaluate the contribution of the network to the power requirements of a server, we devised a set of experiments based on a client-server C script devised on purpose for this task.

There are several aspects that we consider relevant in order to characterize the impact of the NIC on the total power requirements of a server and that led us to choose these two tools. First, the ability of performing tests in which the server under study acts as sender or as receiver during a network connection, and therefore we can observe server's power requirements while sending data or receiving it. To clarify the terms, *sender* is the server which injects traffic to the network, and *receiver* is the server which accepts traffic from the network. Second, the ability of those tools to change several parameters that we consider relevant for the energy characterization of the servers, namely, the packet size and the offered load, jointly with the frequency of the system.

Our experiments consist, then, on measuring the achieved data rate, the CPU active cycles per second (ACPS) and the total power required by the server under study either as sender or as

[5] `http://linux.die.net/man/1/dd`.

Table 3.1: Characteristics of the servers under study

Component	Servers		
	Survivor	Nemesis	Erdos
CPU (# cores)	4	4	64
# freqs	8	11	5
Freqs List (*GHz*)	1.2, 1.333, 1.467, 1.6, 1.733, 1.867, 2, 2.133	1.596, 1.729, 1.862, 1.995, 2.128, 2.261, 2.394, 2.527, 2.666, 2.793, 2.794	1.4, 1.6, 1.8, 2.1, 2.3
RAM	4 *GB*	4 *GB*	512 *GB*
Disk	2 *TB*	$2+3$ *TB*	2×146*GB* 4×1 *TB*
Network	1 *Gbps*	3×1 *Gbps*	4×1 *Gbps* 2×10 *Gbps*
Architecture	Intel	Intel	AMD

receiver using different packet sizes and different transfer rates. We run each experiment multiple times for statistical consistency.

Finally, using the CPU active cycles per second which were measured during the experiment, we identify the power required by the CPU. Subtracting both CPU power requirements and the baseline power from the total energy consumption of the experiment, we can isolate the power requirements of the network.

3.3 Measurements

3.3.1 Devices and Setup

In order to monitor and store the instantaneous power required by a server during the different experiments we used a Voltech PM1000+ power analyzer[6], which is able to measure the total instantaneous power required by the server under test on a per-second basis. In Fig. 3.1 we show a schematic representation of the setup we used and the components under testing. More specificaly, in order to take our measurements we connected the server being measured to the power analyzer and the latter to the power supply.In the case of servers with power redundancy one of the two power sources was unplugged to ensure that the power measurement was correct. In the experiments where the network was not involved (CPU and disk), we disconnected the server from the network, which has an impact on the power requirements as the port goes idle. In the network based experiments we established an Ethernet connection between the server under study and a second machine in order to study the server behavior, both as receiver as well as as sender.

We evaluated three different servers: `Survivor`, `Nemesis`, and `Erdos`. We will now

[6] More information about the PM1000 can be found in `http://www.farnell.com/datasheets/320316.pdf`

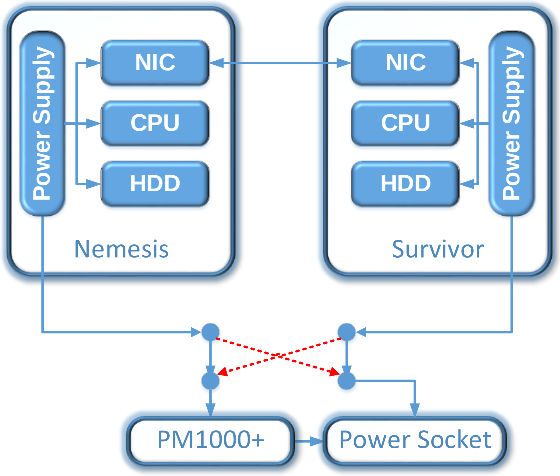

Figure 3.1: Schematic representation of our setup when Nemesis is being measured. Red arrows show the alternative scheme to measure Survivor (or Erdos).

present these servers although their main characteristics, including their sets of available CPU frequencies, can be also found in Table 3.1. `Survivor` has an Intel Xeon E5606 4-core processor, with 4 *GB* of RAM, a 2 *TB* Seagate Barracuda XT hard drive and a 1 *Gigabit* Ethernet card integrated in the motherboard. `Nemesis` is a Dell Precision T3500 with an Intel Xeon W3530 4-core processor, 4 *GB* of RAM, 2 hard drives (a 2 *TB* Seagate Barracuda XT and a 3 *TB* Seagate Barracuda), a 1 *Gigabit* Ethernet card integrated in the motherboard, and a separate Ethernet card with two 1 *Gigabit* ports. In this study we only evaluate the Seagate Barracuda XT disk and the integrated Ethernet card. Both `Survivor` and `Nemesis` use the Ubuntu Server edition 10.4 LTS Linux distribution. Finally, `Erdos` is a Dell PowerEdge R815 with 4 AMD Opteron 6276 16-core processors (i.e., 64 cores in total), 512 *GB* of RAM, two 146 *GB* SAS hard drives configured as a single RAID1 system (which is the “disk” analyzed here) and four 1 *TB* Near-line SAS hard drives. It also includes four 1 *Gigabit* and two 10 *Gigabit* ports. `Erdos` is a high-end server and uses Linux Debian 7 Wheezy.

3.3.2 Baseline and CPU

As we mentioned in Section 3.2, for each server we have measured the power it uses with neither disk accesses nor network traffic. We assume that the power utilization observed is the sum of the baseline consumption plus the power used by the CPU. We have obtained samples of the power consumed under different configurations that vary in the number of active cores used, the frequency at which the CPU operates (all cores operate at the same frequency), and the active cores load (all active cores are equally loaded). The list of available and tested CPU frequencies and cores can be found in Table 3.1. We tune the total load ρ by using `lookbusy`, as described in the previous section. Each experiment lasts 30 s and it is repeated 10 times. Results are summarized in terms of average and standard deviation. Specifically, in the figures reported in this section, the power utilization for each tested configuration is depicted by means of a vertical segment centered on the average power utilization measured, and with segment size equal to two times the standard deviation of the samples.

The results of these experiments for each of the 3 servers are presented in Figure 3.2 (the measurements for some frequencies and some number of cores are omitted for clarity). Here, for each configuration of number of active cores, frequency, and load in ACPS, the mean and standard deviation of all the experiments with that configuration are presented. Also the least squares polynomial fitting curve for the samples is shown for each number of cores and frequency. The curves shown are for polynomials of degree 7, but we observed that using a degree 3 polynomial instead does not reduce drastically the quality of the fit (e.g., the relative average error of the fitting increases from 0.7% with 7-th degree polynomials to 1.5% with degree equal to 3 for `Erdos`, while it remains practically stable and below 0.7% for `Nemesis`). In general, we can use an expression like the following to characterize the CPU power consumption:

$$P_{BC}(\rho) = \sum_{k=0}^{n} \alpha_k \rho^k, \quad n \leq 7, \tag{3.1}$$

where P_{BC} includes both the baseline power utilization of the servers and the power used by the CPU, and ρ is the load expressed in active cycles per second. Therefore, coefficient α_0 in Eq. 3.1 represents the consumption of the system when the CPU activity tends to 0, and we can thereby interpret α_0 as the baseline power utilization of the system. Note that the polynomial fitting, and hence the baseline power utilization α_0, depends on the particular combination of number of cores and frequency adopted. However, for sake of readability, we do not explicitly account for such a dependency in the notation.

A first observation of the fitting curves for each particular server in Figure 3.2 reveals that the power for near-zero load is almost the same in curves (e.g., for `Nemesis` this value is between 84 and 85 W). Observe that it is impossible to run an experiment in which the load of the CPU is actually zero to obtain the baseline power utilization of a server. However, all the fitting curves converge to a similar value for $\rho \to 0$, which can be assumed to represent the baseline power utilization.

A second observation is that for one core the curves grow linearly with the load. However, as soon as two or more cores are used, the curves are clearly concave, which implies that for a fixed frequency the efficiency grows with the load (we will discuss later the efficiency in terms of number of active cycles per energy unit).

A third observation is that frequency does not significantly impact the power consumption when the load is low. In contrast, at high load, the power clearly increases with the CPU frequency. More precisely, the power consumption grows superlinearly with the frequency, for a fixed load and number of cores. This is particularly evident in the curves characterizing `Erdos`, which is the most powerful among our servers.

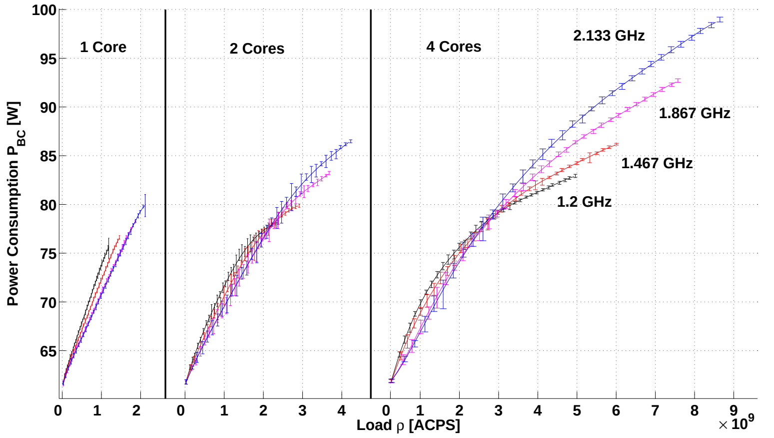

(a) `Survivor`

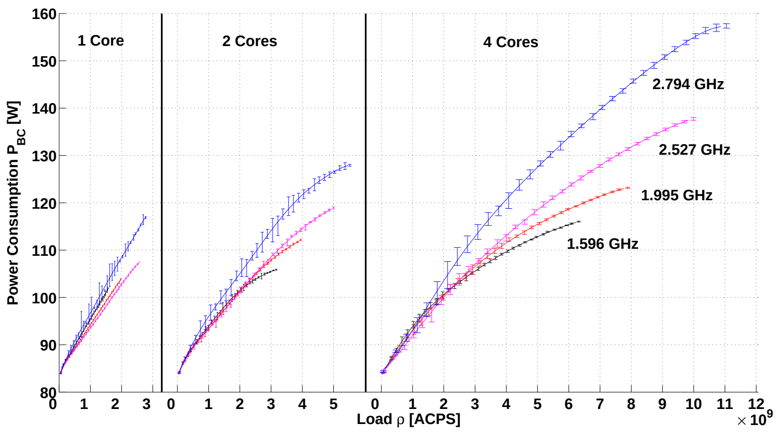

(b) `Nemesis`

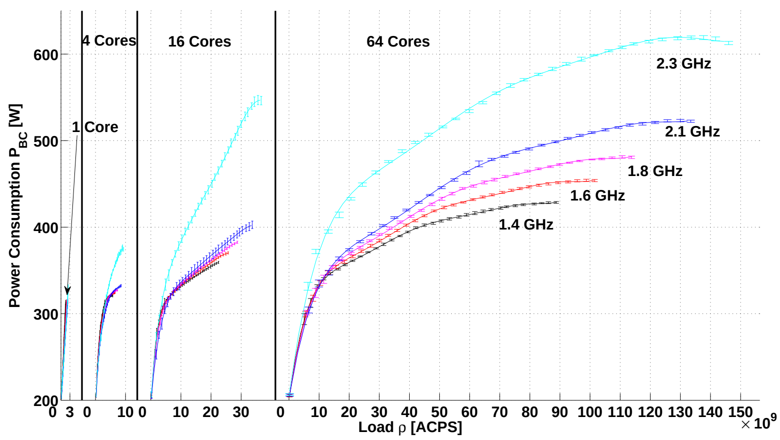

(c) `Erdos`

Figure 3.2: Power consumption of 3 servers (`Survivor`, `Nemesis`, and `Erdos`) for baseline and CPU characterization experiments.

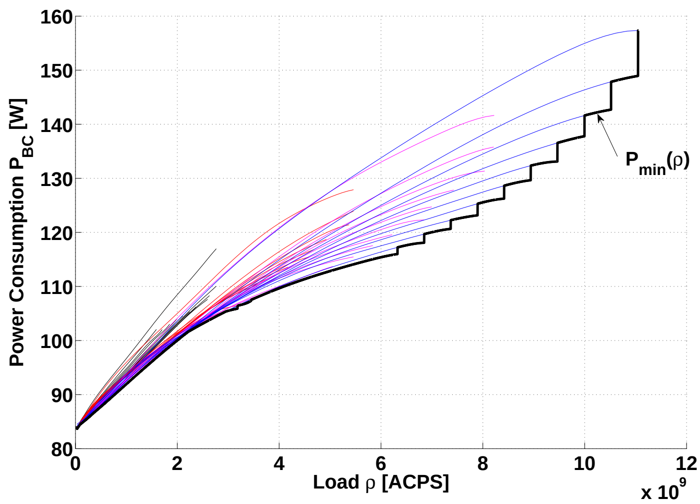

(a) Minimal power.

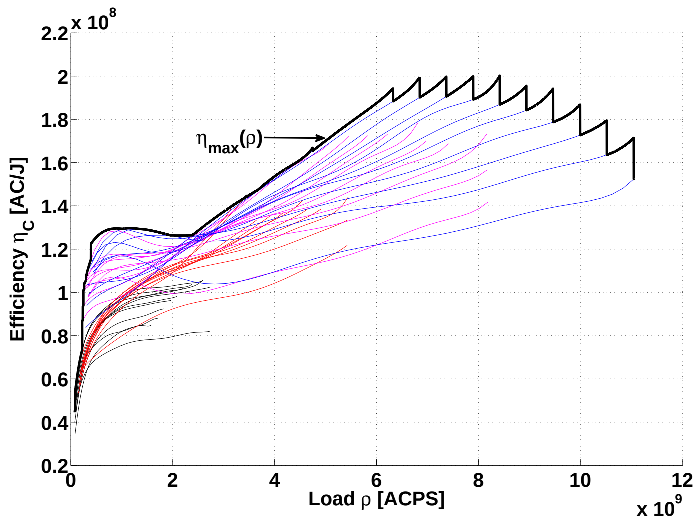

(b) Maximal efficiency.

Figure 3.3: CPU performance bounds of `Nemesis`.

From the previous figures it emerges that the power consumption due to CPU and baseline can be minimized by selecting the right number of active cores and a suitable CPU frequency. Similarly, we can expect that the energy efficiency, defined as number of active cycles per energy unit, can be maximized by tuning the same operational parameters. We graphically represent the impact of operation parameters on power consumption and energy efficiency in Figures 3.3 and 3.5 respectively for `Nemesis` and `Erdos` (results for `Survivor` are similar to the ones shown for `Nemesis` and are omitted).

In particular, Figures 3.3(a) and 3.5(a) report all possible fitting curves for the power consumption measurements, plus a curve marking the lowest achievable power consumption at a given

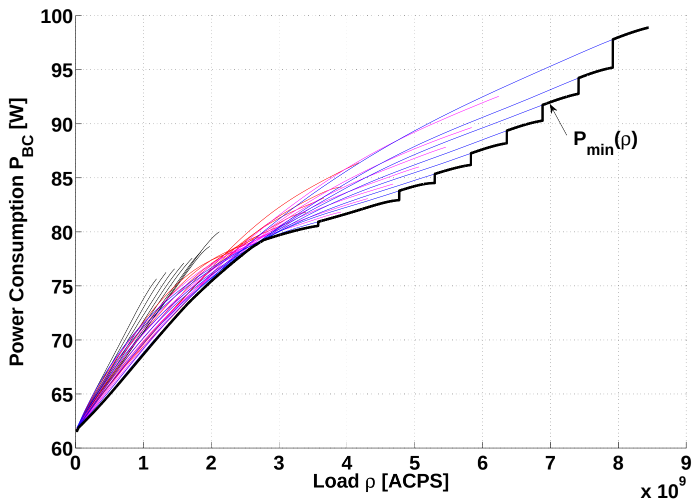

(a) Minimal power.

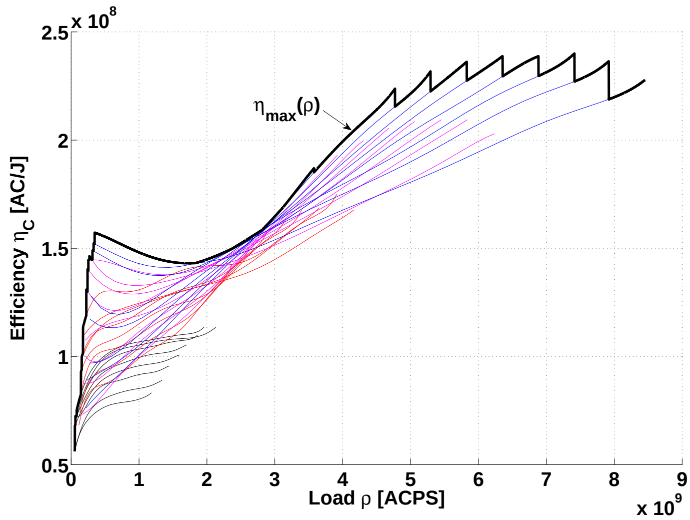

(b) Maximal efficiency.

Figure 3.4: CPU performance bounds of `Survivor`.

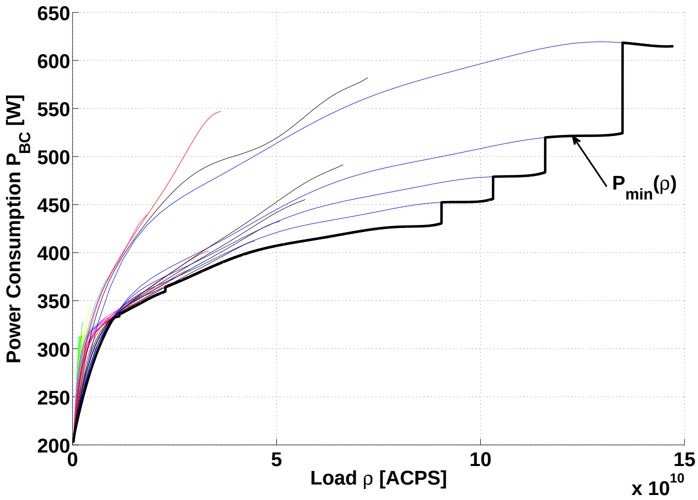

(a) Minimal power.

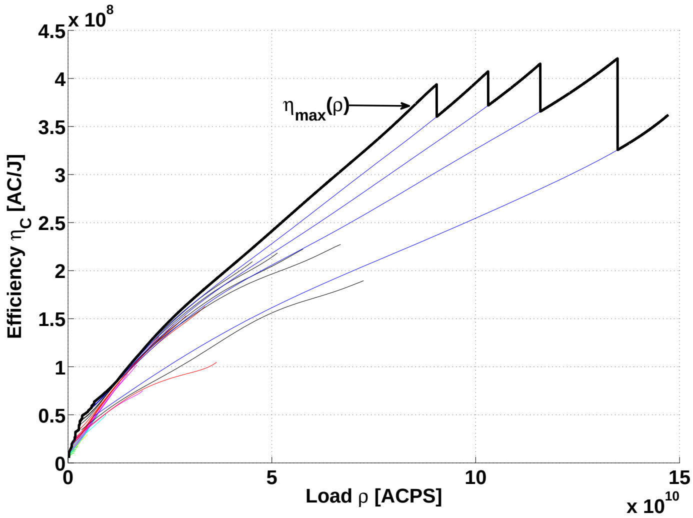

(b) Maximal efficiency.

Figure 3.5: CPU performance bounds of `Erdos`.

load. We name such a curve "minimal power curve" $P_{\min}(\rho)$, and we observe that (i) it only depends on the load ρ, and (ii) it is a piecewise concave function, which makes it suitable to formulate power optimization problems. Finally, to evaluate the energy efficiency of the CPU, we report in Figures 3.3(b) and 3.5(b) the number of active cycles per energy unit obtained from our measurements respectively for `Nemesis` and `Erdos`. We compute the power due to active cycles as the power $P_{BC} - \alpha_0$, i.e., by subtracting the baseline consumption from P_{BC}, and we obtain the efficiency η_C by dividing the load (in active cycles per second) by the power due to active cycles:

$$\eta_C = \frac{\rho}{P_{BC}(\rho) - \alpha_0}. \tag{3.2}$$

Also in this case we show the curve that maximizes the efficiency at a given load, which we name "Maximal efficiency curve" $\eta_{\max}(\rho)$. Interestingly, we observe that (i) $\eta_{\max}(\rho)$ presents multiple local maxima, (ii) for a given configuration of frequency and number of active cores, the efficiency is maximized at the highest achievable load, (iii) all local maxima corresponds to the use of all available active cores, but (iv) the absolute maximum is *not* achieved neither at the highest CPU frequency nor at the lowest.

3.3.3 Disks

We now characterize the power and energy consumption of disk I/O operations. During the experiments, we continuously commit either read or write operations, while keeping the CPU load ρ as low as possible (i.e., we disconnect the network and we do not run other tasks). Still, the power measurements obtained during the disk experiments contain both the power used by the disk and power due to CPU and baseline. Indeed, Figure 3.6 shows, for each experiment, the total measured power P_t, the power P_{BC} computed according to Eq. 3.1 at the load ρ measured during the experiment, and the power due to disk operations, computed as:

$$P_D^x = P_t - P_{BC}(\rho), \quad x \in \{r, w\}, \tag{3.3}$$

where superscripts r and w refer to reading and writing operations, respectively. We test sequentially all the available frequencies for each server (see Table 3.1), and I/O block sizes ranging from 10 *KB* to 100 *MB*. Figure 3.6 shows average and standard deviation of the measures over 10 experiment repetitions for each one of our servers. Indeed, it can be easily seen that `Survivor` and `Nemesis` have similar disks and file systems, while `Erdos` is equipped with SAS disks with RAID. In all cases shown in the figure, the disk power is small but not negligible with respect to the baseline consumption. Furthermore, we can observe that the two servers presented behave differently. Indeed, while the power consumption due to writing is affected both by the block size B for both machines, we observe that both `Nemesis` and `Survivor`' disk writing power P_D^w is not affected by the CPU frequency, while `Erdos`' results show an increase with the frequency. Moreover, the results obtained with `Erdos` are affected by a substantial amount of variability in the measurements, which we believe is due to the caching operations enforced by the RAID mechanism in `Erdos`.

Similarly to what was described for the CPU, we now comment on the energy efficiencies η_D^r and η_D^w of disk reading and writing operations. Figure 3.7 reports efficiency as a function of the I/O block size, and shows one line per each CPU frequency[7]. The efficiency is computed by subtracting the baseline power from the total power, and by measuring the volume V of data read or written in an interval T:

$$\eta_D^x = \frac{V}{P_D^x T}, \quad x \in \{r, w\}. \tag{3.4}$$

We can observe that results are similar for all the servers. Specifically, reading efficiency is almost constant at any frequency and for each block size, while writing is more efficient with large block sizes. We also observe that the efficiency changes very little with the adopted CPU frequency. Another observation is that the efficiency saturates to a disk-dependent asymptotic value, which is due to the mechanical constraints of the disk (e.g., due to the non-negligible *seek* time, the number of read/write operations per second is limited). In addition, although not visible in the figure due to the log-scale adopted, η_D^w is a concave function of the block size B.

3.3.4 Network

The last server component that we characterize via measurements is the network card. Similarly to the cases described previously, we run experiments in which only the operating system and our test scripts are active. In this case, we run a script to either transmit or receive UDP packets over a gigabit Ethernet connection and count the system active cycles ρ. We measure the total power consumption P_t during the experiment, so that the power due to network activity can be then estimated as follows:

$$P_N^x = P_t - P_{BC}(\rho), \quad x \in \{s, r\}, \tag{3.5}$$

where superscripts s and r refer to the sender and the receiver cases, respectively.

In the experiments, we sequentially test all the available frequencies for each server (see Table 3.1), and fix the packet size and the transmission rate within the achievable set of rates (which depends on the packet size, e.g., < 950 *Mbps* for 1470-B packets). We report results for the network energy consumption in terms of efficiencies η_N^s and η_N^r (volume of data transferred per unit of energy). These efficiencies are computed as follows:

$$\eta_N^x = \frac{R}{P_N^x}, \quad x \in \{s, r\}, \tag{3.6}$$

where R is the transmission rate during the experiment.

Figures 3.8, 3.9 and 3.10 show the network efficiencies of `Survivor`, `Nemesis` and

[7] For readability, results for `Survivor` are omitted.

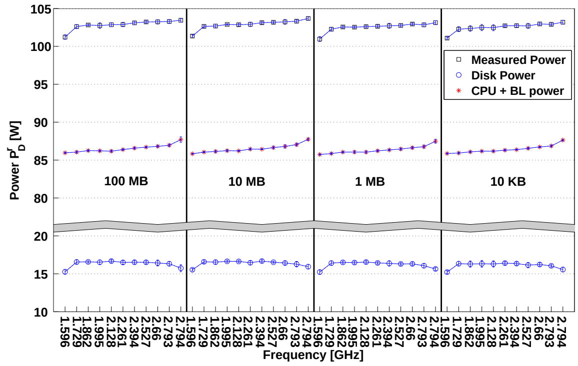

(a) Power consumption during reading (`Nemesis`).

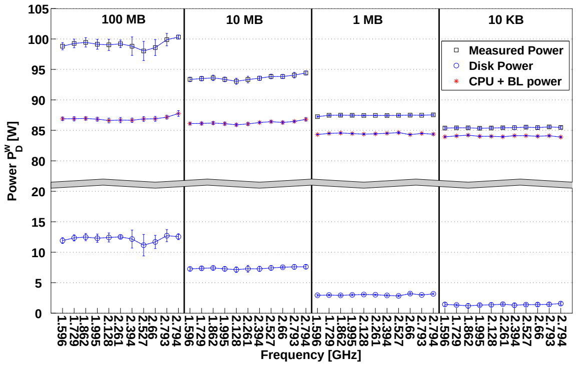

(b) Power consumption during writing (`Nemesis`).

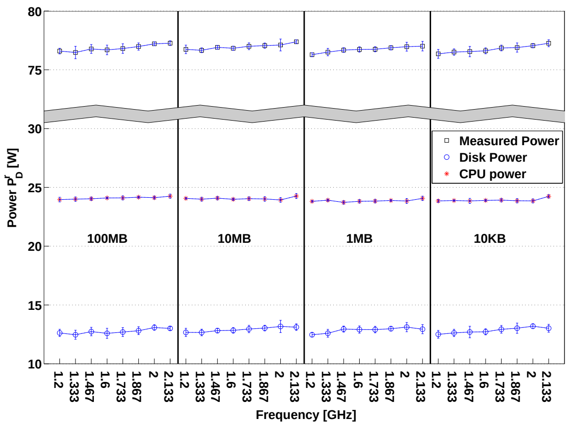

(c) Power consumption during reading (`Survivor`).

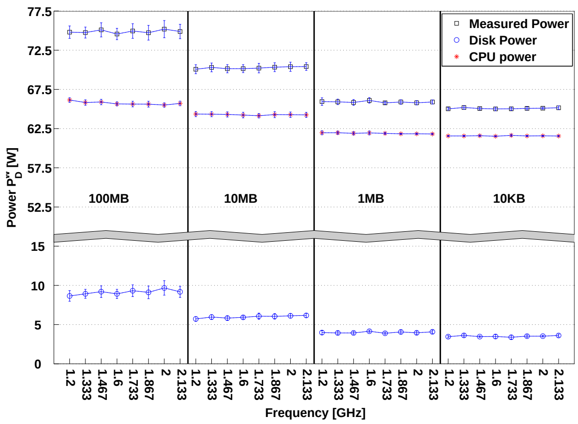

(d) Power consumption during writing (`Survivor`).

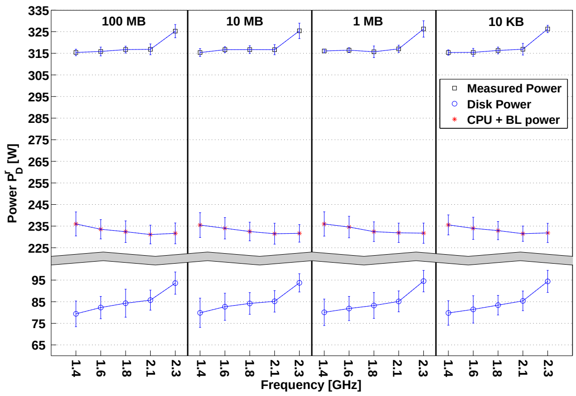

(e) Power consumption during reading (`Erdos`).

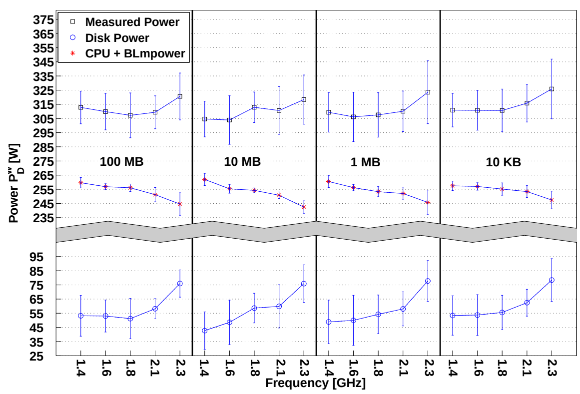

(f) Power consumption during writing (`Erdos`).

Figure 3.6: Instantaneous power consumption for reading/writing operations. Results are presented for every frequency and for 4 different block sizes for each one of our servers.

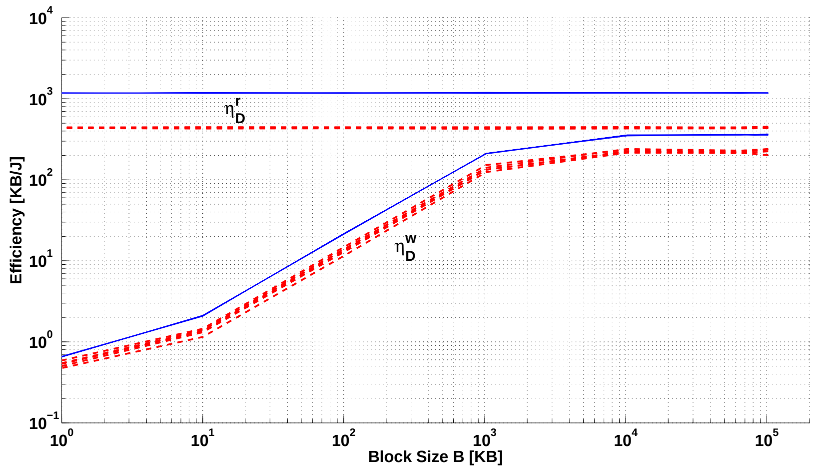

Figure 3.7: Disk reading and writing efficiencies for `Erdos` (red dotted lines) and `Nemesis` (blue solid lines).

`Erdos`, respectively, averaged over 5 samples per transmission rate R.[8] [9] For the sake of readability, the figures only show results for the biggest and smallest packet sizes, i.e., 64-B and 1470-B packets. For `Nemesis` and `Survivor` we report four CPU frequencies: the lowest, the highest, the most efficient (according to Figures 3.3(b) and 3.4(b)) and an intermediate one. For `Erdos` all five available frequencies are shown. The figure also reports the polynomial fitting curves for efficiency, which we found to be at most of second order. Since the efficiency is represented in terms of network activity only, in the fitting we force the zero-order coefficient of the polynomials to be 0. Therefore, we can use the following expression to characterize the network efficiencies of our servers:

$$\eta_N^x = \beta_1 R + \beta_2 R^2, \quad x \in \{s, r\}, \tag{3.7}$$

where the β_i coefficients are computed by minimizing the least square error of the fitting.

It can observed in Figures 3.8, 3.9 and 3.10 that efficiencies are almost linear or slightly superlinear with the transfer rate, e.g., the receiving efficiency of `Survivor` exhibits an evident quadratic behavior. Indeed, our measurements show that the network power consumption is independent from the throughput, which is a well known result for legacy Ethernet devices. In fact, the NICs of our servers are not equipped with power saving features like, e.g., the recently standardized IEEE 802.3az [60].

In all cases, the efficiency is strongly affected by the selected CPU frequency. Moreover, efficiency is also affected by packet size, although the impact of packet size changes from server to server, e.g., `Survivor` sending efficiency is only slightly affected by it.

Another observation is that, depending on the packet size and frequency used, sending can

[8] Network results are obtained by using a point-to-point Ethernet connection between two controlled servers.

[9] Due to technical and regulation reasons it was only possible to complete the *sender* part for `Erdos`, obtaining only partial results which, because of this partiality, are not published.

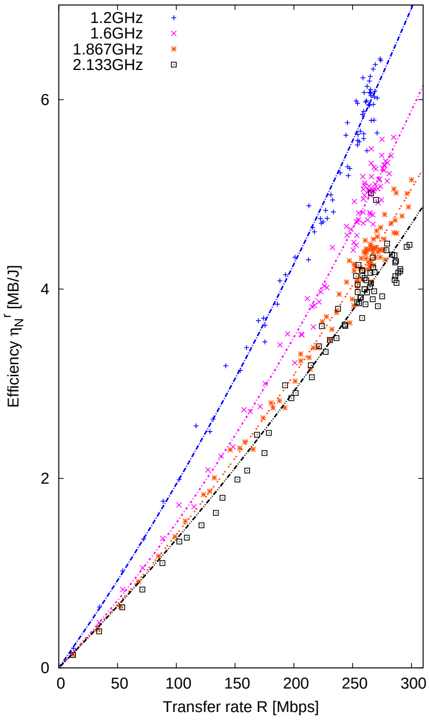

(a) Receiver efficiency in `Survivor` when using 64-B packets.

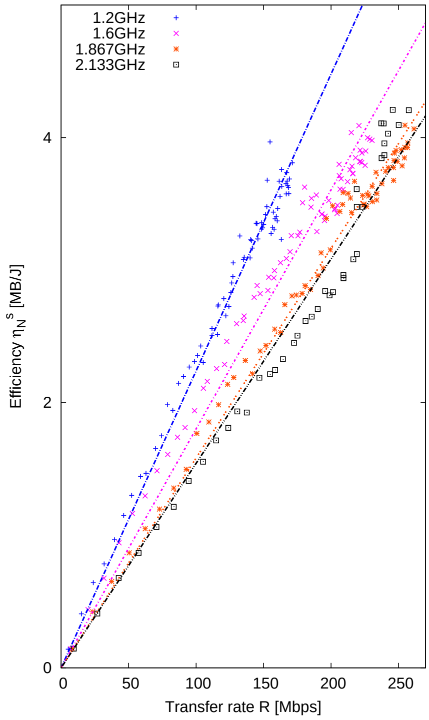

(b) Sender efficiency in `Survivor` when using 64-B packets.

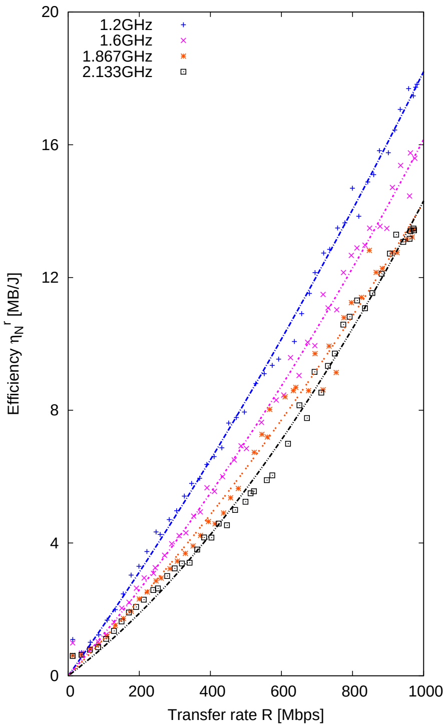

(c) Receiver efficiency in `Survivor` when using 1470-B packets.

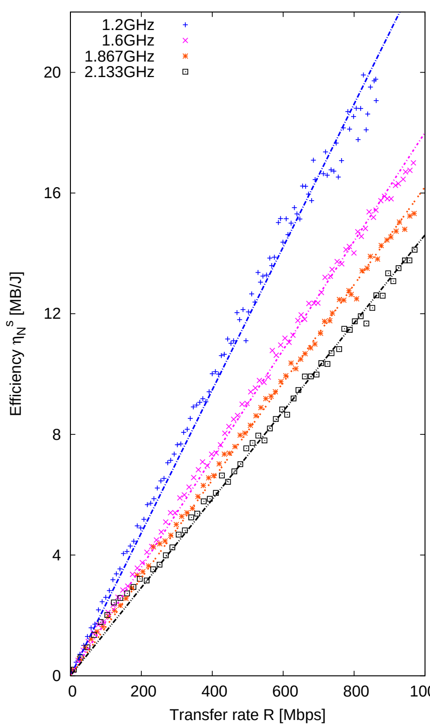

(d) Sender efficiency `Survivor` when using 1470-B packets.

Figure 3.8: Network efficiencies for `Survivor` under different frequencies and 64-B and 1470-B packets.

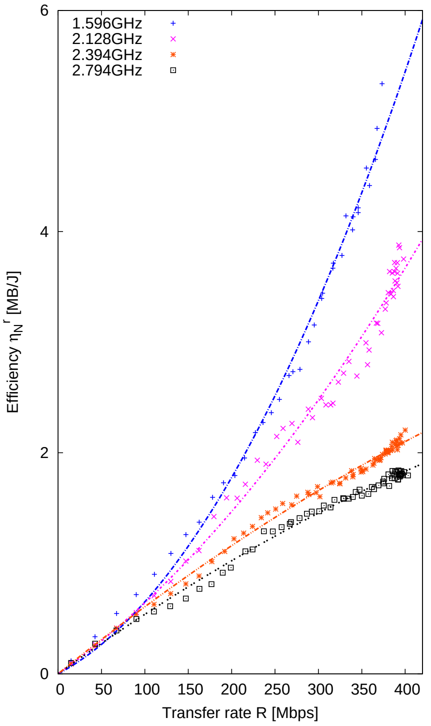

(a) Sender efficiency in `Nemesis` when using 64-B packets.

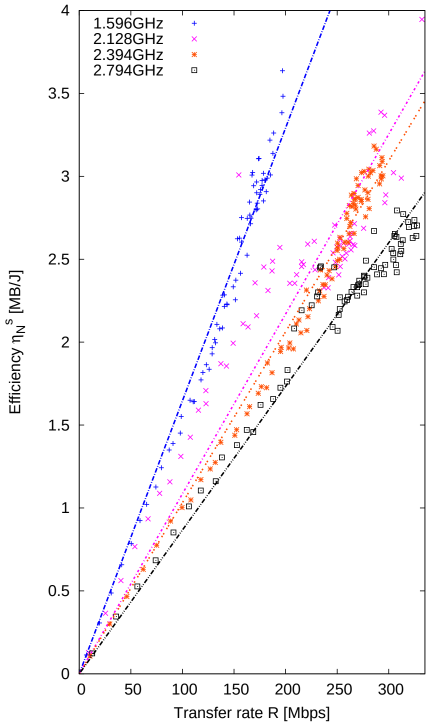

(b) Sender efficiency in `Nemesis` when using 64-B packets.

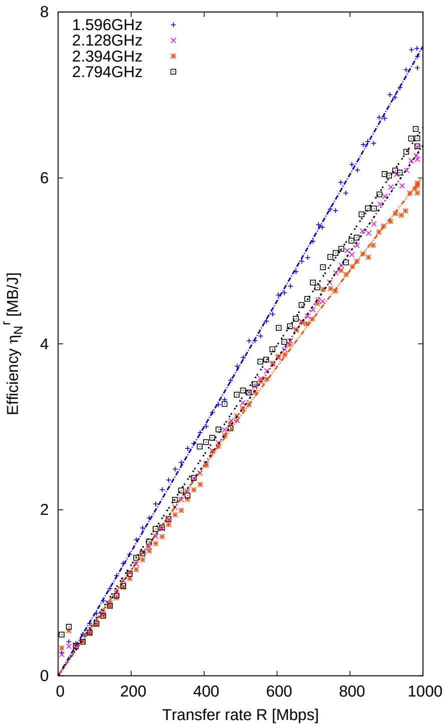

(c) Sender efficiency in `Nemesis` when using 1470-B packets.

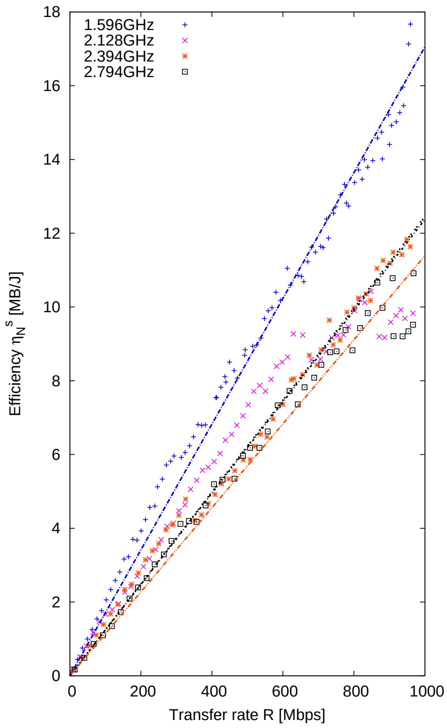

(d) Sender efficiency in `Nemesis` when using 1470-B packets.

Figure 3.9: Network efficiencies for `Nemesis` under different frequencies and 64-B and 1470-B packets.

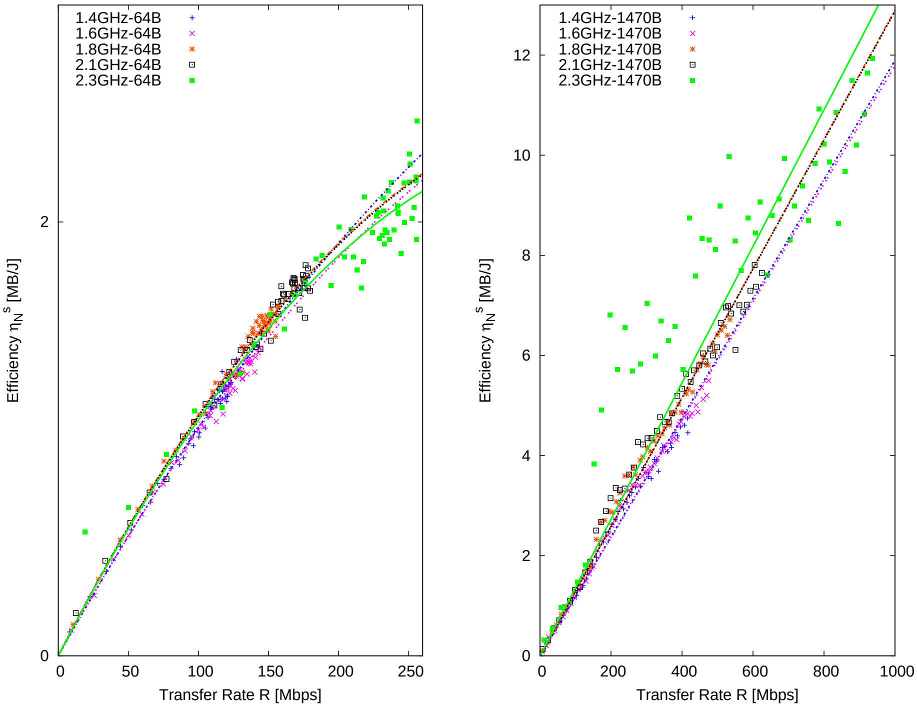

(a) Sender efficiency in `Erdos` when using 64-B packets. (b) Sender efficiency in `Erdos` when using 1470-B packets.

Figure 3.10: Network efficiencies for `Erdos` under different frequencies and 64-B and 1470-B packets.

be more energy efficient than receiving at a given transmission rate, and using the highest CPU frequency is never the most efficient solution. Note also that the efficiency decreases with the packet size, although this effect is particularly evident at the receiver side, while it only slightly impacts the efficiency of the packet sender. However, network activity also causes non-negligible CPU activity, as shown in Figure 3.11 for a few experiment configurations for all three servers. Overall, the lowest CPU frequency yields the lowest total power consumption during network activity periods.

3.4 Estimating Energy Consumption

While the results presented in the previous sections are useful to understand the energy consumption pattern of CPU, disk and network, we believe that a much more important use of these results is to estimate the energy consumption of applications. In this section we describe how this can be done from simple data about the application. Moreover, we validate the proposed approach by estimating the energy consumed by several *map-reduce Hadoop* computations.

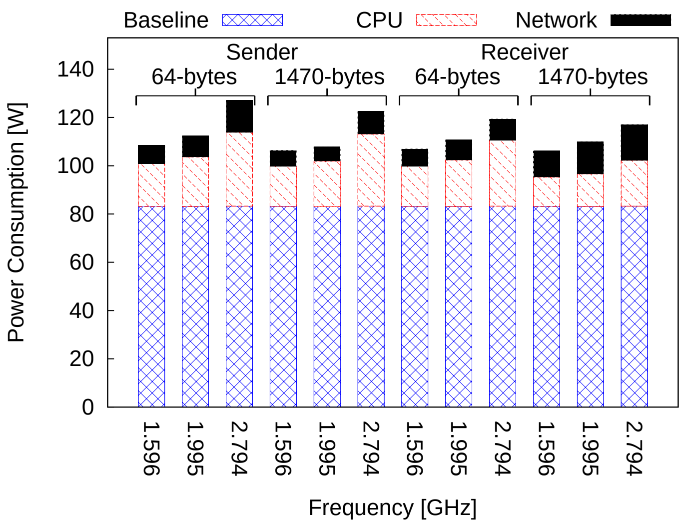

(a) Network efficiency (`Nemesis`).

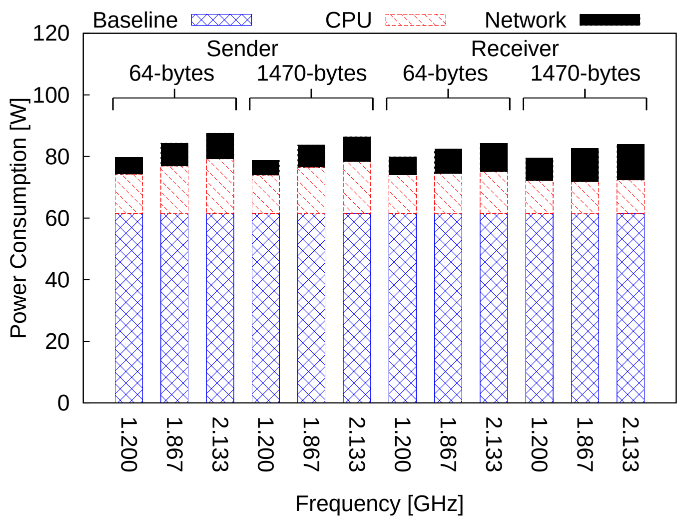

(b) Network efficiency (`Survivor`).

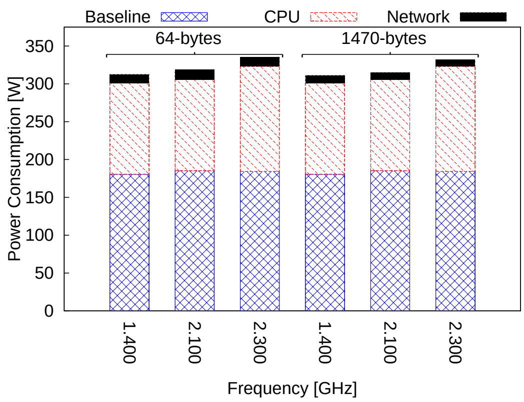

(c) Network efficiency (`Erdos`, only sender side).

Figure 3.11: Power utilization with network activity for `Erdos`, `Nemesis` and `Survivor` (64-B experiments were run with a transmission rate $R = 150$ *Mbps*, while $R = 400$ *Mbps* for the experiments with 1470-B packets).

3.4.1 Energy Estimation Hypothesis

The approach we propose to estimate the energy E_{app} consumed by an application lays on the basic assumption that the energy is essentially the sum of the baseline energy E_B (baseline power times application running time), the energy consumed by the CPU E_C, the energy consumed by the disk E_D, and the energy consumed by the network interface E_N:

$$E_{app} = E_B + E_C + E_D + E_N. \tag{3.8}$$

Hence, the process of estimating E_{app} is reduced to estimating these four terms. In order to estimate the first two terms, we need to know the total number of active cycles that the application will execute, C_{app}, and the load ρ_{app} (in ACPS) that the execution will incur in the CPU. From this, the total running time T_{app} can be computed as

$$T_{app} = C_{app}/\rho_{app}.$$

Then, once the number of cores and the frequency that will be used have been defined, it is also possible to estimate the baseline power plus CPU power, P_{BC}, from the fitting curves of Fig. 3.2. This allows to estimate the sum of the first two terms of Eq. 3.8 as

$$E_B + E_C = P_{BC}T_{app} = P_{BC}C_{app}/\rho_{app}. \tag{3.9}$$

The energy consumed by the disk is simply the energy consumed while reading and writing, i.e., $E_D = E_D^r + E_D^w$. To estimate these latter values, the block size to be used has to be decided, from which we can obtain an estimate of the efficiency of reading, η_D^r, and writing, η_D^w (see Figure 3.7). These, combined with the total volume of data read and written by the application, denoted as V_D^r and V_D^w respectively, allow to obtain the estimate energy as

$$E_D = \frac{V_D^r}{\eta_D^r} + \frac{V_D^w}{\eta_D^w}. \tag{3.10}$$

Finally, to estimate E_N, the transfer rate R and the packet size S have to be chosen, which combined with the frequency used, yield sending and receiving efficiencies η_N^s and η_N^r (see Figures 3.8, 3.9 and 3.10). Then, if the total volumes of data to be sent and received are V_N^s and V_N^r, respectively, the energy spent due to network is as follows:

$$E_N = \frac{V_N^s}{\eta_N^s} + \frac{V_N^r}{\eta_N^r}. \tag{3.11}$$

All is left to do to obtain the estimate E_{app} is to add up the values obtained in Equations 3.9, 3.10, and 3.11.

3.4.2 Applications and Scenarios for Validation

In this subsection we present the applications and scenarios we experimented with in order to validate the model presented in Section 3.4.1. Our goal was to be able to estimate the energy consumed by an application deployed on a data center based on the usage of its different components. For that, we executed two different Hadoop applications, PageRank and WordCount, in three different scenarios: first with an Isolated Server (no network), second with a server connected to the network, and finally with a two-server cloud. For the first two scenarios we used `Nemesis`, whereas, for the cloud case, we used both `Nemesis` and `Survivor`. We describe applications and scenarios in detail below.

Our first application is a **Hadoop Map-Reduce PageRank** based application that follows the approach from Castagna [30]. This application, that we denote PageRank for simplicity, computes several iterations of the pagerank algorithm on an Erdos-Renyi random (directed) graph with 1 million nodes and average degree 5 (Our PageRank algorithm assigns one input graph to each mapper so, in order to have one map task in each machine, two instances of this graph had to be used in the cluster scenario).

The execution of the PageRank application has three phases: preprocessing, map-reduce, and postprocessing. On its side, the map-reduce phase is a sequence of several homogeneous iterations of the PageRank algorithm that runs until a certain threshold is met. For simplicity, we only estimate the energy consumed during the map-reduce phase of the pagerank algorithm, which we force to run 10 times.

Our second application is the **Hadoop Map-Reduce WordCount**. This is a simple program that reads text files and counts how often words occur. For WordCount we use a few hundreds of books as input and estimate the energy consumed for the whole map-reduce process.

As we have mentioned above, these applications are run in 3 different scenarios. In the first scenario, denoted as **Isolated Server**, we run Hadoop in `Nemesis` keeping it disconnected from the network. When we run our applications in this scenario we are basically measuring the impact on the energy consumption of the baseline, CPU and hard disk.

In the second scenario, denoted as **Connected Server**, we run Hadoop in `Nemesis` while it exchanges data on a gigabit LAN. In order to measure the effect of the network on the energy consumption, we evaluate 4 different cases for each application. These cases result from combining 2 different behaviors, depending on whether `Nemesis` acts as a sender or as a receiver of data, with 2 different packet sizes, 64 and 1470 bytes. To do so, we run Iperf, as a server or as a client according to the case, in parallel with Hadoop.

Finally, in the third scenario, denoted as **Cloud**, we set up a two-server Hadoop cluster with `Nemesis` and `Survivor`. In this scenario `Nemesis` is configured as the master node of the cluster and `Survivor` as a slave node. The execution of the applications is shared by both nodes so Hadoop itself exchanges traffic between both servers, and we do not insert additional network traffic in this case. Finally, in order to have a better control of the experiment, we force the reduce tasks to be mandatorily run in `Nemesis`, which also conditions the way the data is exchanged between `Nemesis` and `Survivor`.

Observe that all 3 scenarios are based on Hadoop. This implies that, apart from the map and reduce tasks due to the applications being run, there are some extra processes executed in the servers we are using. The most important processes that we can find in `Nemesis` are *NameNode* (the process that keeps the directory tree of all files in the file system, and tracks where across the cluster the file data is kept), *Secondary NameNode* (that performs periodic checkpoints of the *NameNode*), *DataNode* (the process that is in charge of storing data in the Hadoop File System (HDFS)), *JobTracker* (that receives the jobs and submits MapReduce tasks to the cluster nodes) and *TaskTracker* (a per node process that can accept a determined number of MapReduce tasks). On its side, `Survivor` runs, in the cloud scenario, DataNode and TaskTracker.

3.4.3 Experiments and Observed Results

For the sake of consistency in the results, we ran both applications 10 times per frequency for each one of the considered scenarios and averaged the results.

We start by describing the *Isolated Server* scenario. For each run i we record the total number of active cycles executed C^i_{app}, the time spent T^i_{app} and the volume of data read (written), $V^{r,i}_D$ ($V^{w,i}_D$). Since we cannot measure the instantaneous CPU load, we assume that the CPU load is the same during the run for a given frequency. Hence, the CPU load can be estimated as

$$\rho^i_{app} = C^i_{app}/T^i_{app}.$$

Then, from ρ^i_{app} we obtain the estimate of the instantaneous power P^i_{BC} using the fitting curves as described in Section 3.3. Finally, using Eq. 3.9 we compute the estimate $E^i_B + E^i_C$. In order to estimate the energy consumed by the disk operations, we use the fact that Hadoop uses a block size of 64 MB. This allows us to estimate the reading (writing) efficiencies, $\eta^{r,i}_D$ ($\eta^{w,i}_D$) that we compute, in Joules per byte. Combining these values with the measured volume of data read and written ($V^{r,i}_D$ and $V^{w,i}_D$), as described in Eq. 3.10, we obtain E^i_D.

The total estimated energy of the application in run i, E^i_{app}, is obtained by summing up the energy of the different components used in run i, as stated in Eq. 3.8 (remember that, in the *Isolated Server* the network is not used). Then, by summing the values of the ten runs of an experiment, we obtain the total estimated energy as

$$E_{app} = \sum_{i=1}^{10} E^i_{app}.$$

The (approximated) total *real* energy $\hat{E}^i_{app}$ consumed in run i is computed by the average value of the power samples which we registered with the power analyzer during the run, and we multiply it with the run time T_{app}. Finally, the total energy consumed by the experiment is obtained as

$$\hat{E}_{app} = \sum_{i=1}^{10} \hat{E}^{i}_{app}.$$

The estimation error for each experiment is then computed as $\hat{E}_{app} - E_{app}$.

We show the results obtained for the *Isolated Server* scenario with the minimum, the maximum, and the most efficient[10] frequencies (the results for the remaining frequencies are similar) in Figure 3.12. The figure shows the results for both PageRank and WordCount. As can be seen, the error is relatively small, except for the case when we run WordCount at the maximum frequency. Errors are of 4%, 4%, 7%, 5%, 7% and 10% respectively, following the same order as in Figure 3.12.

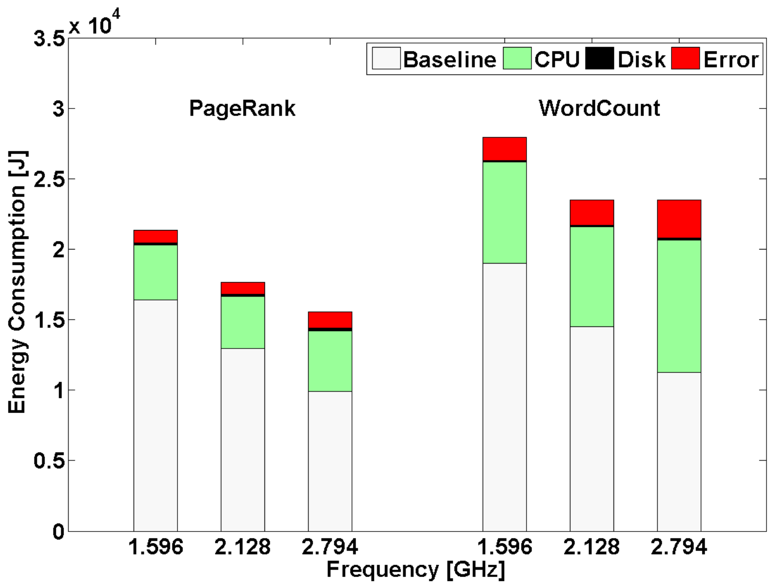

Figure 3.12: Energy consumption of `Nemesis` in the Isolated Server scenario.

We move now to the *Connected Server* scenario. As we described in the previous section, this scenario is studied in 4 different cases depending on whether `Nemesis` acts as sender or receiver and whether the size of the packets is of 64 or 1470 bytes. Of course, another relevant parameter is the rate at which these packets are sent. The rates used are 150 and 400 Mbps when using packets of 64 or 1470 bytes, respectively.

The total energy consumed in these cases is computed in the same way as we did for the *Isolated Server* scenario but adding the contribution of the network. In order to estimate the network consumption in one run with `Nemesis` sending traffic (resp., receiving traffic), the sending efficiency η^s_N, (resp., receiving efficiency η^r_N) is obtained from the transfer rate R, the frequency and packet size used (see Figures 3.9). The amount of data sent (resp., received) can be obtained

[10] Respectively 1.596, 2.128 and 2.794 GHz, according to the results shown in Section 3.3.

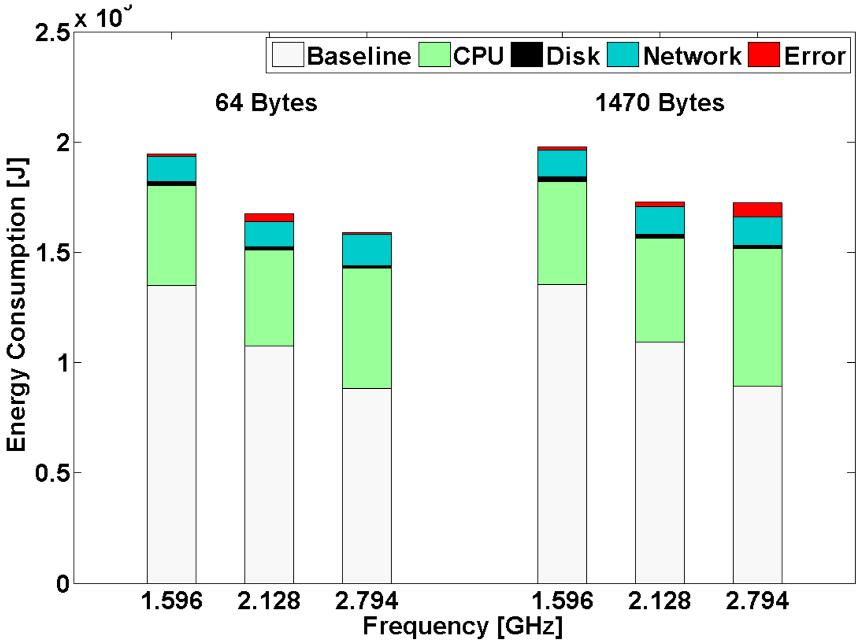

(a) PageRank, sender side.

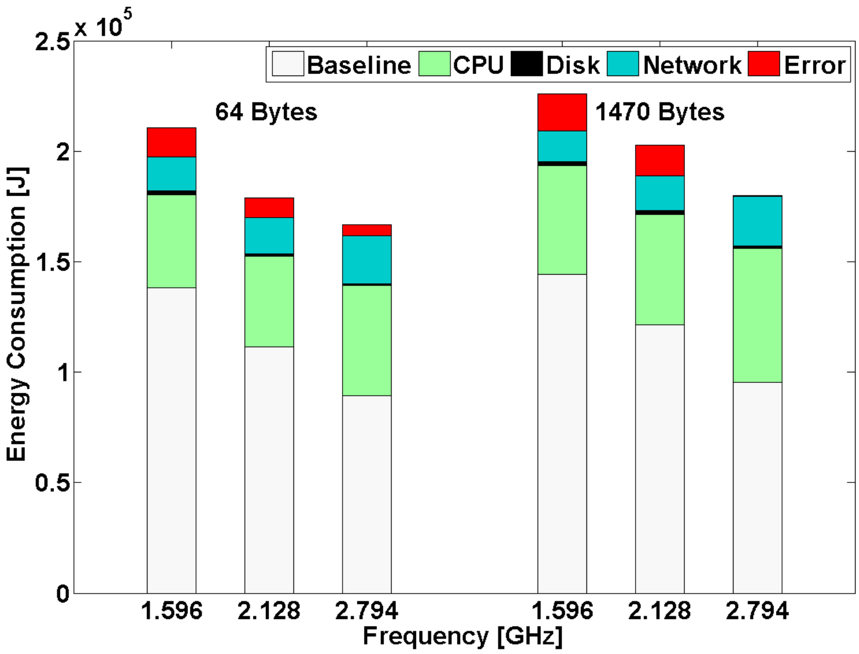

(b) PageRank, receiver side.

Figure 3.13: Energy consumption of `Nemesis` running PageRank in the Connected Server scenario, with either small or big packets.

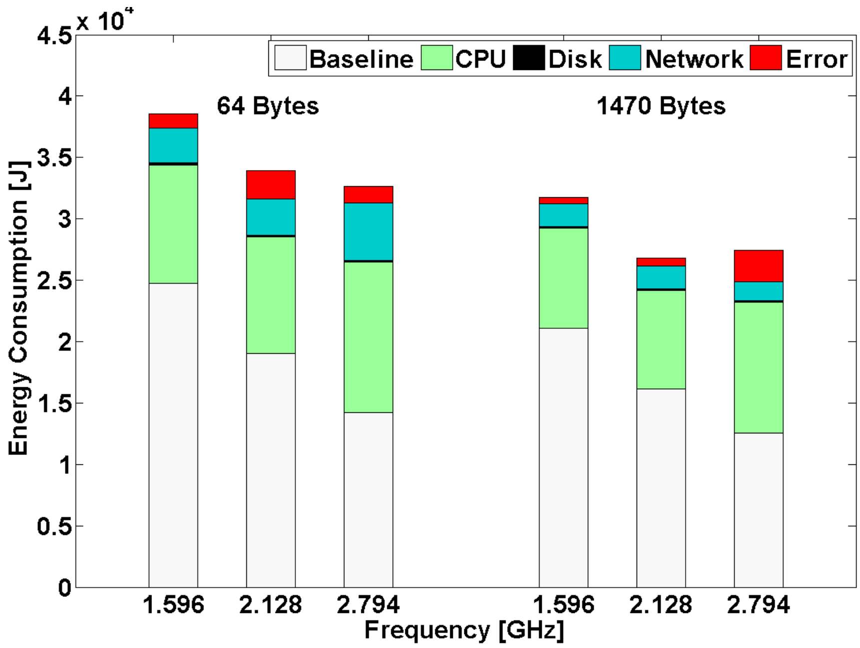

(a) WordCount, sender side.

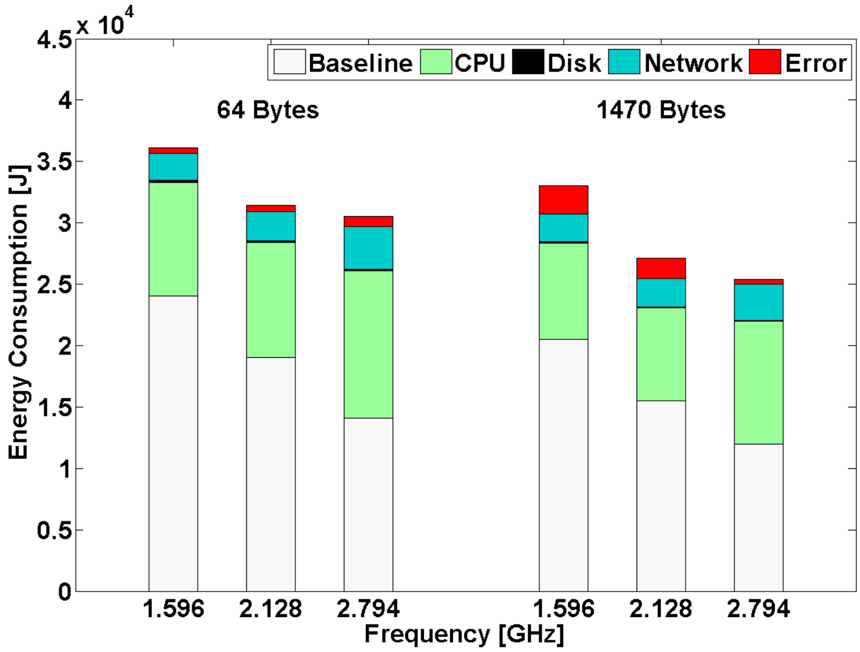

(b) WordCount, receiver side.

Figure 3.14: Energy consumption of `Nemesis` running WordCount in the Connected Server scenario, with either small or big packets.

Table 3.2: Error measured in the different cases of the *Connected Server* scenario.

Packet Size	Freq	Cases			
		PR - Send	PR - Rec	WC - Send	WC - Rec
64-B	1.596	0.5%	6.0%	2.9%	2.7%
	2.128	2.0%	4.7%	6.4%	1.5%
	2.794	0.5%	2.9%	4.0%	2.9%
1470-B	1.596	0.7%	6.9%	1.6%	6.5%
	2.128	1.1%	6.5%	5.8%	5.8%
	2.794	3.8%	0.3%	0.9%	1.5%

from the server itself by consulting the OS registers[11]. Therefore, the energy of the network for an run i, E_N^i, is obtained using Eq. 3.11. Then, including E_N^i for each run in the computation of E_{app}^i we can obtain the total energy consumed by the application. Following the same steps as in the previous scenario, we get the results shown in Figure 3.13 and 3.14. The error measured is again relatively smaller for PageRank than for WordCount. The error measured for each of the cases can be found in Table 3.2.

We finally analyze the *Cloud* scenario. In this scenario we set up a cluster with two servers, `Nemesis` and `Survivor`, and run the 2 aforementioned Hadoop applications in it. This scenario may seem relatively similar to the *Connected Server* scenario, but it has is a major difference. While in the previous scenario we were the ones controlling the network traffic, here the traffic is controlled by Hadoop. Specifically, we know that, in this scenario, there are two main sources of traffic: requesting input data when it is not present in a server, and sending the mapper tasks outputs to the reducer tasks. The only condition we impose in the server to have some control over the traffic is related to this later aspect, we force the reducers to be always in `Nemesis`.

Although we are able to retrieve the total amount of data received or sent by each server, we know neither the size of the packets used nor the rate. Therefore, we can compute neither the sending efficiency η_N^s nor the receiving efficiency η_N^r. In order to be able to compute both the sending and receiving efficiencies we analyze the traffic exchanged by both servers for each one of the applications. Figure 3.15 shows the amount of packets of each size that were exchanged by both servers (and the direction of the exchange) for both applications. The results show the vast majority of packets are either small (64 bytes) or big (1470 bytes). Moreover, it shows that most of the packets sent from `Nemesis` to `Survivor` are small packets for both applications, while big packets are sent in the opposite direction.

Given these results, we approximate the energy consumed by the network assuming that all the packets exchanged are of the same size and that the rate is the maximum achievable rate for each packet size according to the results from Section 3.3. For instance, we consider roughly 30 Mbps when `Survivor` receives 64-Byte packets and roughly 970 Mbps if it sends 1470-Byte packets. These assumptions allow us to compute now η_N^s and η_N^r. The remaining parameters are computed

[11] We can read the registers rx_bytes, rx_packets, tx_bytes or tx_packets from `/sys/class/net/eth0/statistics`.

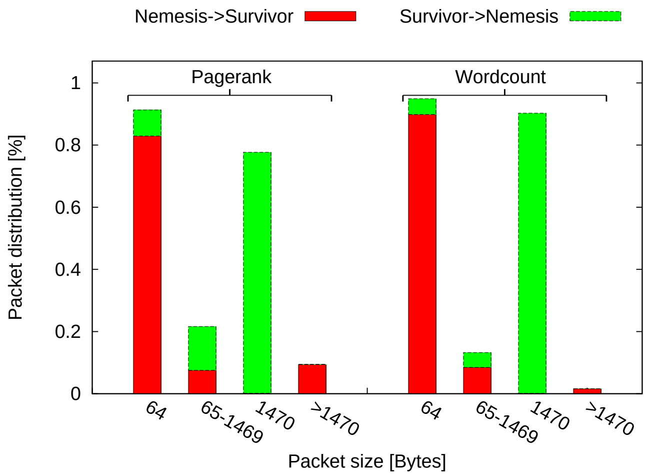

Figure 3.15: Distribution of the sizes of the packets exchanged between `Nemesis` and `Survivor` for both PageRank and WordCount in the Cloud scenario.

as for the other scenarios, so to determine $\hat{E}_{app}$ and E_{app}. The results are shown in Figure 3.16. As in the previous scenarios, errors are relatively low. In particular, the error in `Nemesis` when running PageRank is 3.1% and 1.4% for 2.128 GHz and 2.794 GHz, respectively, and of a 9.7% and a 6.5% for 2.128 GHz and 2.794 GHz when running WordCount. On the other hand, the measured errors for `Survivor` are 3.3% and 3.6% for 1.867 GHz and 2.133 GHz when running PageRank and 5.1% and 5.2%, respectively, when running WordCount.

3.5 Discussion

We discuss now some of the implications of our results. We start with consolidation as a technique for energy saving. It has been often assumed that the best way of saving energy is by using the highest frequency available and applying consolidation (which is to fill servers as much as possible). This reduces the total number of servers being used, allowing to switch off the rest. This assumption has led to proposing bin-packing based solutions [24, 76, 82, 95]. However, the results presented in Figures 3.3(b), 3.4(b) and 3.5(b) show that the highest frequency is not always the most efficient one, and this has been found to be true for two different architectures (Intel and AMD). This implies that, by running servers at the optimal amount of load, and the right frequency, a considerable amount of energy could be saved.

A second relevant aspect is the baseline consumption of servers. The results presented for all 3 servers show that their baselines are within a 30-50$\%$ of the maximum consumption. Then, it is obvious that more effort has to be done for reducing baseline consumption. For instance, a solution could consist in switching off cores in real time, not just disabling them, or in introducing very fast transitions between active and lower energy states, i.e., to achieve real *suspension* in idle state.

There is another relevant issue related to the CPU load associated to disk and network activity.

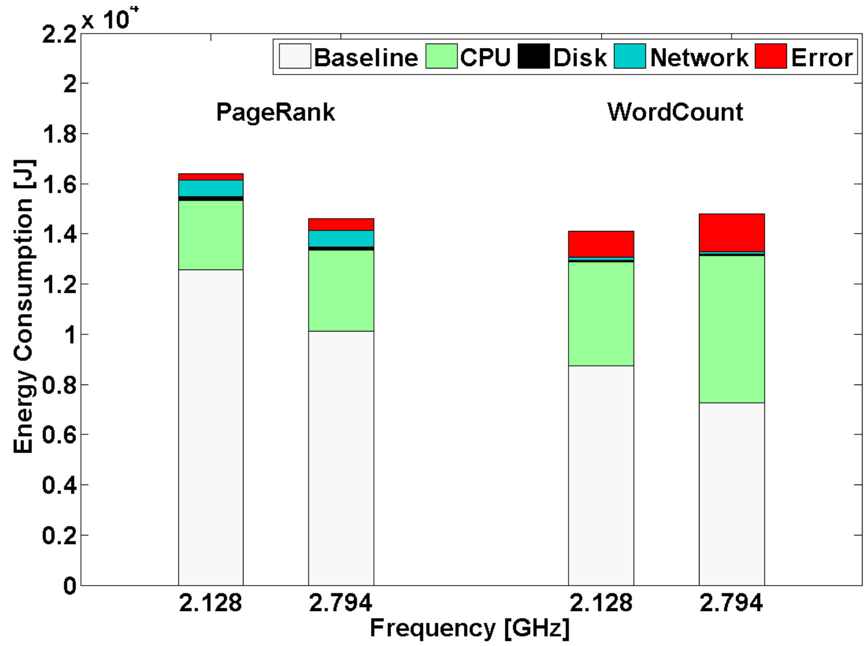

(a) Nemesis

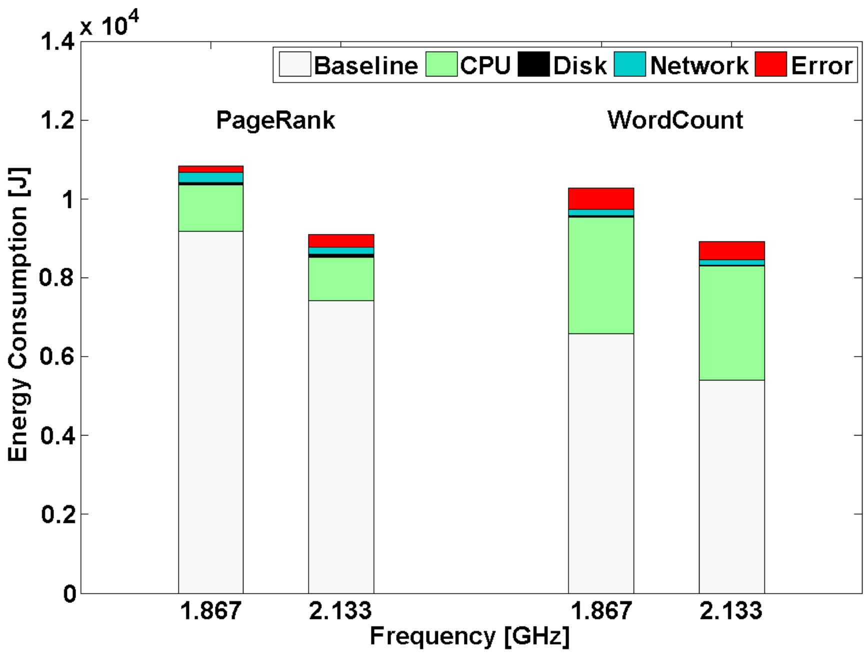

(b) Survivor

Figure 3.16: Energy consumption of `Nemesis` and `Survivor` in the Cloud scenario.

It can be observed in Figure 3.6 that disks do not incur much CPU overhead. In fact, the power used by the CPU plus baseline does not change much across the experiments. Instead, the energy consumed by the CPU due to network operations is even larger than the energy consumed by the NIC (see Figure 3.11). Some works, like [47], have already pointed out that the way packets are handled by the protocol stack is not energy efficient. Our results reinforce this feeling and point out that building a more efficient protocol stack would certainly reduce the amount of energy consumed due to the network.

Finally, it is worth to mention that in this work we have assumed that the power utilization of the RAM memory is included in the *baseline*. The characterization experiments have been run in such a way that there were few memory accesses, so its power utilization did not affect our measurements. However, RAM memory became an uncontrolled source of power utilization in Section 3.4.3 when we validated our proposed model. In fact, all the Hadoop processes that run in the servers consume significant RAM memory. This impacts more significantly the memory used by the cluster's master node, since it runs internal Hadoop processes (such as the NameNode or the JobTracker) whose memory requirement increases with the number of mappers and reducers. This cost is, therefore, paid only in `Nemesis`, the master node of our cluster, and not in `Survivor`, which explains the different accuracy of the model for the two servers. This error is particularly evident when when WordCount is run, due to the fact that the required number of mappers for WordCount is larger than for PageRank and, therefore, the RAM required in `Nemesis` increases and so does the uncontrolled energy consumption.

3.6 Conclusions

In this chapter we have reported our measurement-based characterization of energy and power consumption in a server. We have exhaustively measured the power consumed by CPU, disk, and NIC under different configurations, identifying the optimal operational levels, which usually do not correspond to the static system configurations commonly adopted. We found that, besides the *baseline component*, which does not changes significantly with the operational parameters, the CPU has the largest impact on energy consumption among all the three components. We observe that CPU consumption is neither linear nor concave with the load, i.e., the systems are not *energy proportional*. Disk I/O is the second larger contributor to power consumption, although performance changes sensibly with the I/O block size used by the applications. Finally, the NIC activity is responsible for a small but not negligible fraction of power consumption, which scales almost linearly with the network transmission rate. In general, most of the energy/power performance figures do not scale linearly with the utilization, in contrast to what is commonly assumed in the literature. We have then shown how to predict and optimize the energy consumed by an application via a concrete example using 2 different Hadoop applications, PageRank and WordCount, in three different scenarios. First We ran both applications without network activity, next with bulky network activity, and finally in a two-server cluster. Our model achieves very accurate energy estimates, within 4.1% from the measured total power consumption on average and never worse than a 10%.

Chapter 4

Efficient Assignment of Virtual Machines to Physical Machines

4.1 Overview

Having studied how nowadays servers use power, we now show a way in which this knowledge can be useful. In this chapter we will apply the concept of the optimal operational point of a server to the Virtual Machine Assignment problem (VMA), which basically consists in deciding in which server we want to place a new task arriving to a system. Having an optimal operational point conditions drastically the way we must load a server in order to optimize the energy which is being consumed.

Having this in mind, we study, in particular, the hardness and online competitiveness of the four versions of the VMA problem that we described in Chapter 2. This 4 versions of VMA differed in whether we considered PMs with bounded or unbounded capacity C or a limited or unlimited number of PMs m in the system. We denoted these versions as $(\cdot,\cdot)$-VMA, $(C,\cdot)$-VMA, $(\cdot,m)$-VMA and (C,m)-VMA[1].

We start by showing various lower and upper bounds on the offline approximation of VMA. The first fact we observe is that there is a hard decision version of (C,m)-VMA: Is there a feasible partition π of the set D of VMs? By reduction from the 3-Partition problem, it can be shown that this decision problem is strongly NP-complete. We then show that the $(\cdot,\cdot)$-VMA, $(C,\cdot)$-VMA, and $(\cdot,m)$-VMA problems are NP-hard in the strong sense, even if α is constant. This result implies that VMA problems do not have a fully polynomial time approximation scheme (FPTAS), even if α is constant. Nevertheless, using previous results derived for more general objective functions, we notice that $(\cdot,\cdot)$-VMA and $(\cdot,m)$-VMA have a polynomial time approximation scheme (PTAS), while the $(C,\cdot)$-VMA problem can not be approximated beyond a ratio of $\frac{3}{2}\cdot\frac{\alpha-1+(\frac{2}{3})^{\alpha}}{\alpha}$ (unless P $=$ NP). On the positive side, we show how to use an existing Asymptotic PTAS [45] to obtain algorithms that approximate the optimal solution of $(C,\cdot)$-VMA. All our

[1]The central dots $(\cdot)$ imply unboundedness.

offline results as well as the online ones can be seen in Table 4.1.

Then we move on to online VMA algorithms. We show various upper and lower bounds on the competitive ratio of the four versions of the problem. Observe that the results are often different depending on whether x^* is smaller than C or not. In fact, when $x^* < C$, there is a lower bound of $\frac{(3/2)2^\alpha - 1}{2^\alpha - 1}$ that applies to all versions of the problem. Rather than attempting to obtain tight bounds for particular instances of the parameters of the problem (C, m, α, b) we focus on obtaining *general bounds*, whose parameters can be instantiated for the specific application. The bounds obtained show interesting trade-offs between the PM capacity and the fixed cost of adding a new PM to the system. For clarity, we will consider $\mu = 1$ throughout the whole chapter. All the results presented apply for other values of μ.

The resulting bounds are shown in Table 4.1. As can be observed, the resulting upper and lower bounds are not very far in general. To give some intuition on the tightness of these bounds, we instantiate them for a realistic value of $\alpha = 3$, and normalized values of $b = 2$ and $C \in \{1, 2\}$. These values for α are obtained from the servers we studied in Chapter 3, in particular from the ones denoted as `Erdos` and `Nemesis`. In them the values for α are close to 1.5 and 3 and x^* values of $0.76C$ and $0.9C$ respectively (x^* denotes the load that minimizes the ratio power consumption against load.). The bounds based on these realistic values are shown in Table 4.2.

RoadMap The rest of the chapter is organized as follows. Section 4.2 includes some preliminary results that will be used throughout the chapter. The offline and online analyses are included in Section 4.3 and 4.4 respectively. In Section 4.5 we compare different state of the art allocation policies and compare them with the algorithms proposed in Section 4.4. Section 4.6 discusses some practical issues and provides some useful insights regarding real implementation. Section 4.7 concludes the Chapter.

4.2 Preliminaries

The following claims will be used in the analysis. We call ***power rate*** the power consumed per unit of load in a PM. Let x be the load of a PM. Then, its power rate is computed as $f(x)/x$. The load at which the power rate is minimized, denoted x^*, is the ***optimal load***, and the corresponding rate is the ***optimal power rate*** $\varphi^* = f(x^*)/x^*$. Using calculus we get the following observation.

Observation 1. *The optimal load is* $x^* = (b/(\alpha - 1))^{1/\alpha}$. *Additionaly, for any* $x \neq x^*$, $f(x)/x > \varphi^*$.

The following lemmas will be used in the analysis.

Lemma 1. *Consider two solutions* $\pi = \{A_1, \ldots, A_m\}$ *and* $\pi' = \{A'_1, \ldots, A'_m\}$ *of an instance of the VMA problem, such that for some* $x, y \in [1, m]$ *it holds that*

- $A_x \neq \emptyset$ *and* $A_y \neq \emptyset$*;*

VMA subprob.	$x^* < C$	$x^* \geq C$
$(C, \cdot)$ offline	$\rho \geq \frac{3}{2}\frac{\alpha-1+(2/3)^\alpha}{\alpha}$	$\rho \geq \frac{3}{2}\frac{\alpha-1+(2/3)^\alpha}{\alpha}$
	$\rho < \frac{\overline{m}}{m^*}\left(1+\epsilon+\frac{1}{\alpha-1}+\frac{1}{\overline{m}}\right)$	$\rho < 1+\epsilon+\frac{C^\alpha}{b}+\frac{1}{\overline{m}}$
$(C, \cdot)$ online	$\rho \geq \frac{(3/2)2^\alpha-1}{2^\alpha-1}$	$\rho \geq \frac{C^\alpha+2b}{b+\max\{C^\alpha, 2(C/2)^\alpha+b\}}$
	$\rho = 1$ if $D_s = \emptyset$, else $\rho \leq \left(1-\frac{1}{\alpha}\left(1-\frac{1}{2^\alpha}\right)\right)\left(2+\frac{x^*}{\ell(D_s)}\right)$	$\rho \leq \frac{2b}{C}\left(1+\frac{1}{(\alpha-1)2^\alpha}\right)\left(2+\frac{C}{\ell(D)}\right)$
(C, m) online	$\rho \geq \frac{(3/2)2^\alpha-1}{2^\alpha-1}$	$\rho \geq \frac{C^\alpha+2b}{b+\max\{C^\alpha, 2(C/2)^\alpha+b\}}$
$(\cdot, \cdot)$ online	$\rho \geq \frac{(3/2)2^\alpha-1}{2^\alpha-1}$	not applicable
	$\rho = 1$ if $D_s = \emptyset$, else $\rho \leq \left(1-\frac{1}{\alpha}\left(1-\frac{1}{2^\alpha}\right)\right)\left(2+\frac{x^*}{\ell(D_s)}\right)$	
$(\cdot, m)$ online	$\rho \geq \max\{\frac{(3/2)2^\alpha-1}{2^\alpha-1}, \frac{3^\alpha}{2^{\alpha+2}+\epsilon}\}$	not applicable
	$\rho \leq O(\alpha)^\alpha$ In [53]	
$(\cdot, 2)$ online	$\rho \geq \max\{\frac{3^\alpha}{2^{\alpha+1}}, \frac{(3/2)2^\alpha-1}{2^\alpha-1}, \frac{3^\alpha}{2^{\alpha+2}+\epsilon}\}$	not applicable
	$\rho = 1$ if $\ell(D) \leq \sqrt[\alpha]{b/(2^\alpha-2)}$, else $\rho \leq \max\{2, \left(\frac{3}{2}\right)^{\alpha-1}\}$	

Table 4.1: Summary of bounds on the approximation/competitive ratio ρ. All lower bounds are existential. The number of PMs in an optimal $(C, \cdot)$-VMA solution is denoted as m^*. The number of PMs in an optimal Bin Packing solution is denoted as $\overline{m}$. The load that minimizes the ratio power consumption against load is denoted as x^*. The subset of VMs with load smaller than x^* is denoted as D_s.

- *$A'_x = A_x \cup A_y$, $A'_y = \emptyset$, and $A_i = A'_i$, for all $i \neq x$ and $i \neq y$; and*
- *$\ell(A_x) + \ell(A_y) \leq \min\{x^*, C\}$.*

Then, $P(\pi') < P(\pi)$.

Proof: Let $\ell(A_i) = x$ and $\ell(A_j) = y$. First we notice that π' is feasible because $x + y \leq C$. Now, using that $x + y \leq x^*$, we have

$$b = (x^*)^\alpha(\alpha-1) \geq (x+y)^\alpha(\alpha-1) > (x+y)^\alpha \geq (x+y)^\alpha - (x^\alpha + y^\alpha)$$

where the second inequality comes from the fact that $\alpha > 1$. The above inequality is equivalent to

$$2b + x^\alpha + y^\alpha > b + (x+y)^\alpha,$$

VMA subprob.	$x^* < C$	$x^* \geq C$
$(C, \cdot)$ offline	$\rho \geq \frac{11}{9}$	$\rho \geq \frac{11}{9}$
	$\rho < \frac{\overline{m}}{m^*}\left(\frac{3}{2} + \epsilon + \frac{1}{\overline{m}}\right)$	$\rho < \frac{3}{2} + \epsilon + \frac{1}{\overline{m}}$
$(C, \cdot)$ online	$\rho \geq \frac{11}{7}$	$\rho \geq \frac{20}{17}$
	$\rho \leq \frac{17}{12}\left(1 + \frac{1}{2\ell(D_s)}\right)$	$\rho \leq \frac{17}{2}\left(1 + \frac{1}{2\ell(D)}\right)$
(C, m) online	$\rho \geq \frac{11}{7}$	$\rho \geq \frac{20}{17}$
$(\cdot, \cdot)$ online	$\rho \geq \frac{11}{7}$	not applicable
	$\rho \leq \frac{17}{12}\left(1 + \frac{1}{2\ell(D_s)}\right)$	
$(\cdot, m)$ online	$\rho \geq \frac{11}{7}$	not applicable
$(\cdot, 2)$ online	$\rho \geq \frac{11}{7}$	not applicable
	$\rho \leq \frac{9}{4}$	

Table 4.2: Summary of bounds on the approximation/competitive ratio ρ for $\alpha = 3$, $b = 2$, and $C = 2$ on the left and $C = 1$ on the right.. All lower bounds are existential. The number of PMs in an optimal $(C, \cdot)$-VMA solution is denoted as m^*. The number of PMs in an optimal Bin Packing solution is denoted as $\overline{m}$. The load that minimizes the ratio power consumption against load is denoted as x^*. The subset of VMs with load smaller than x^* is denoted as D_s.

which implies the lemma. ■

From this lemma, it follows that the global power consumption can be reduced by having 2 VMs together in the same PM, when its aggregated load is smaller than $\min\{x^*, C\}$, instead of moving one VM to an unused PM. When we keep VMs together in a given partition we say that we are *using* Lemma 1.

Lemma 2. *Consider two solutions $\pi = \{A_1, \ldots, A_m\}$ and $\pi' = \{A'_1, \ldots, A'_m\}$ of an instance of the VMA problem, such that for some $x, y \in [1, m]$ it holds that*

- *$A_x \cup A_y = A'_x \cup A'_y$, while $A_i = A'_i$, for all $x \neq i \neq y$;*
- *none of A_x, A_y, A'_x, and A'_y is empty; and*
- *$|\ell(A_x) - \ell(A_y)| < |\ell(A'_x) - \ell(A'_y)|$.*

Then, $P(\pi) < P(\pi')$.

Proof: From the definition of $P(\cdot)$, to prove the claim is it enough to prove that

$$\ell(A_x)^\alpha + \ell(A_y)^\alpha < \ell(A'_x)^\alpha + \ell(A'_y)^\alpha.$$

Let us assume wlog that $\ell(A_x) \leq \ell(A_y)$ and $\ell(A'_x) \leq \ell(A'_y)$. Let us denote

$$L = \ell(A_x) + \ell(A_y) = \ell(A'_x) + \ell(A'_y),$$

and assume that $\ell(A_x) = \delta_1 L$ and $\ell(A'_x) = \delta_2 L$. Note that $\delta_2 < \delta_1 \leq 1/2$. Then, the claim to be proven becomes

$$(\delta_1 L)^\alpha + ((1-\delta_1)L)^\alpha < (\delta_2 L)^\alpha + ((1-\delta_2)L)^\alpha$$
$$\delta_1^\alpha + (1-\delta_1)^\alpha < \delta_2^\alpha + (1-\delta_2)^\alpha$$

Which holds because the function $f(x) = x^\alpha + (1-x)^\alpha$ is decreasing in the interval $(0, 1/2)$. ■

This lemma carries the intuition that balancing the load among the used PMs as much as possible reduces the power consumption.

Corollary 1. *Consider a solution $\pi = \{A_1, \ldots, A_m\}$ of an instance of the VMA problem with total load $\ell(D)$, such that exactly k of the A_x sets, $x \in [1, m]$, are non-empty (hence it uses k PMs). Then, the power consumption is lower bounded by the power of the (maybe unfeasible) solution that balances the load evenly, i.e.,*

$$P(\pi) \geq kb + k(\ell(D)/k)^\alpha.$$

4.3 Offline Analysis

4.3.1 NP-hardness

As was mentioned, it can be shown that deciding whether there is a feasible solution for an instance of the (C, m)-VMA problem is NP-complete or not, by a direct reduction from the 3-Partition problem. However, this result does not apply directly to the $(C, \cdot)$-VMA, $(\cdot, m)$-VMA, and $(\cdot, \cdot)$-VMA problems. We show now that these problems are NP-hard. We first prove the following lemma.

Lemma 3. *Given an instance of the VMA problem, any solution $\pi = \{A_1, \ldots, A_m\}$ where $\ell(A_i) \neq x^*$ for some $i \in [1, m] : A_i \neq \emptyset$, has power consumption $P(\pi) > \rho^* \ell(D) = \rho^* \sum_{d \in D} \ell(d)$.*

Proof: The total cost of π is $P(\pi) = \sum_{i \in [1,m]} f(\ell(A_i))$ which, from Observation 1, satisfies

$$\begin{aligned} P(\pi) &> \sum_{i \in [1,m]: A_i \neq \emptyset} \ell(A_i) \rho^* \\ &= \rho^* \sum_{i \in [1,m]: A_i \neq \emptyset} \sum_{d \in A_i} \ell(d) = \rho^* \sum_{d \in D} \ell(d). \end{aligned}$$

■

We show now in the following theorem that the different versions of the (C, m)-VMA problem with unbounded C or m are NP-hard.

Theorem 1. *The $(C, \cdot)$-VMA, $(\cdot, m)$-VMA and $(\cdot, \cdot)$-VMA problems are strongly NP-hard, even if α is constant.*

Proof: We show a reduction from 3-Partition defined as follows [48], which is strongly NP-complete.

INSTANCE: Set A of $3k$ elements, a bound $B \in \mathbb{Z}^+$ and, for each $a \in A$, a size $s(a) \in \mathbb{Z}^+$ such that $B/4 < s(a) < B/2$ and $\sum_{a \in A} s(a) = kB$.

QUESTION: can A be partitioned into k disjoint sets $\{A_1, A_2, \ldots, A_k\}$ such that $\sum_{a \in A_i} s(a) = B$ for each $1 \leq i \leq k$?

The reduction is as follows. Given an instance of 3-Partition on a set $A = \{a_1, \ldots, a_{3k}\}$ with bound B, and given a fixed value $\alpha > 1$, we define an instance $\mathcal{I}$ of $(\cdot, \cdot)$-VMA as follows: $D = \{a_1, \ldots, a_{3k}\}$, $\ell(\cdot) = s(\cdot)$, and $b = B^\alpha(\alpha - 1)$ (i.e., $x^* = B$). (For the proof of the $(C, \cdot)$-VMA and $(\cdot, m)$-VMA problems it is enough to set $C = B$ and $m = k$ when required.) We show now that the answer to the 3-Partition problem is YES if and only if the output $\pi = \{A_1, A_2, \ldots, A_m\}$ of the $(\cdot, \cdot)$-VMA problem on input $\mathcal{I}$ is such that $\sum_{i=1}^{m} f(\ell(A_i)) = kf(B)$.

For the direct implication, assume that there exists a partition $\{A_1, A_2, \ldots, A_k\}$ of A such that for each $i \in [1, k]$, $\sum_{a \in A_i} s(a) = B$. Then, in the context of the $(\cdot, \cdot)$-VMA problem, such partition has cost $\sum_{i=1}^{m} f(\ell(A_i)) = kf(B)$. We claim that any partition has at least cost $kf(B)$. In order to prove it, assume for the sake of contradiction that there is a partition $\pi' = \{A'_1, A'_2, \ldots, A'_m\}$ of $(\cdot, \cdot)$-VMA on input $\mathcal{I}$ with cost less than $kf(B)$. Then, there is some $i \in [1, m]$ such that $A'_i \neq \emptyset$ and $\ell(A'_i) \neq B$. From Lemma 3, $P(\pi') > \rho^* \ell(D) = (f(x^*)/x^*)kB$. Since $B = x^*$, we have that $P(\pi') > kf(B)$, which is a contradiction.

To prove the reverse implication, assume an output $\pi = \{A_1, A_2, \ldots, A_m\}$ of the $(\cdot, \cdot)$-VMA problem on input $\mathcal{I}$ such that $P(\pi) = \sum_{i=1}^{m} f(\ell(A_i)) = kf(B)$. Then, it must be $\forall i \in [1, m]$: $A_i \neq \emptyset, \ell(A_i) = B$. Otherwise, from Lemma 3, $P(\pi) > kf(B)$, a contradiction. ■

It is known that strongly NP-hard problems cannot have a fully polynomial-time approximation scheme (FPTAS) [92]. Hence, the following corollary.

Corollary 2. *The $(C, \cdot)$-VMA, $(\cdot, m)$-VMA and $(\cdot, \cdot)$-VMA problems do not have fully polynomial-time approximation schemes (FPTAS), even if α is constant.*

In the following sections we show that, while the $(\cdot, m)$-VMA and $(\cdot, \cdot)$-VMA problems have polynomial-time approximation schemes (PTAS), the $(C, \cdot)$-VMA problem cannot be approximated below $\frac{3}{2} \cdot \frac{\alpha - 1 + (2/3)^\alpha}{\alpha}$.

4.3.2 The $(\cdot, m)$-VMA and $(\cdot, \cdot)$-VMA Problems Have PTAS

We have proved that the $(\cdot, m)$-VMA and $(\cdot, \cdot)$-VMA problems are NP-hard in the strong sense and that, hence, there exists no FPTAS for them. However, Alon *et al.* [8], proved that if a function $f(\cdot)$ satisfies a condition denoted $F*$, then the problem of scheduling jobs in m identical machines so that $\sum_i f(M_i)$ is minimized has a PTAS, where M_i is the load of the jobs allocated to machine i. This result implies that if our function $f(\cdot)$ satisfies condition $F*$, the same PTAS can be used for the $(\cdot, m)$-VMA and $(\cdot, \cdot)$-VMA problems. From Observation 6.1 in [42], it can be derived that, in fact, our power consumption function $f(\cdot)$ satisfies condition $F*$. Hence, the following theorem.

Theorem 2. *There are polynomial-time approximation schemes (PTAS) for the* $(\cdot, m)$*-VMA and* $(\cdot, \cdot)$*-VMA problems.*

4.3.3 Bounds on the Approximability of the $(C, \cdot)$-VMA Problem

We study now the $(C, \cdot)$-VMA problem, where we consider an unbounded number of machines with bounded capacity C. We will provide a lower bound on its approximation ratio, independently on the relation between x^* and C; and upper bounds for the cases when $x^* \geq C$ and $x^* < C$.

4.3.3.1 Lower bound on the approximation ratio

The following theorem shows a lower bound on the approximation ratio of any offline algorithm for $(C, \cdot)$-VMA.

Theorem 3. *No algorithm achieves an approximation ratio smaller than* $\frac{3}{2} \cdot \frac{\alpha - 1 + (\frac{2}{3})^\alpha}{\alpha}$ *for the* $(C, \cdot)$*-VMA problem unless* $\mathrm{P} = \mathrm{NP}$.

Proof: The claim is proved showing a reduction from the partition problem [48]. In the partition problem there is a set $A = \{a_1, a_2, \ldots, a_n\}$ of n elements, there is a size $s(a)$ for each element $a \in A$, and the sum $M = \sum_{a \in A} s(a)$ of the sizes of the elements in A. The problem decides whether there is a subset $A' \subset A$ such that $\sum_{a \in A'} s(a) = M/2$.

From an instance of the partition problem, we construct an instance of the $(C, \cdot)$-VMA problem as follows. The set of VMs in the system is $D = \{a_1, a_2, \ldots, a_n\}$, the load function is $\ell(\cdot) = s(\cdot)$, the capacity of each PM is set to $C = M/2$, and b is set to $b = C^\alpha(\alpha - 1)$ (i.e., $x^* = C$). Let us study the optimal partition π^* such that the total power consumption $P(\pi^*)$ is minimized. If there is a partition of D such that each subset in this partition has load $M/2$ then, from Observation 1, π^* has all the VMs assigned to two PMs. Otherwise, π^* needs at least 3 PMs to allocate all the VMs. From Corollary 1, the power consumption of this solution is lower bounded by the power of a (maybe unfeasible) partition that balances the load among the 3 PMs as evenly as possible. Formally,

$$\exists A' : \sum_{a \in A'} s(a) = M/2 \quad \Rightarrow \quad P(\pi^*) = 2b + 2\left(\frac{M}{2}\right)^\alpha = 2b + 2C^\alpha$$
$$\nexists A' : \sum_{a \in A'} s(a) = M/2 \quad \Rightarrow \quad P(\pi^*) \geq 3b + 3\left(\frac{M}{3}\right)^\alpha = 3b + 3\left(\frac{2C}{3}\right)^\alpha.$$

Comparing both values we obtain the following ratio.

$$\begin{aligned}\rho &= \frac{3b + 3\left(\frac{2C}{3}\right)^{\alpha}}{2b + 2C^{\alpha}} \\ &= \frac{3C^{\alpha}(\alpha - 1) + 3\left(\frac{2C}{3}\right)^{\alpha}}{2C^{\alpha}(\alpha - 1) + 2C^{\alpha}} \\ &= \frac{3}{2} \cdot \frac{\alpha - 1 + (\frac{2}{3})^{\alpha}}{\alpha}.\end{aligned}$$

Therefore, given any $\epsilon > 0$, having a polynomial-time algorithm $\mathcal{A}$ with approximation ratio $\rho - \epsilon$ would imply that this algorithm could be used to decide if there is a subset $A' \subset A$ such that $\sum_{a \in A'} s(a) = M/2$. In other words, this algorithm would be able to solve the partition problem. This contradicts the fact that the partition problem is NP-hard and no polynomial time algorithm solves it unless P $=$ NP. Therefore, there is no algorithm that achieves a $\rho - \epsilon = \frac{3}{2} \cdot \frac{\alpha - 1 + (\frac{2}{3})^{\alpha}}{\alpha} - \epsilon$ approximation ratio for the $(C, \cdot)$-VMA problem unless P $=$ NP. ■

4.3.3.2 Upper bound on the approximation ratio for $x^* \geq C$

We study now an upper bound on the competitive ratio of the $(C, \cdot)$-VMA problem for the case when $x^* \geq C$. Under this condition, the best is to load each PM to its full capacity. Intuitively, an optimal solution should load every machine up to its maximum capacity or, if not possible, should balance the load among PMs to maximize the average load. The following lemma formalizes this observation.

Lemma 4. *For any system with unbounded number of PMs where $x^* \geq C$ the power consumption of the optimal assignment π^* is lower bounded by the power consumption of a (possibly not feasible) solution where $\ell(D)$ is evenly distributed among $\overline{m}$ PMs, where $\overline{m}$ is the minimum number of PMs required to allocate all VMs (i.e., the optimal solution of the packing problem). That is,*

$$P(\pi^*) \geq \overline{m} \cdot b + \overline{m}(\ell(D)/\overline{m})^{\alpha}.$$

Proof: Denote the number of PMs used in an optimal $(C, \cdot)$-VMA solution π^* by m^*. By Corollary 1, we know that

$$P(\pi^*) \geq m^* b + m^*(\ell(D)/m^*)^{\alpha}.$$

Given that $\overline{m} \leq m^*$, we know that $\ell(D)/m^* \leq \ell(D)/\overline{m} \leq C \leq x^*$. Thus, for evenly-balanced loads the power consumption is reduced if the number of PMs is reduced, that is

$$m^* b + m^*(\ell(D)/m^*)^{\alpha} \geq \overline{m} \cdot b + \overline{m}(\ell(D)/\overline{m})^{\alpha}.$$

Hence, the claim follows. ■

Now we prove an upper bound on the approximation ratio showing a reduction to bin packing [48]. The reduction works as follows. Let each PM be seen as a bin of capacity C, and each

VM be seen as an object to be placed in the bins, whose size is the VM load. Then, a solution for this bin packing problem instance yields a feasible (perhaps suboptimal) solution for the instance of $(C, \cdot)$-VMA. Moreover, using any bin-packing approximation algorithm, we obtain a feasible solution for $(C, \cdot)$-VMA that approximates the minimal number of PMs used. The power consumption of this solution approximates the power consumption of the optimal solution π^* of the instance of $(C, \cdot)$-VMA. In order to compute an upper bound on the approximation ratio of this algorithm, we will compare the power consumption of such solution against a lower bound on the power consumption of π^*. The following theorem shows the approximation ratio obtained.

Theorem 4. *For every $\epsilon > 0$, there exists an approximation algorithm for the $(C, \cdot)$-VMA problem when $x^* \geq C$ that achieves an approximation ratio of*

$$\rho < 1 + \epsilon + \frac{C^\alpha}{b} + \frac{1}{\overline{m}},$$

where $\overline{m}$ is the minimum number of PMs required to allocate all the VMs.

Proof: Consider an instance of the $(C, \cdot)$-VMA problem. If $\ell(D) \leq C$, the optimal solution is to place all the VMs in one single PM. Hence, we assume in the rest of the proof that $\ell(D) > C$. Define the corresponding instance of bin packing following the reduction described above. Let the optimal number of bins to accommodate all VMs be $\overline{m}$. As shown in [45], for every $\epsilon > 0$, there is a polynomial-time algorithm that fits all VMs in $\widehat{m}$ bins, where $\widehat{m} \leq (1+\epsilon)\overline{m} + 1$. From Lemma 2, once the number of PMs used $\widehat{m}$ is fixed, the power consumption is maximized when the load is unbalanced to the maximum. I.e., the power consumption of the assignment is at most $\widehat{m}b + (\ell(D)/C)C^\alpha$. On the other hand, as shown in Lemma 4, the power consumption of the optimal $(C, \cdot)$-VMA solution is at least $\overline{m} \cdot b + \overline{m}\left(\frac{\ell(D)}{\overline{m}}\right)^\alpha$. Then, we compute a bound on the approximation ratio as follows.

$$\rho \leq \frac{\widehat{m}b + \left(\frac{\ell(D)}{C}\right) C^\alpha}{\overline{m} \cdot b + \overline{m}\left(\frac{\ell(D)}{\overline{m}}\right)^\alpha} \tag{4.1}$$

$$< \frac{\widehat{m}b + \left(\frac{\ell(D)}{C}\right) C^\alpha}{\overline{m} \cdot b + \overline{m}\left(\frac{C}{2}\right)^\alpha}, \tag{4.2}$$

where the second inequality comes from $\ell(D)/\overline{m} > C/2$. (If $\ell(D)/\overline{m} \leq C/2$, there must be two PMs whose loads add up to less than C, which contradicts the fact that $\overline{m}$ is the number of bins used in the optimal solution of bin packing.) Let $\gamma = (x^*/C)^\alpha$. Then, replacing $b = \gamma C^\alpha(\alpha - 1)$, in Eq. (4.1) we have

$$\begin{aligned}
\rho &< \frac{\widehat{m}\gamma C^\alpha(\alpha-1) + \left(\frac{\ell(D)}{C}\right) C^\alpha}{\overline{m}\gamma C^\alpha(\alpha-1) + \overline{m}\left(\frac{C}{2}\right)^\alpha} \\
&= \frac{\widehat{m}\gamma(\alpha-1) + \left(\frac{\ell(D)}{C}\right)}{\overline{m}\gamma(\alpha-1) + \overline{m}\left(\frac{1}{2}\right)^\alpha} \leq \frac{\widehat{m}\gamma(\alpha-1) + \overline{m}}{\overline{m}\gamma(\alpha-1) + \left(\frac{\overline{m}}{2^\alpha}\right)} \qquad (4.3) \\
&\leq \frac{(\overline{m}(1+\epsilon)+1))\gamma(\alpha-1) + \overline{m}}{\overline{m}\gamma(\alpha-1) + \left(\frac{\overline{m}}{2^\alpha}\right)} \qquad (4.4) \\
&= \frac{(1+\epsilon)\gamma(\alpha-1)+1}{\gamma(\alpha-1) + \left(\frac{1}{2^\alpha}\right)} + \frac{\gamma(\alpha-1)}{\overline{m}\gamma(\alpha-1) + \left(\frac{\overline{m}}{2^\alpha}\right)} \\
&= \frac{2^\alpha((1+\epsilon)\gamma(\alpha-1)+1)}{2^\alpha\gamma(\alpha-1)+1} + \frac{2^\alpha\gamma(\alpha-1)}{\overline{m}(2^\alpha\gamma(\alpha-1)+1)} \\
&< \frac{(1+\epsilon)\gamma(\alpha-1)+1}{\gamma(\alpha-1)} + \frac{1}{\overline{m}} \\
&= 1+\epsilon+\frac{1}{\gamma(\alpha-1)} + \frac{1}{\overline{m}} = 1+\epsilon+\frac{C^\alpha}{b} + \frac{1}{\overline{m}}
\end{aligned}$$

Inequality (4.3) follows from $\ell(D)/C \leq \overline{m}$, Inequality (4.4) from the approximation algorithm for bin packing, and the last inequality is because $\overline{m} > 0$. ■

4.3.3.3 Upper bound on the approximation ratio for $x^* < C$

We study now the $(C, \cdot)$-VMA problem when $x^* < C$. In this case, the optimal load per PM is less than its capacity, so an optimal solution would load every PM to x^* if possible, or try to balance the load close to x^*. In this case we slightly modify the bin packing algorithm described above, reducing the bin size from C to x^*. Then, using an approximation algorithm for this bin packing problem, the following theorem can be shown.

Theorem 5. *For every $\epsilon > 0$, there exists an approximation algorithm for the $(C, \cdot)$-VMA problem when $x^* < C$ that achieves an approximation ratio of*

$$\rho < \frac{\overline{m}}{m^*}\left((1+\epsilon) + \frac{1}{\alpha-1}\right) + \frac{1}{m^*},$$

where m^* is the number of PMs used by the optimal solution of $(C, \cdot)$-VMA, and $\overline{m}$ is the minimum number of PMs required to allocate all the VMs without exceeding load x^* (i.e., the optimal solution of the bin packing problem).

Proof: Consider an instance of the $(C, \cdot)$-VMA problem. If $\ell(D) \leq x^*$ then the optimal solution is to assign all the VMs to one single PM. Then, in the rest of the proof we assume that $\ell(D) > x^*$. Assuming m^* to be the number of PMs of an optimal $(C, \cdot)$-VMA solution π^* for load $\ell(D)$, from Corollary 1, we can claim that the power consumption $P(\pi^*)$ can be bounded as $P(\pi^*) \geq m^* b + m^*(\ell(D)/m^*)^\alpha$.

Now, let $\overline{m}$ be the minimum number of PMs required to allocate all the VMs of the $(C, \cdot)$-VMA problem without exceeding load x^*. As shown in [45], for every $\epsilon > 0$, there is a polynomial-time algorithm that fits all VMs in $\widehat{m}$ bins, where $\widehat{m} \leq (1+\epsilon)\overline{m}+1$. From Lemma 2, this approximation results in a power consumption no larger than $\widehat{m}b + (\ell(D)/x^*)(x^*)^\alpha$. Hence, the approximation ratio ρ of the solution obtained wit this algorithm can be bounded as follows.

$$\rho \leq \frac{\widehat{m}b + \left(\frac{\ell(D)}{x^*}\right)(x^*)^\alpha}{m^*b + m^*\left(\frac{\ell(D)}{m^*}\right)^\alpha}. \tag{4.5}$$

Since $\ell(D) > x^*$, we know that $\ell(D)/m^* > x^*/2$, since otherwise there are two used PMs whose load is no larger than x^*, contradicting by Lemma 1 the definition of m^*. Also, from the definition of $\overline{m}$, it follows that $\ell(D) \leq \overline{m} \cdot x^*$. Finally, recall that $b = (x^*)^\alpha(\alpha - 1)$. Applying these results to Eq. (4.5) we have the following.

$$\begin{aligned}
\rho &< \frac{\widehat{m}(x^*)^\alpha(\alpha-1) + \left(\frac{x^*\overline{m}}{x^*}\right)(x^*)^\alpha}{m^*(x^*)^\alpha(\alpha-1) + m^*\left(\frac{x^*}{2}\right)^\alpha} \\
&= \frac{\widehat{m}(\alpha-1)+\overline{m}}{m^*(\alpha-1)+m^*\left(\frac{1}{2}\right)^\alpha} \leq \frac{(\overline{m}(1+\epsilon)+1)(\alpha-1)+\overline{m}}{m^*(\alpha-1)+\frac{m^*}{2^\alpha}} \\
&= \frac{\overline{m}(1+\epsilon)(\alpha-1)+\overline{m}}{m^*(\alpha-1)+\frac{m^*}{2^\alpha}} + \frac{\alpha-1}{m^*(\alpha-1)+\frac{m}{2^\alpha}} \\
&= \frac{\overline{m}}{m^*}\frac{2^\alpha((1+\epsilon)(\alpha-1)+1)}{2^\alpha(\alpha-1)+1} + \frac{2^\alpha(\alpha-1)}{2^\alpha m^*(\alpha-1)+m^*} \\
&\leq \frac{\overline{m}}{m^*}\left((1+\epsilon)+\frac{1}{\alpha-1}\right)+\frac{1}{m^*},
\end{aligned}$$

where the first inequality comes from applying the results aforementioned, and second one from using $\widehat{m} = \overline{m}(1+\epsilon)+1$, while the last one results from simplifying the previous equation. ■

4.4 Online Analysis

In this section, we study the online version of the VMA problem, i.e., when the VMs are revealed one by one. We first study lower bounds and then provide online algorithms and prove upper bounds on their competitive ratio.

4.4.1 Lower Bounds

In this section, we compute lower bounds on the competitive ratio for $(\cdot, \cdot)$-VMA, $(C, \cdot)$-VMA, $(\cdot, m)$-VMA, (C, m)-VMA and $(\cdot, 2)$-VMA problems. We start with one general construction that is used to obtain lower bounds on the first four cases. Then, we develop special constructions for $(\cdot, m)$-VMA and $(\cdot, 2)$-VMA that improve the lower bounds for these two problems.

4.4.1.1 General Construction

We prove lower bounds on the competitive ratio of $(\cdot,\cdot)$-VMA, $(C,\cdot)$-VMA, $(\cdot, m)$-VMA and (C, m)-VMA problems. These lower bounds are shown in the following two theorems. In Theorem 6, we prove a lower bound on the competitive ratio that is valid in the cases when C is unbounded and when it is larger or equal than x^*. The case $C \leq x^*$ is covered in Theorem 7.

Theorem 6. *There exists an instance of problems $(\cdot,\cdot)$-VMA, $(\cdot, m)$-VMA, $(C,\cdot)$-VMA and (C, m)-VMA when $C > x^*$, such that no online algorithm can guarantee a competitive ratio smaller than $\frac{(3/2)2^\alpha - 1}{2^\alpha - 1}$.*

Proof: We consider a scenario where, for any online algorithm, an adversary injects VMs of size ϵx^* ($\epsilon > 0$ is an arbitrarily small constant) to the system until the algorithm starts up a new PM. Let us assume that the total number of VMs injected is k. According to the adversary's behavior, the assignment of the VMs should be that all the VMs except one are allocated to a single PM while the second PM has only one VM. Depending on what the optimal solution is, we discuss the following two cases:

Case 1: $k \leq \frac{1}{\epsilon}\left(\frac{\alpha-1}{1-2^{1-\alpha}}\right)^{1/\alpha}$. The optimal solution will allocate all the VMs to a single PM. Consequently, the competitive ratio of the online algorithm satisfies

$$\rho(k) \geq \lim_{\epsilon \to 0}\left(\frac{((k-1)\epsilon x^*)^\alpha + (\epsilon x^*)^\alpha + 2b}{(k\epsilon x^*)^\alpha + b}\right).$$

It can be easily verified that function $\rho(k)$ is monotone decreasing with k. That is, $\rho(k)$ is minimized when $k = \frac{1}{\epsilon}\left(\frac{\alpha-1}{1-2^{1-\alpha}}\right)^{1/\alpha}$. As a result, we obtain,

$$\begin{aligned}
\rho(k) &\geq \lim_{\epsilon \to 0}\left(\frac{\left(\left(\frac{\alpha-1}{1-2^{1-\alpha}}\right)^{1/\alpha} x^*\right)^\alpha + (\epsilon x^*)^\alpha + 2b}{\left(\left(\frac{\alpha-1}{1-2^{1-\alpha}}\right)^{1/\alpha} x^*\right)^\alpha + b}\right) \\
&= \frac{\left(\left(\frac{\alpha-1}{1-2^{1-\alpha}}\right)^{1/\alpha} x^*\right)^\alpha + 2(x^*)^\alpha(\alpha-1)}{\left(\left(\frac{\alpha-1}{1-2^{1-\alpha}}\right)^{1/\alpha} x^*\right)^\alpha + (x^*)^\alpha(\alpha-1)} \\
&= \frac{3 - 2^{1-\alpha}}{2 - 2^{1-\alpha}} = \frac{(3/2)2^\alpha - 1}{2^\alpha - 1}.
\end{aligned}$$

Case 2: $k > \frac{1}{\epsilon}\left(\frac{\alpha-1}{1-2^{1-\alpha}}\right)^{1/\alpha}$. The optimal solution will use two PMs with $k/2$ PMs assigned to each PM. Accordingly, the competitive ratio of the online algorithm satisfies

$$\rho(k) \geq \lim_{\epsilon \to 0} \left(\frac{((k-1)\epsilon x^*)^\alpha + (\epsilon x^*)^\alpha + 2b}{2\left(\frac{k\epsilon x^*}{2}\right)^\alpha + 2b} \right).$$

Similarly, we observe that $\rho(k)$ is monotone increasing with k. Consequently, the following inequality applies.

$$\begin{aligned}
\rho(k) &\geq \lim_{\epsilon \to 0} \left(\frac{\left(\left(\frac{\alpha-1}{1-2^{1-\alpha}}\right)^{1/\alpha} x^*\right)^\alpha + (\epsilon x^*)^\alpha + 2b}{2\left(\frac{1}{2}\left(\frac{\alpha-1}{1-2^{1-\alpha}}\right)^{1/\alpha} x^*\right)^\alpha + 2b} \right) \\
&= \frac{\left(\left(\frac{\alpha-1}{1-2^{1-\alpha}}\right)^{1/\alpha} x^*\right)^\alpha + 2(x^*)^\alpha(\alpha-1)}{2\left(\frac{1}{2}\left(\frac{\alpha-1}{1-2^{1-\alpha}}\right)^{1/\alpha} x^*\right)^\alpha + 2(x^*)^\alpha(\alpha-1)} \\
&= \frac{3-2^{1-\alpha}}{2-2^{1-\alpha}} = \frac{(3/2)2^\alpha - 1}{2^\alpha - 1}
\end{aligned}$$

Note that it can also happen that $C < \left(\frac{\alpha-1}{1-2^{1-\alpha}}\right)^{1/\alpha} x^*$. In this case, k is smaller than $\frac{1}{\epsilon}\left(\frac{\alpha-1}{1-2^{1-\alpha}}\right)^{1/\alpha}$. Therefore, the competitive ratio is always larger than $\frac{(3/2)2^\alpha-1}{2^\alpha-1}$, proving the lower bound. ■

Theorem 7. *There exists an instance of problems $(C, \cdot)$-VMA and (C, m)-VMA when $C \leq x^*$ such that no online algorithm can guarantee a competitive ratio smaller than $(C^\alpha + 2b)/(b + \max(C^\alpha, 2(C/2)^\alpha + b))$.*

Proof: Similarly to the proof of Theorem 6, we prove the result by considering an adversarial injection of VMs of size ϵC. This injection stops when a new PM started up by an online algorithm. We discuss the following two cases:
Case 1: $k \leq 1/\epsilon$. In this case, the optimal algorithm will assign all the VMs to a single PM. The competitive ratio of the online algorithm satisfies

$$\begin{aligned}
\rho(k) &\geq \lim_{\epsilon \to 0} \frac{((k-1)\epsilon C)^\alpha + (\epsilon C)^\alpha + 2b}{(k\epsilon C)^\alpha + b} \\
&\geq \lim_{\epsilon \to 0} \frac{(1-\epsilon)^\alpha C^\alpha + 2b}{C^\alpha + b} \\
&\geq \frac{C^\alpha + 2b}{C^\alpha + b} \geq 2 - \frac{1}{\alpha}
\end{aligned}$$

The second inequality results from applying $k \leq 1/\epsilon$, which is observed from the monotone decreasing property of function $\rho(k)$. The last inequality comes from computing the limit when ϵ goes to 0 and by applying $b \geq C^\alpha(\alpha - 1)$.
Case 2: $k > 1/\epsilon$. In this case, the adversary stops injecting VMs as there will be, mandatorily,

two active PMs, one of them not capable to allocate more VMs and the second one hosting one single VM. Since all the VMs can not be consolidated to a single PM. The optimal solution would use also two PMs but evenly balancing the loads among them. The competitive ratio of the online algorithm satisfies

$$\begin{aligned}\rho(k) &= \lim_{\epsilon\to 0}\left(\frac{((k-1)\epsilon C)^\alpha + (\epsilon C)^\alpha + 2b}{2\left(\frac{k\epsilon C}{2}\right)^\alpha + 2b}\right)\\ &= \lim_{\epsilon\to 0}\left(\frac{C^\alpha + (\epsilon C)^\alpha + 2b}{2\left(\frac{C+\epsilon C}{2}\right)^\alpha + 2b}\right)\\ &= \frac{C^\alpha + 2b}{2\left(\frac{C}{2}\right)^\alpha + 2b}.\end{aligned}$$

Hence, combining the results from both cases 1 and 2 we obtain the bound presented in Theorem 7. ■

4.4.1.2 Special Constructions for $(\cdot, m)$-VMA and $(\cdot, 2)$-VMA

We show first that for m PMs there is a lower bound on the competitive ratio that improves the previous lower bound when $\alpha > 4.5$. Secondly, we prove a particular lower bound for problem $(\cdot, 2)$-VMA, that improves the previous lower bound when $\alpha > 3$.

Theorem 8. *There exists an instance of problem $(\cdot, m)$-VMA such that no online algorithm can guarantee a competitive ratio smaller than $3^\alpha/(2^{\alpha+2}+\epsilon)$ for any $\epsilon > 0$.*

Proof: We prove the result by giving an adversarial arrival of VMs. We evaluate the competitive ratio of any online algorithm ALG with respect to an algorithm OPT that distributes the VMs among all the PMs "as evenly as possible". We define a value $\beta > 1$ such that $\epsilon \geq (\alpha-1)/\beta^\alpha$ for some value $\epsilon > 0$. Note that such value β can be defined for any $\epsilon > 0$. The adversarial arrival follows. In a first phase, m VMs arrive, each with load βx^*.

Let π be the partition given by ALG. We show first that if π uses less than $3m/4$ PMs[2] or some PM is assigned more than 2 VMs there exists another partition that can be obtained from π, it uses exactly $3m/4$ PMs, no PM is assigned more than 2 VMs, and the power consumption is not worse.

If π uses less than $3m/4$ PMs, then there exists another partition π' that uses exactly $3m/4$ PMs with a power consumption that is not worse than $P(\pi)$. To see why, notice that there are PMs in π that are assigned more than one VM and that each load is $\beta x^* > x^*$. Then, applying repeatedly Lemma 1 until $3m/4$ PMs are used, where ℓ_1 and ℓ_2 are the loads of any pair of VMs assigned to the same PM, a partition π' such that $P(\pi') \leq P(\pi)$ can be obtained.

If in π' some PM is assigned more than 2 VMs, then there exists another partition π'' where no PM is assigned more than 2 VMs with a power consumption that is not worse than $P(\pi')$. To

[2]For clarity we omit floors and ceilings in the proof.

see why, consider the following reassignment procedure. Repeatedly until there is no such PM, locate a PM s_i with at least 3 VMs. Then, locate a PM s_j with one single VM (which exists by the pigeonhole principle). Then, move one VM from s_i to s_j. From Lemma 2 each movement decreases the power consumed. Hence, π'' is still a partition that uses $3m/4$ PMs, each PM has at most 2 VMs assigned, and $P(\pi'') \leq P(\pi')$.

Then, we know that $P(\pi)$ is not smaller than the power consumption of a partition where exactly $3m/4$ PMs are used and no PM is assigned more than 2 VMs. On the other hand, OPT would have assigned each VM to a different PM. Thus, using that $x^* = (b/(\alpha-1))^{1/\alpha}$, the competitive ratio is

$$\begin{aligned} \rho &\geq \frac{(2\beta x^*)^\alpha m/4 + (\beta x^*)^\alpha m/2 + 3mb/4}{m(\beta x^*)^\alpha + mb} \\ &\geq \frac{(2^{\alpha-2}+1/2)\beta^\alpha}{\beta^\alpha + (\alpha-1)} \geq 2^{\alpha-3} + 1/4, \end{aligned}$$

where the last inequality follows from $\beta^\alpha \geq (\alpha-1)$. Finally, observe that $2^{\alpha-3} + 1/4 \geq 3^\alpha/(2^{\alpha+2}+\epsilon)$ for $\alpha > 1$. No more VMs arrive in this case.

Let us consider now the the case where ALG assigns the m initial VMs to more than $3m/4$ PMs. Then, after ALG has assigned the first m VMs, a second batch of $m/2$ VMs arrive, each VM with load $2\beta x^*$. Let π be the partition output by ALG after this second batch is assigned. If in π two of the second batch VMs are assigned to the same PM s_i, by the pigeonhole principle there is at least one PM s_j with at most load βx^*. Then, from Lemma 2, the power consumed is reduced if one of the new VMs is moved from s_i to s_j. After repeating this process as many times as possible, a partition π' is obtained where each of the VMs of the second batch is assigned to a different PM, and $P(\pi') \leq P(\pi)$. Since ALG used more than $3m/4$ PMs in the first batch, in π', there are at least $m/4$ PMs with load $3\beta x^*$. On the other hand, OPT can distribute all the VMs in such a way that each PM has a load of $2\beta x^*$. Thus, the bound on the competitive ratio is as follows.

$$\rho \geq \frac{m(3\beta x^*)^\alpha/4}{m(2\beta x^*)^\alpha + mb} \geq \frac{3^\alpha}{2^{\alpha+2}+\epsilon},$$

where the last inequality follows from $\epsilon \geq (\alpha-1)/\beta^\alpha$. ■

Now, we show a stronger lower bound on the competitive ratio for $(\cdot, 2)$-VMA problem.

Theorem 9. *There exists an instance of problem* $(\cdot, 2)$*-VMA such that no online algorithm can guarantee a competitive ratio smaller than* $3^\alpha/2^{\alpha+1}$.

Proof: We prove the result by showing an adversarial arrival of VM. We evaluate the competitive ratio of any online algorithm ALG with respect to an optimal algorithm OPT that knows the future VM arrivals. The adversarial arrival follows. In a first phase two VM d_1 and d_2 arrive, with loads $\ell(d_1) = \ell(d_2) = 6x^*$ (Recall from Section 4.2 that $x^* = (b/(\alpha-1))^{1/\alpha}$).

If ALG assigns both VMs to the same PM, the power consumed will be $(12x^*)^\alpha + b$, whereas OPT would assign them to different PMs, with a power consumption of $2((6x^*)^\alpha + b)$. Hence, the ratio ρ would be

$$\begin{aligned}\rho &= \frac{(12x^*)^\alpha + b}{2((6x^*)^\alpha + b)} > \frac{12^\alpha}{2(6^\alpha + \alpha - 1)} \\ &> \frac{12^\alpha}{2(6^\alpha + 2^\alpha)} = \frac{6^\alpha}{2(3^\alpha + 1)},\end{aligned}$$

where the first inequality follows from $\alpha > 1$ and the second from $\alpha - 1 < 2^\alpha$ for any $\alpha > 1$. It is enough to prove that $6^\alpha/(2(3^\alpha + 1)) \geq (3/2)^\alpha/2$, or equivalently $4^\alpha \geq 3^\alpha + 1$, which is true for any $\alpha > 1$. Then, there are no new VM arrivals.

If, otherwise, ALG assigns each VM d_1 and d_2 to a different PM, then a third VM d_3 arrives, with load $\ell(d_3) = 12x^*$. Then, ALG must assign it to one of the PMs. Independently of which PM is used, the power consumption of the final configuration is $(18x^*)^\alpha + (6x^*)^\alpha + 2b$. On its side, OPT assigns d_1 and d_2 to one PM, and d_3 to the other, with a power consumption of $2((12x^*)^\alpha + b)$. Hence, the competitive ratio ρ is

$$\begin{aligned}\rho &= \frac{(18x^*)^\alpha + (6x^*)^\alpha + 2b}{2((12x^*)^\alpha + b)} > \frac{18^\alpha + 6^\alpha}{2(12^\alpha + \alpha - 1)} \\ &> \frac{18^\alpha + 6^\alpha}{2(12^\alpha + 4^\alpha)} \geq (3/2)^\alpha/2,\end{aligned}$$

where the first inequality follows from $\alpha > 1$, the second from $\alpha - 1 < 4^\alpha$ for any $\alpha > 1$, and the third from $(9^\alpha + 3^\alpha)/(6^\alpha + 2^\alpha) \geq (3/2)^\alpha$, what can be checked to be true. Then, there are no new VM arrivals and the claim follows. ■

4.4.2 Upper Bounds

Now, we study upper bounds for $(\cdot,\cdot)$-VMA, $(C,\cdot)$-VMA, and $(\cdot,2)$-VMA problems. We start giving two online VMA algorithms, that we denote as *VMA1* and *VMA2* and that can be found in Algorithm 1 and Algorithm 2. We study algorithm *VMA1* for both $(\cdot,\cdot)$-VMA and $(C,\cdot)$-VMA problems while *VMA2* is only studied for the $(\cdot,\cdot)$-VMA case. These algorithms use the load of the new revealed VM in order to decide the PM where it will be assigned. The behaviour of both algorithms is similar. In algorithm *VMA1*, if the load of the revealed VM is strictly larger than $\min\{x^*, C\}/2$, the algorithm assigns this VM to a new PM without any other VM already assigned to it. Otherwise, the algorithm schedules the revealed VM to any loaded PM whose current load is smaller or equal than $\frac{\min\{x^*,C\}}{2}$. Hence, when this new VM is assigned, the load of this PM remains smaller than $\min\{x^*, C\}$. If there is no such loaded PM, the revealed VM is assigned to a new PM. On the other hand, algorithm *VMA2* shares the same behaviour but uses a different threshold. In this case, if the load of the revealed VM is strictly larger than $\min\{x^*, C\}$, the algorithm assigns this VM to a new PM. If the load is smaller than x^*, *VMA2* schedules the new VM to the most loaded PM whose load is smaller than x^*.

Note that, since the case under consideration assumes the existence of an unbounded number of PMs, there exists always one new PM. As mentioned, a detailed description of these algorithms is shown in Algorithm 1 and Algorithm 2. As before, A_j denotes the set of VMs assigned to PM s_j at a given time.

Algorithm 1: Online algorithm *VMA1* for $(\cdot,\cdot)$-VMA and $(C,\cdot)$-VMA problems.

for *each VM* d_i **do**
  if $\ell(d_i) > \frac{\min\{x^*,C\}}{2}$ **then**
    d_i is assigned to a new PM
  else
    d_i is assigned to any loaded PM s_j where $\ell(A_j) \leq \frac{\min\{x^*,C\}}{2}$. If such loaded PM does not exist, d_i is assigned to a new PM.

Algorithm 2: Online algorithm *VMA2* for the $(\cdot,\cdot)$-VMA problem.

for *each VM* d_i **do**
  if $\ell(d_i) \geq x^*$ **then**
    d_i is assigned to a new PM
  else
    d_i is assigned to the PM s_j such that $\ell(A_k) \leq \ell(A_j) < x^*$ for all k. If such loaded PM does not exist, d_i is assigned to a new PM.

4.4.2.1 Upper Bounds for the $(\cdot,\cdot)$-VMA and $(C,\cdot)$-VMA problems

We prove the approximation ratio of Algorithm 1 in the following two theorems.

Theorem 10. *There exists an online algorithm for $(\cdot,\cdot)$-VMA and $(C,\cdot)$-VMA when $x^* < C$ that achieves the following competitive ratio:*

$$\begin{aligned}\rho &= 1, \text{ if no VM } d_i \text{ has load such that } \ell(d_i) < x^*,\\ \rho &\leq \left(1 - \frac{1}{\alpha}\left(1 - \frac{1}{2^\alpha}\right)\right)\left(2 + \frac{x^*}{\ell(D_s)}\right), \text{ otherwise.}\end{aligned}$$

Proof: We proceed with the analysis of the competitive ratio of Algorithm 1 shown above. Let us first consider an optimal algorithm, that is, an algorithm that gives an optimal solution for any instance. Let us denote by π^* the optimal solution obtained by the optimal algorithm, and A_i the load assigned to PM s_i in that solution, for a particular instance of VMA problem. Furthermore, load A_i is decomposed in $d_{i_1}, d_{i_2}, \ldots, d_{i_{k_i}}$, where each d_{i_j} is a VM that π^* assigns to s_i. Using simple algebra, it holds:

$$f(\ell(A_i)) = \frac{f(\ell(A_i))}{\ell(A_i)}(\ell(d_{i_1}) + \ell(d_{i_2}) + \cdots + \ell(d_{i_{k_i}})).$$

It is possible now to split the set A_i in two sets, one with those VMs assigned to s_i whose load is strictly smaller than x^* and a second set that contains those VMs assigned to s_i whose load is bigger than x^*. In terms of notation, we say that A_i is split in B_i and S_i (where B stands for Big loads and S stands for Small loads). Therefore, it also holds:

$$f(\ell(A_i)) = \sum_{d_{i_j} \in B_i} \frac{f(\ell(A_i))}{\ell(A_i)} \ell(d_{i_j}) + \sum_{d_{i_j} \in S_i} \frac{f(\ell(A_i))}{\ell(A_i)} \ell(d_{i_j}).$$

On the other hand, by definition of x^*, it holds that:

$$f(\ell(A_i))/\ell(A_i) \geq f(x^*)/x^*$$

for all i (indeed, for any load). Moreover, if a PM has been assigned with a load $\ell(d_{i_j})$ bigger than x^*, it also holds that

$$f(\ell(A_i))/\ell(A_i) \geq f(\ell(d_{i_j}))/\ell(d_{i_j}).$$

Hence, we obtain the following inequality:

$$f(\ell(A_i)) \geq \sum_{d_{i_j} \in B_i} f(\ell(d_{i_j})) + \sum_{d_{i_j} \in S_i} \frac{f(x^*)}{x^*} \ell(d_{i_j}).$$

In order to lower bound the power consumption of the solution π^*, we plug the above inequality into the corresponding equation:

$$\begin{aligned} P(\pi^*) &= \sum_{A_i \neq \emptyset} f(\ell(A_i)) \\ &\geq \sum_{A_i \neq \emptyset} \sum_{d_{i_j} \in B_i} f(\ell(d_{i_j})) + \frac{f(x^*)}{x^*} \sum_{A_i \neq \emptyset} \sum_{d_{i_j} \in S_i} \ell(d_{i_j}), \end{aligned}$$

or, equivalently expressed in more compact notation:

$$P(\pi^*) \geq \sum_{d_i : \ell(d_i) \geq x^*} f(\ell(d_i)) + \frac{f(x^*)}{x^*} \sum_{d_i : \ell(d_i) < x^*} \ell(d_i). \tag{4.6}$$

Consider now Algorithm 1. Let us denote by π a solution that Algorithm 1 gives for a particular instance. Also, let us denote by $\hat{A}_i$ the load assigned by Algorithm 1 to PM s_i. Note that due to the design of the algorithm, after the last VM has been assigned, either there is only one loaded PM whose current load is smaller than $x^*/2$, or every loaded PM has a load at least $x^*/2$. We study these two cases separately.

Case 1: $\ell(\hat{A}_i) \geq x^*/2$ for all i. In this case, in a solution provided by π there are PMs with two types of load: those that are loaded with one VM whose load is no smaller than x^*, and those that are loaded with VMs whose load is strictly smaller than x^*, nonetheless, their total load is bigger than $x^*/2$. Note that due to the design of the algorithm, none of the PMs in the second group has a load bigger than x^*. Let us denote by B the set of VMs with load at least x^*, and D_s the set of VMs with load less than x^*. Therefore, it holds:

$$\begin{aligned} P(\pi) &= \sum_{d \in B} f(\ell(d)) + \sum_{\frac{x^*}{2} \leq \ell(\hat{A}_i) \leq x^*} f(\ell(\hat{A}_i)) \\ &\leq \sum_{d \in B} f(\ell(d)) + \frac{f(\frac{x^*}{2})}{\frac{x^*}{2}} \ell(D_s). \end{aligned}$$

Computing the ratio ρ between $P(\pi)$ and $P(\pi^*)$, we obtain the following inequality:

$$\begin{aligned} \rho &\leq \frac{\sum_{d \in B} f(\ell(d)) + \frac{f(\frac{x^*}{2})}{\frac{x^*}{2}} \ell(D_s)}{\sum_{d \in B} f(\ell(d)) + \frac{f(x^*)}{x^*} \ell(D_s)} \\ &\leq \frac{\frac{f(\frac{x^*}{2})}{\frac{x^*}{2}} \ell(D_s)}{\frac{f(x^*)}{x^*} \ell(D_s)} \\ &= 2 \frac{f(\frac{x^*}{2})}{f(x^*)} = 2 \left(1 - \frac{1}{\alpha} \left(1 - \frac{1}{2^\alpha}\right)\right). \end{aligned}$$

Case 2: there exists s_i such that $\ell(\hat{A}_i) < x^*/2$. In this case, π gives solutions with three types of loaded PMs: those that are loaded with one VM whose load is bigger than x^*, those that are loaded with VMs whose load is strictly smaller than x^*, but which total load is at least $x^*/2$, and one PM whose total load is is strictly smaller than $x^*/2$. Let us denote such a PM by s'. Therefore, it holds:

$$\begin{aligned} P(\pi) &= \sum_{d \in B} f(\ell(d)) + \sum_{\frac{x^*}{2} \leq \ell(\hat{A}_i) \leq x^*} f(\ell(\hat{A}_i)) + f(\ell(\hat{A}_{s'})) \\ &\leq \sum_{d \in B} f(\ell(d)) + \frac{f(\frac{x^*}{2})}{\frac{x^*}{2}} \Big(\ell(D_s) - \ell(\hat{A}_{s'})\Big) + f(\ell(\hat{A}_{s'})) \\ &= \sum_{d \in B} f(\ell(d)) + \frac{f(\frac{x^*}{2})}{\frac{x^*}{2}} \Big(\ell(D_s) - \ell(\hat{A}_{s'})\Big) + \ell(\hat{A}_{s'})^\alpha + b. \end{aligned}$$

Let us denote the latter expression by $\Pi(\pi)$. Computing the ratio ρ between $P(\pi)$ and $P(\pi^*)$, we obtain the following inequality:

$$\begin{aligned}
\rho &\leq \frac{\Pi(\pi)}{\sum_{d\in B} f(\ell(d)) + \frac{f(x^*)}{x^*}\ell(D_s)} \\
&\leq 2\left(1-\frac{1}{\alpha}\left(1-\frac{1}{2^\alpha}\right)\right) + \frac{\ell(\hat{A}_{s'})^\alpha - \ell(\hat{A}_{s'})\frac{f(\frac{x^*}{2})}{\frac{x^*}{2}} + b}{\frac{f(x^*)}{x^*}\ell(D_s)} \\
&\leq 2\left(1-\frac{1}{\alpha}\left(1-\frac{1}{2^\alpha}\right)\right) + \frac{\ell(\hat{A}_{s'})^\alpha + b}{\frac{f(x^*)}{x^*}\ell(D_s)} \\
&\leq 2\left(1-\frac{1}{\alpha}\left(1-\frac{1}{2^\alpha}\right)\right) + \frac{(\frac{x^*}{2})^\alpha + b}{\frac{f(x^*)}{x^*}\ell(D_s)} \\
&= \left(1-\frac{1}{\alpha}\left(1-\frac{1}{2^\alpha}\right)\right)\left(2+\frac{x^*}{\ell(D_s)}\right).
\end{aligned}$$

Since $x^*/\ell(D_s)$ is always positive, the competitive ratio of Algorithm 1 is equal to $2^{\alpha-1} + x^*/\ell(D_s)$. Observe that, when no VM d has load $\ell(d) < x^*$, i,e., $S = \emptyset$, $P(\pi)$ and $P(\pi^*)$ are equal. Hence, the competitive ratio is 1. ■

Theorem 11. *There exists an online algorithm for $(C,\cdot)$-VMA when $x^* \geq C$ that achieves competitive ratio $\rho \leq \frac{2b}{C}\left(1+\frac{1}{(\alpha-1)2^\alpha}\right)\left(2+\frac{C}{\ell(D)}\right)$.*

Proof: We proceed with the analysis of the competitive ratio of Algorithm 1 in the case when $x^* \geq C$. The analysis uses the same technique used in the proof for the previous theorem. Hence, we just show the difference.

On the one hand, when $x^* \geq C$, it holds that $f(\ell(A_i))/\ell(A_i) \geq f(C)/C$ due to the fact that $f(x)/x$ is monotone decreasing in interval $(0, C]$. It is also obvious that all the PMs will be loaded no more C. As a result, the optimal power consumption for $(C,\cdot)$-VMA can be bounded by

$$P(\pi^*) \geq \frac{f(C)}{C}\ell(D).$$

On the other hand, the solution given by Algorithm 1 can also be upper bounded. We consider the following two cases.

Case 1: $\ell(\hat{A}_i) \geq C/2$ for all i. In this case, every PM will be loaded between $C/2$ and C. Consequently,

$$P(\pi) = \sum_{\frac{C}{2}\leq \ell(\hat{A}_i)\leq C} f(\ell(\hat{A}_i)) \leq \frac{f(\frac{C}{2})}{\frac{C}{2}}\ell(D).$$

The competitive ratio ρ then satisfies

$$\rho \leq \frac{\frac{f(\frac{C}{2})}{\frac{C}{2}}\ell(D)}{\frac{f(C)}{C}\ell(D)} = 2\frac{f(\frac{C}{2})}{f(C)} \leq \frac{2b}{C}\left(1+\frac{1}{(\alpha-1)2^\alpha}\right).$$

Case 2: there exists s_i such that $\ell(\hat{A}_i) < C/2$. In this case, it holds:

$$\begin{aligned} P(\pi) &= \sum_{\frac{C}{2} \leq \ell(\hat{A}_i) \leq C} f(\ell(\hat{A}_i)) + f(\ell(\hat{A}_{s'})) \\ &\leq \frac{f(\frac{C}{2})}{\frac{C}{2}} \Big(\sum_{d_i : \ell(d_i) \leq C} \ell(d_i) - \ell(\hat{A}_{s'}) \Big) + f(\ell(\hat{A}_{s'})) \\ &= \frac{f(\frac{C}{2})}{\frac{C}{2}} \Big(\ell(D) - \ell(\hat{A}_{s'}) \Big) + \ell(\hat{A}_{s'})^\alpha + b. \end{aligned}$$

The competitive ratio ρ then satisfies

$$\begin{aligned} \rho &\leq \frac{P(\pi)}{\frac{f(C)}{C}\ell(D)} \leq \frac{2b}{C}\left(1 + \frac{1}{(\alpha-1)2^\alpha}\right) + \\ &+ \frac{\ell(\hat{A}_{s'})^\alpha - \ell(\hat{A}_{s'})\frac{f(\frac{C}{2})}{\frac{C}{2}} + b}{\frac{f(C)}{C}\ell(D)} \\ &\leq \frac{2b}{C}\left(1 + \frac{1}{(\alpha-1)2^\alpha}\right) + \frac{\ell(\hat{A}_{s'})^\alpha + b}{\frac{f(C)}{C}\ell(D)} \\ &\leq \frac{2b}{C}\left(1 + \frac{1}{(\alpha-1)2^\alpha}\right) + \frac{(\frac{C}{2})^\alpha + b}{\frac{f(C)}{C}\ell(D)} \\ &= \frac{2b}{C}\left(1 + \frac{1}{(\alpha-1)2^\alpha}\right)\left(2 + \frac{C}{\ell(D)}\right). \end{aligned}$$

■

4.4.2.2 Approximation Ratio for Algorithm *VMA2*

We prove the approximation ratio of Algorithm 2 in the following theorem.

Theorem 12. *For a system with unbounded number of PMs, Algorithm 2 achieves a competitive ratio of* 1 *if no VM* d_i *has* $\ell(d_i) < x^*$*, and of* $2^{\alpha-1} + x^* \big/ \sum_{d_i : \ell(d_i) < x^*} \ell(d_i)$*, otherwise.*

Proof: Maintaining the same notation we used for proving Theorem 10, let us denote by π a solution that Algorithm 2 gives for a particular instance and by $\hat{A}_i$ the load assigned by Algorithm 2 to PM s_i. Note that due to the design of the algorithm, after the last VM has been assigned, either there is only one loaded PM whose current load is smaller than x^*, or every loaded PM has a load bigger than x^*. We study these two cases separately.

Case 1: $\ell(\hat{A}_i) \geq x^*$ for all i. In this case, in a solution provided by π there are PMs with two types of load: those that are loaded with one VM whose load is no smaller than x^*, and those that are loaded with VMs whose load is strictly smaller than x^*, nonetheless, their total load is bigger than x^*. Note that due to the design of the algorithm, none of the PMs in the second group has a load bigger than $2x^*$. Let us denote by B the set of demands with load at least x^*, and S the set of demands with load less than x^*. Therefore, it holds:

$$\begin{aligned} P(\pi) &= \sum_{d\in B} f(\ell(d)) + \sum_{x^*\le \ell(\hat{A}_i)\le 2x^*} f(\ell(\hat{A}_i)) \\ &\le \sum_{d\in B} f(\ell(d)) + \frac{f(2x^*)}{2x^*} \sum_{\ell(d\in S} \ell(d). \end{aligned}$$

Computing the ratio ρ between $P(\pi)$ and the expression for $P(\pi^*)$ from Eq. (4.6), we obtain the following inequality:

$$\rho \le \frac{\sum_{d\in B} f(\ell(d)) + \frac{f(2x^*)}{2x^*}\sum_{d\in S}\ell(d)}{\sum_{d\in B} f(\ell(d)) + \frac{f(x^*)}{x^*}\sum_{d\in S}\ell(d)} \le \frac{\frac{f(2x^*)}{2x^*}\sum_{d\in S}\ell(d)}{\frac{f(x^*)}{x^*}\sum_{d\in S}\ell(d)} \le 2^{\alpha-1}. \tag{4.7}$$

Case 2: there exists s_i such that $\ell(\hat{A}_i) < x^*$. In this case, π gives solutions with three types of loaded PMs: those that are loaded with one VM whose load is bigger than x^*, those that are loaded with VMs whose load is strictly smaller than x^*, but which total load is bigger than x^*, and one PM whose total load is is strictly smaller than x^*. Let us denote such a PM by s'. Therefore, it holds:

$$\begin{aligned} P(\pi) &= \sum_{d\in B} f(\ell(d)) + \sum_{x^*\le \ell(\hat{A}_i)\le 2x^*} f(\ell(\hat{A}_i)) + f(\ell(\hat{A}_{s'})) \\ &\le \sum_{d\in B} f(\ell(d)) + \frac{f(2x^*)}{2x^*}\Big(\sum_{d\in S}\ell(d) - \hat{A}_{s'}\Big) + f(\ell(\hat{A}_{s'})) \\ &= \sum_{d\in B} f(\ell(d)) + \frac{f(2x^*)}{2x^*}\Big(\sum_{d\in S}\ell(d) - \hat{A}_{s'}\Big) + \ell(\hat{A}_{s'})^\alpha + b. \end{aligned}$$

Let us denote the latter expression by $\Pi(\pi)$. Computing the ratio ρ between $P(\pi)$ and $P(\pi^*)$, we obtain the following inequality:

$$\begin{aligned} \rho &\le \frac{\Pi(\pi)}{\sum_{d\in B} f(\ell(d)) + \frac{f(x^*)}{x^*}\sum_{d\in S}\ell(d)} \\ &\le 2^{\alpha-1} + \frac{\ell(\hat{A}_{s'})^\alpha - \hat{A}_{s'}\frac{f(2x^*)}{2x^*} + b}{\frac{f(x^*)}{x^*}\sum_{d\in S}\ell(d)} \\ &\le 2^{\alpha-1} + \frac{\ell(\hat{A}_{s'})^\alpha + b}{\frac{f(x^*)}{x^*}\sum_{d\in S}\ell(d)} \\ &\le 2^{\alpha-1} + \frac{(x^*)^\alpha + b}{\frac{f(x^*)}{x^*}\sum_{d\in S}\ell(d)} \\ &= 2^{\alpha-1} + \frac{x^*}{\sum_{d\in S}\ell(d)}. \end{aligned} \tag{4.8}$$

Since $x^*/\sum_{d\in S}\ell(d)$ is always positive, the competitive ratio of Algorithm 1 is equal to $2^{\alpha-1} + x^*/\sum_{d\in S}\ell(d)$. Observe that, when no VM d has load $\ell(d) < x^*$, i,e., $S = \emptyset$, equations (4.7) and (4.8) become $\frac{P(\pi)}{P(\pi^*)} \leq 1$. ■

4.4.2.3 Upper Bounds for $(\cdot, 2)$-VMA problem

We now present an algorithm (detailed in Algorithm 3) for $(\cdot, 2)$-VMA problem and show an upper bound on its competitive ratio. A_1 and A_2 are the sets of VMs assigned to PMs s_1 and s_2, respectively, at any given time.

Algorithm 3: Online algorithm for $(\cdot, 2)$-VMA.

```
for each VM d_i do
    if ℓ(d_i) + ℓ(A_1) ≤ (b/(2^α − 2))^{1/α} or ℓ(A_1) ≤ ℓ(A_2) then
        d_i is assigned to s_1;
    else
        d_i is assigned to s_2;
```

We prove the approximation ratio of Algorithm 3 in the following theorem.

Theorem 13. *There exists an online algorithm for $(\cdot, 2)$-VMA that achieves the following competitive ratios.*

$$\rho = 1, \qquad \text{for } \ell(D) \leq \left(\frac{b}{2^\alpha - 2}\right)^{1/\alpha},$$
$$\rho \leq \max\left\{2, \left(\frac{3}{2}\right)^{\alpha-1}\right\}, \qquad \text{for } \ell(D) > \left(\frac{b}{2^\alpha - 2}\right)^{1/\alpha}.$$

Proof: Consider Algorithm 3 shown above. If $\ell(D) \leq (b/(2^\alpha - 2))^{1/\alpha}$, then the competitive ratio is 1 as we show. Algorithm 3 assigns all the VMs to PM s_1. On the other hand, the optimal offline algorithm also assigns all the VMs to one PM. To prove it, it is enough to show that $\ell(D)^\alpha + b < \ell(A_1)^\alpha + \ell(A_2)^\alpha + 2b$. Using that $\ell(A_1)^\alpha + \ell(A_2)^\alpha > 2\,(\ell(D)/2)^\alpha$ and manipulating, it is enough to prove $\ell(D) < 2\,(b/(2^\alpha - 2))^{1/\alpha}$. This is true for $\ell(D) \leq (b/(2^\alpha - 2))^{1/\alpha}$.

We consider now the case $(b/(2^\alpha - 2))^{1/\alpha} < \ell(D) < 2\,(b/(2^\alpha - 2))^{1/\alpha}$. Within this range, for the optimal algorithm is still better to assign all VMs to one PM, as shown. Then, the competitive ratio ρ is

$$\rho = \frac{\ell(A_1)^\alpha + \ell(A_2)^\alpha + 2b}{\ell(D)^\alpha + b} \leq \frac{\ell(D)^\alpha + 2b}{\ell(D)^\alpha + b} < 2. \tag{4.9}$$

Consider any given step after $\ell(D) \geq 2\,(b/(2^\alpha - 2))^{1/\alpha}$. Within this range, the optimal algorithm may assign the VMs to one or both PMs. If the optimal algorithm assigns to one PM, Inequality 4.9 applies. Otherwise, the competitive ratio ρ is

$$\begin{aligned}\rho &= \frac{\ell(A_1)^\alpha + \ell(A_2)^\alpha + 2b}{2(\ell(D)/2)^\alpha + 2b} \leq 2^{\alpha-1}\frac{\ell(A_1)^\alpha + \ell(A_2)^\alpha}{\ell(D)^\alpha} \\ &= 2^{\alpha-1}\frac{\ell(A_1)^\alpha/\ell(A_2)^\alpha + 1}{(\ell(A_1)/\ell(A_2)+1)^\alpha}.\end{aligned}$$

Then, in order to obtain a ratio at most $x^\alpha/2$, where x will be set later, it is enough to guarantee

$$\begin{aligned}2^{\alpha-1}\frac{\ell(A_1)^\alpha/\ell(A_2)^\alpha + 1}{(\ell(A_1)/\ell(A_2)+1)^\alpha} &\leq \frac{x^\alpha}{2} \\ \frac{(\ell(A_1)/\ell(A_2))^\alpha + 1}{(\ell(A_1)/\ell(A_2)+1)^\alpha} &\leq \left(\frac{x}{2}\right)^\alpha.\end{aligned}$$

Without loss of generality, assume $\ell(A_1) \leq \ell(A_2)$. This implies that $(\ell(A_1)/\ell(A_2))^\alpha \leq \ell(A_1)/\ell(A_2)$. Then, it is enough to have

$$\frac{\ell(A_1)/\ell(A_2) + 1}{(\ell(A_1)/\ell(A_2)+1)^\alpha} \leq \left(\frac{x}{2}\right)^\alpha.$$

Let us now define $\ell(A_1) + \ell = \ell(A_2)$ for some $\ell \geq 0$. Manipulating and replacing, it is enough to show

$$\frac{\ell}{\ell(A_1)} \leq \frac{2-(2/x)^{\alpha/(\alpha-1)}}{(2/x)^{\alpha/(\alpha-1)}-1}. \tag{4.10}$$

If Inequality 4.10 holds the theorem is proved. Otherwise, the following claim is needed.

Claim 1. *If $\ell(D) \geq 2\left(b/(2^\alpha-2)\right)^{1/\alpha}$, then there must exist a VM d_i in D such that $\ell(d_i) \geq |\ell(A_2) - \ell(A_1)|$.*

Proof: If $\ell(A_2) = \ell(A_1)$ the claim follows trivially. Assume that $\ell(A_2) \neq \ell(A_1)$. Consider any given time when $\ell(D) \geq 2\left(b/(2^\alpha-2)\right)^{1/\alpha}$. For the sake of contradiction, assume that for all $d_i \in D$ it is $\ell(d_i) < |\ell(A_2) - \ell(A_1)|$. Let $d_1, d_2, \ldots, d_r$ be the order in which the VMs were revealed to Algorithm 3. And let the respective sets of VMs be called $D_i = \{d_j | j \in [1,i]\}$, that is $D_r = D$. Given that $\ell(D) \geq 2\left(b/(2^\alpha-2)\right)^{1/\alpha} > \left(b/(2^\alpha-2)\right)^{1/\alpha}$, the VM d_r was assigned to the PM with smaller load. Then, either $\ell(d_r) \geq |\ell(A_2) - \ell(A_1)|$ which would be a contradiction, or if $\ell(d_r) < |\ell(A_2) - \ell(A_1)|$ the PM with the smaller load before and after assigning d_r is the same. The argument can be repeated iteratively backwards for each d_{r-1}, d_{r-2}, etc. until, for some $j \in [1,r)$, either it is $\ell(d_j) \geq |\ell(A_2) - \ell(A_1)|$ reaching a contradiction, or the total load is $\ell(D_j) < \left(b/(2^\alpha-2)\right)^{1/\alpha}$. If the latter is the case, we know that for $i \in [1,j]$ every d_i was assigned to s_1. Recall that for $i \in (j,r]$ each d_i was assigned to the same PM. And, given that d_{j+1} is the first VM for which the total load is at least $\left(b/(2^\alpha-2)\right)^{1/\alpha}$, that PM is s_2. But then, we have $\ell(A_2) < \ell(A_1) < \left(b/(2^\alpha-2)\right)^{1/\alpha}$, which is a contradiction with the assumption that $\ell(D) \geq 2\left(b/(2^\alpha-2)\right)^{1/\alpha}$. ■

Using Claim 1 we know that there exists a d_i in the input such that

$$\ell(d_i) \geq \ell > \ell(A_1)\frac{2-(2/x)^{\alpha/(\alpha-1)}}{(2/x)^{\alpha/(\alpha-1)}-1}.$$

From the latter, it can be seen that if $x \geq 2(3/4)^{\frac{\alpha-1}{\alpha}}$, then we have that $\ell > 2\ell(A_1)$. Then, the competitive ratio ρ is

$$\begin{aligned} \rho &= \frac{\ell(A_1)^\alpha + (\ell(A_1)+\ell)^\alpha + 2b}{(2\ell(A_1))^\alpha + \ell^\alpha + 2b} \\ &\leq \frac{\ell(A_1)^\alpha + (\ell(A_1)+\ell)^\alpha}{(2\ell(A_1))^\alpha + \ell^\alpha}. \end{aligned}$$

Using calculus, this ratio is maximized for $\ell = 2\ell(A_1)$ for $\ell \geq 2\ell(A_1)$. Then, we have $\rho \leq (1+3^\alpha)/(2\cdot 2^\alpha)$. Then, in order to obtain a ratio at most $x^\alpha/2$, it is enough to guarantee $(1+3^\alpha)/(2\cdot 2^\alpha) \leq x^\alpha/2$ which yields $x \geq ((1+3^\alpha)/2^\alpha)^{1/\alpha}$.

Given that, for any $\alpha \geq 1$, it holds:

$$2(3/4)^{1-1/\alpha} \geq ((1+3^\alpha)/2^\alpha)^{1/\alpha}.$$

Then, the competitive ratio is $\rho \leq (2(3/4)^{1-1/\alpha})^\alpha/2 = (3/2)^{\alpha-1}$. ■

4.5 Experimental Evaluation

In this section we experimentally evaluate the performance of two of the online algorithms proposed in Section 4.4.2, namely *VMA1* and *VMA2*, and compare them to other state-of-the-art online placement algorithms.

Both *VMA1* and *VMA2* have been extended to handle a bounded number of PMs when necessary. This is achieved by assigning the incoming task to the least loaded machine when no more new PMs are available. Similarly, *VMA2* has also been extended to deal with the bounded capacity case, when required, by using a threshold of $\min\{x^*, C\}$ and checking whether a new load fits into the targeted PM. In case a new load does not fit, it is assigned to the first possible PM with available resources. Although the competitive ratios of *VMA2* are worse than the ones from *VMA1*, it exhibits a better behavior in simulation.

4.5.1 Experimental Setup

The performance of both *VMA1* and *VMA2* is first compared with a lower bound, denoted *LBVMA*, that is obtained as follows. The input VMs are sorted in non-increasing order of their loads. Then, using this order, as many VMs as possible with load at least x^* are assigned to different PMs. Let L be total load of the VMs still unassigned. If there are $\lfloor L/x^* \rfloor$ still unused

PMs they will be used. Otherwise all PMs will be used. Finally, the load L is assigned among all used PMs as if it could be infinitely divided (i.e., as a fluid), using a water-filling algorithm [27].

We evaluate both $(\cdot, m)$-VMA and (C, m)-VMA problems, therefore, in the (C, m)-VMA case a VM can only be assigned to a PM if the latter has sufficient capacity to host it. We test both algorithms *VMA1* and *VMA2* and also compare them with the following algorithms proposed in the literature:

- Random Fit (RF) [74]: It chooses a PM for each VM uniformly at random among the PM. If the chosen PM cannot allocate the load of the VM, the process is repeated, until the VM is assigned to a PM.
- Next Fit (NF) [74]: Starting initially at the first PM, each new VM is assigned to the next PM after the latest PM to which a VM was assigned (in a cyclic fashion) and with sufficient capacity to host it.
- Least Full First (LFF) [74]: Each new VM is assigned to (one of) the least loaded PM(s) in the system with enough capacity to host it.
- Striping (S) [63]: Each new VM is assigned to (one of) the PM(s) with the smallest number of VMs assigned and with enough capacity to host it.
- Watts per Core (WC) [63]: Assigns each new VM to the PM whose power would suffer the smallest increase and with enough capacity to host it.
- First Fit (FF) [74]: Starting initially at the first PM, each new VM is placed in the first PM that can host it.
- Round Robin (RR) [63]: Like FF but, after the first VM is assigned, the search starts from the latest PM in which a VM was allocated.
- Packing (P) [63]: Each new VM is assigned to (one of) the PM(s) with the largest number of VMs assigned provided that the PM can host it.
- Most Full First (MFF) [74]: Each new VM is assigned to the most loaded PM in which it fits.

Observe that, given its nature, First Fit, Round Robin, Packing and Most Loaded First can be only considered for the (C, m)-VMA problem. If PMs had infinite capacity, these algorithms would place all VMs in only one PM. The remaining algorithms are evaluated for both $(\cdot, m)$-VMA and (C, m)-VMA.

The behavior of the aforementioned algorithms is evaluated by inputting the two sets of traces, synthetic and real, shown in Figure 4.1. We call them *Trace A* and *Trace B*, respectively. Trace A is generated by randomly choosing the load of each VM following a power-law distribution with exponential cutoff, which has been chosen so 100% is the maximum task load of a VM.

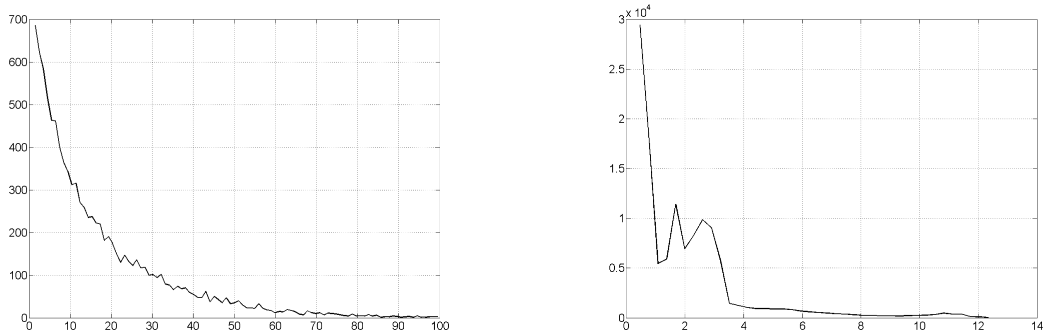

(a) Trace A (synthetic traces)

(b) Trace B (Google traces)

Figure 4.1: VM load distributions used in the evaluations.

We randomly select 10000 integer loads using this distribution. This leads us to the VM load distribution shown in Figure 4.1(a).

Trace B is obtained from public Google traces [57]. We extract all the tasks from these traces, assuming that each task is an independent VM. The VMs (tasks) are sorted by the time at which they join the system. The task load of a VM is the maximum CPU load of the task. The trace then contains 124885 VMs with loads varying between 0.31% and 12.5%. The resulting VM load distribution can be seen in Figure 4.1(b). The load values of the VMs for both distributions are given in percentage in order to scale them up depending on the maximum capacity of a PM in the (C, m)-VMA case or to an appropriate value in the $(\cdot, m)$-VMA case.

Each execution of the algorithms is run with a fixed number of PMs. This number of PMs increases from 1 to the number of VMs in the trace being used. This allows us to see how the power consumption and how the algorithms behavior evolve when the number of available PMs in the system varies. Finally, in order to evaluate both $(\cdot, m)$-VMA and (C, m)-VMA, we *emulate* different PMs by determining their α, b, μ and x^* parameters as well as the PM maximum capacity or the maximum task load when it corresponds. Then we run the proposed algorithms for each one of these emulated PMs and compare the influence of the different values for these parameters on the final results.

4.5.2 Experimental Results for $(\cdot, m)$-VMA

We first evaluate $(\cdot, m)$-VMA. The first step is to define the set of PMs that we are going to use to evaluate it. To do so, we fix the values of α, b and x^* and compute μ depending on the previous parameters. In particular, we used $\alpha = \{1.5, 2, 2.5, 3\}$ and $x^* = \{10, 30, 50, 75, 100, 130, 150, 300, 500, 750, 1000\}$ (given in (Giga)*Cycles per Second* (GCPS) following the conclusions from Chapter 3). The values of b are determined by interpolation of the baseline costs of `Nemesis`, `Survivor` and `Erdos`, whose values for b (~ 85 W, 67 W and 215 W) are known from the experiments performed in Chapter 3. As we mentioned, the values

Table 4.3: Simulation parameters for a set of machines for the $(\cdot, m)$-VMA case.

$(\cdot,\cdot)$-VMA case					
b	x^* [GCPS]	$\alpha = 1.5$	$\alpha = 2$	$\alpha = 2.5$	$\alpha = 3$
73.05	10	1.46E-13	7.31E-19	4.87E-24	3.65E-29
92.15	30	3.55E-14	1.02E-19	3.94E-25	1.71E-30
111.25	50	1.99E-14	4.45E-20	1.33E-25	4.45E-31
135.125	75	1.32E-14	2.40E-20	5.85E-26	1.60E-31
159	100	1.01E-14	1.59E-20	3.35E-26	7.95E-32
187.65	130	8.01E-15	1.11E-20	2.05E-26	4.27E-32
206.75	150	7.12E-15	9.19E-21	1.58E-26	3.06E-32
350	300	4.26E-15	3.89E-21	4.73E-27	6.48E-33
541	500	3.06E-15	2.16E-21	2.04E-27	2.16E-33
779.75	750	2.40E-15	1.39E-21	1.07E-27	9.24E-34
1018.5	1000	2.04E-15	1.02E-21	6.79E-28	5.09E-34

of b, α and x^* determine the value of μ. These combination of parameters results in 44 different instances of PMs which are shown in Table 4.3.

Additionally, taking advantage of the fact that the task loads from Trace A and B are given in percentage, and in order to study the importance of the x^* to task load ratio, we define the maximum VM load λ as the maximum load that a VM arriving to the system can have. Therefore, the load of the VMs arriving to the system will be the product of the task load (in percentage) and λ. λ will take the following values: 10, 30 and 100 GCPS.

We study three different scenarios. In the first one we study the effect of α for different values of λ and x^* when using *VMA1*, *VMA2*, and the lower bound *LBVMA*. Second and third scenarios are devoted to compare our proposed algorithms with the state-of-the-art algorithms, always keeping *LBVMA* as a reference. In the second scenario we study the relevance of λ while keeping α and x^* constant. Finally, in the third one, we study the effect of having different values of x^* while λ and α remain unaltered.

Scenario 1 compares the power consumed by partitions obtained with *VMA1* or *VMA2* and for Trace A and Trace B. We compare these results to the ones achieved by *LBVMA*, that lower bounds the optimal power consumption. The results obtained are presented as graphs in which the power consumed is represented as a function of the number of PMs used.

Figure 4.2 and Figure 4.3 show the results for Trace A and Trace B, for 2 different values of x^*, 30 and 300 GCPS, and 2 different values of λ, 10 and 100 GCPS. We can clearly see how the power consumption is smaller for larger values of α once the optimal number of used PMs is reached, mainly conditioned by how μ decreases as α increases (See Table 4.3. Also, as it can be observed, there is no qualitative difference in the solutions when α varies for a given configuration (Similar results are obtained with other values of x^* and maximum task load).

Regarding the performance of the algorithms, we can see how the power consumed by the partitions found with *VMA2* is lower, in all cases, than the ones obtained by *VMA1* and is always

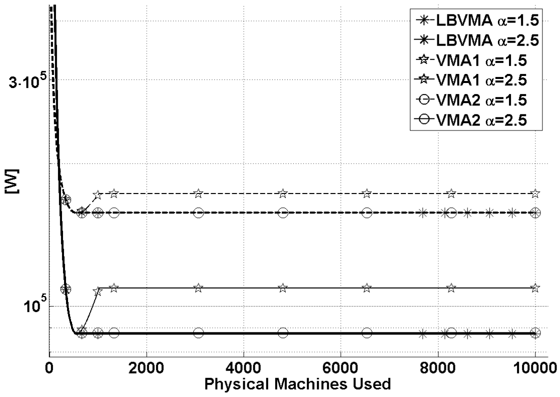

(a) 10 GCPS, $x^* = 30$ GCPS

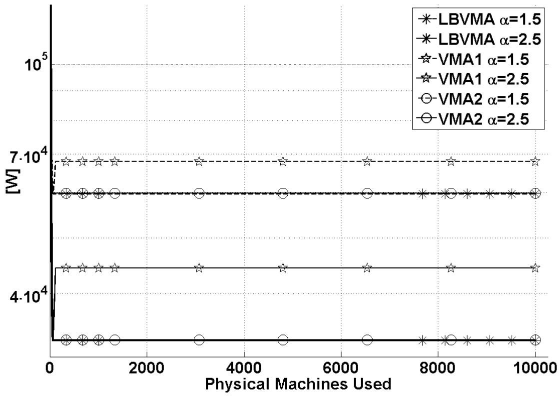

(b) 10 GCPS, $x^* = 300$ GCPS

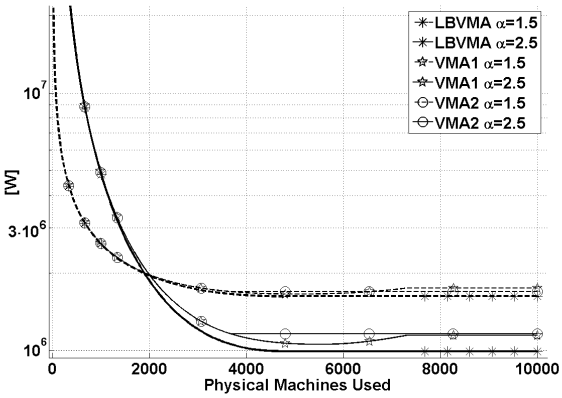

(c) 100 GCPS, $x^* = 30$ GCPS

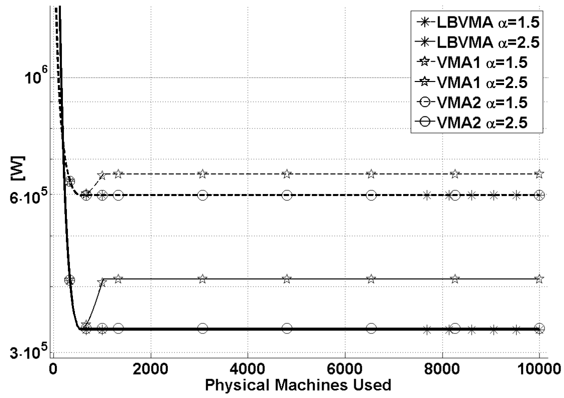

(d) 100 GCPS, $x^* = 300$ GCPS

Figure 4.2: $(\cdot, m)$-VMA: Comparing the power consumed by *VMA1* and *VMA2* with the lower bound $LBVMA$ for $x^* = \{30, 300\}$ GCPS, $\alpha = \{1.5, 2.5\}$ and $\lambda = \{10, 100\}$ GCPS for Trace A (Synthetic traces).

closer to the lower bound obtained by *LBVMA*. This shows that the performance of *VMA2* is close to the optimal for $(\cdot, m)$-VMA. We can see how, in general, *VMA1* is able to match the results of *VMA2* when the number of PMs is relatively low. However, due to the threshold imposed on the load of the PMs for each algorithm, *VMA2* is able to pack the load in less PMs. We only find an exception in Figure 4.2(c), when $x^*/\lambda < 1/3$ and we are using synthetic traces. In this case *VMA2* exhibits a behavior relatively similar to *VMA1*, not being able to hold to its best power consumption and reducing the quality of the solution when the number of PMs increases. However, this flaw is not replicated when using Trace B, due to the smaller amount of big loads in comparison with Trace B.

Scenario 2 compares the performance of *LBVMA*, *VMA1* and *VMA2* with the other assignment algorithms proposed in the literature. Here, the values of x^* and α are fixed to 30 GCPS and 2, respectively, while the value of λ varies. In particular we use $\lambda = \{10, 30, 100\}$ GCPS. Figures

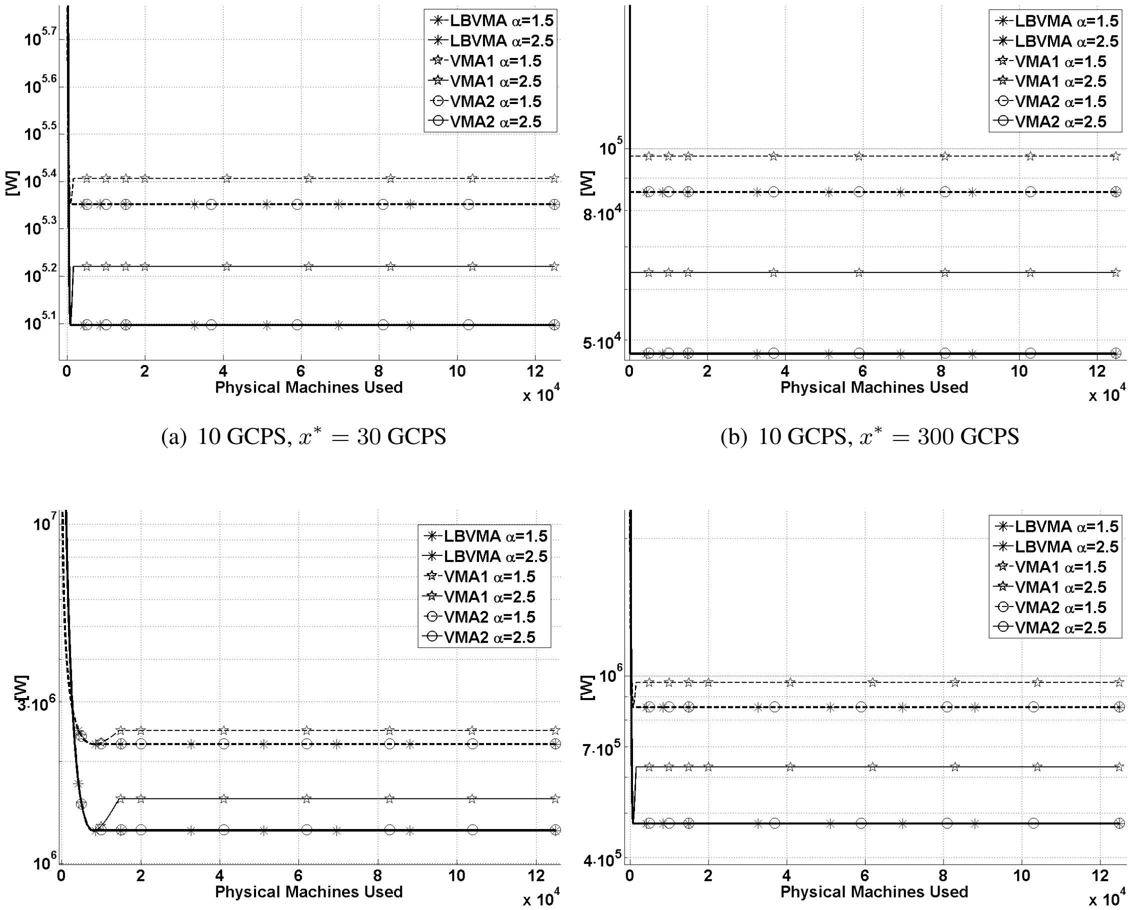

(a) 10 GCPS, $x^* = 30$ GCPS

(b) 10 GCPS, $x^* = 300$ GCPS

(c) 100 GCPS, $x^* = 30$ GCPS

(d) 100 GCPS, $x^* = 300$ GCPS

Figure 4.3: $(\cdot, m)$-VMA: Comparing the power consumed by *VMA1* and *VMA2* with the lower bound $LBVMA$ for $x^* = \{30, 300\}$ GCPS, $\alpha = \{1.5, 2.5\}$ and $\lambda = \{10, 100\}$ GCPS for Trace B (Google traces).

4.4 and 4.5 present the results for Trace A and Trace B[3].

We can easily see, in general, 3 different trends in Figures 4.4 and 4.5. The first trend would include *LBVMA*, *VMA1* and *VMA2*; then we have a second one including *Striping*, *RF*, *NF* and *LLF*; and, finally, in some sort of no-man's land, we have *WC*. These trends have their origin in power awareness. While *LBVMA*, *VMA1*, *VMA2* and *WC* are power aware, the rest are not. According to Figures 4.4 and 4.5, power aware algorithms outperform the non power aware ones.

It is interesting to see how *WC* reduces its power consumption (for Trace A) as the ratio x^*/λ decreases and even performs better than *VMA1* for $\lambda = 100$ GCPS. This does not happen, though, for Trace B due to the smaller average task load, that gives advantage to *VMA1* and *VMA2*. With Trace B, due to its nature, *WC* obtains partitions that use less PMs and hence, because of the

[3]For the sake of clarity, we do not show the power consumption resulting of using only one (or a few) machines and center the figure into more relevant cases

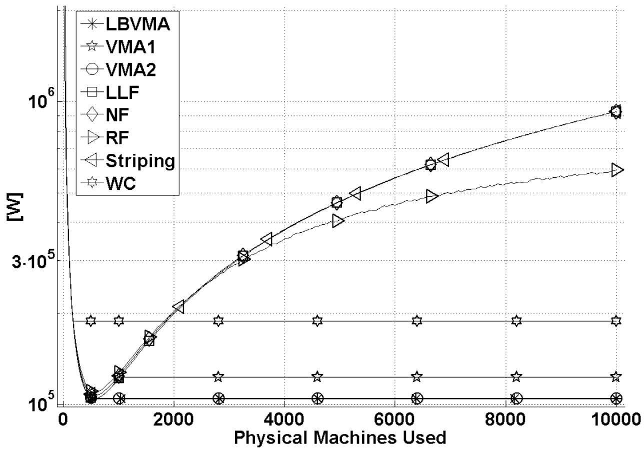

(a) $\lambda = 10$ GCPS

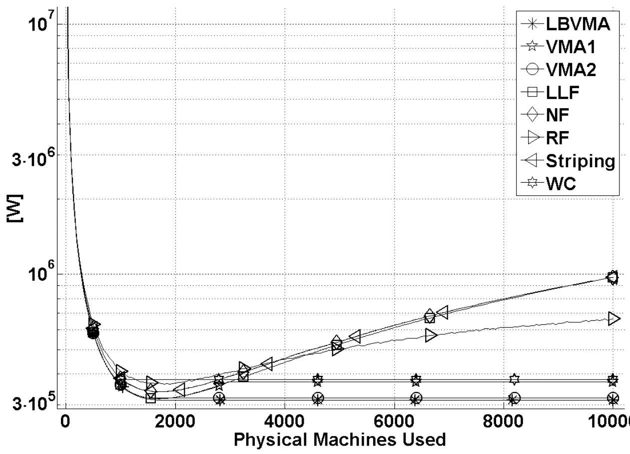

(b) $\lambda = 30$ GCPS

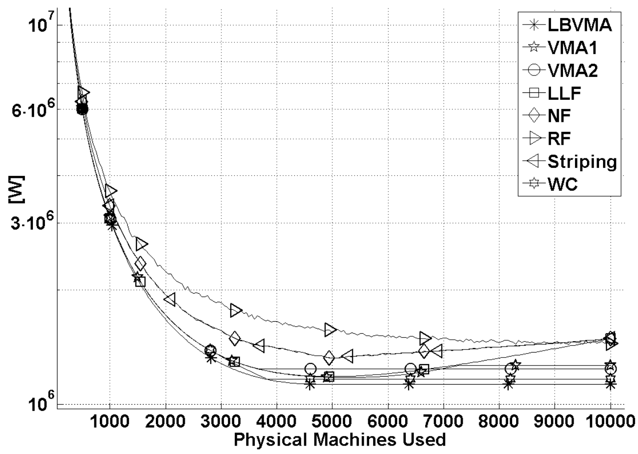

(c) $\lambda = 100$ GCPS

Figure 4.4: $(\cdot, m)$-VMA: Comparing the power consumed by the different assignment algorithms for $x^* = 30$ GCPS, $\alpha = 2$ and $\lambda = \{10, 30, 100\}$ GCPS for Trace A (Synthetic traces).

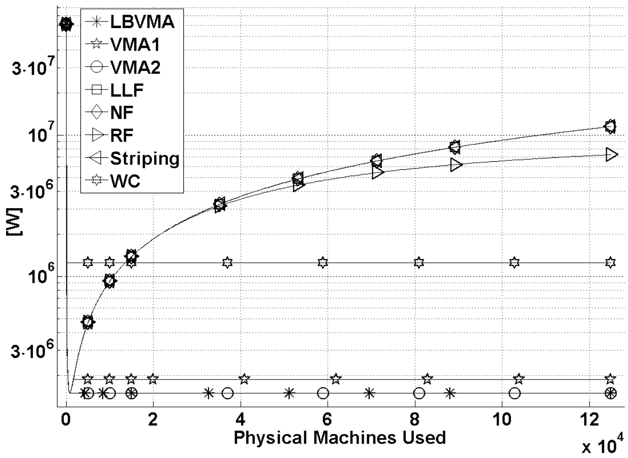

(a) $\lambda = 10$ GCPS

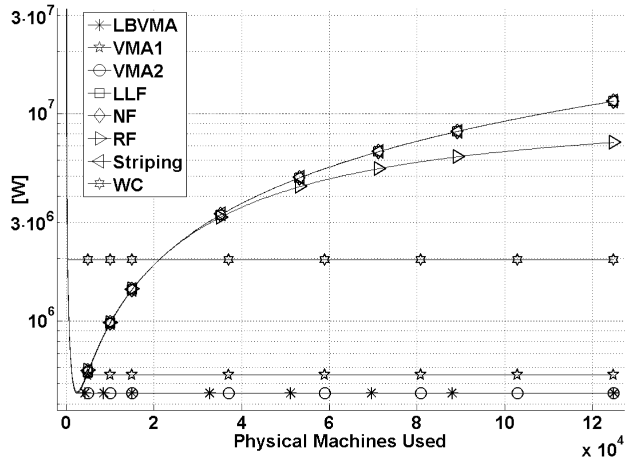

(b) $\lambda = 30$ GCPS

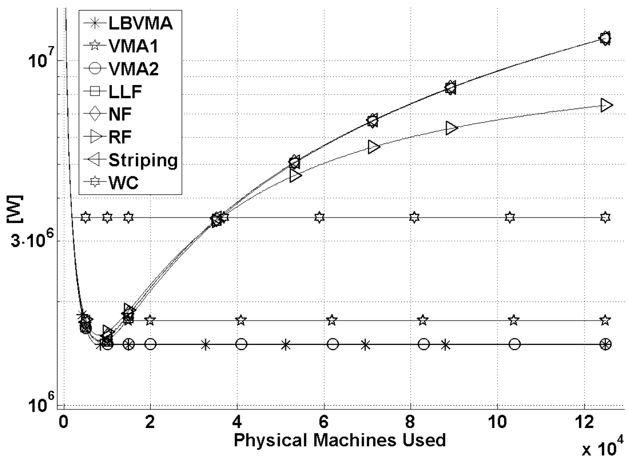

(c) $\lambda = 100$ GCPS

Figure 4.5: $(\cdot, m)$-VMA: Comparing the power consumed by the different assignment algorithms for $x^* = 30$ GCPS, $\alpha = 2$ and $\lambda = \{10, 30, 100\}$ GCPS for Trace B (Google traces).

superlinear dependence of the power consumption on the load, have higher power consumptions.

Similarly, we can observe that the results in Figure 4.4(c) are tighter. This is a consequence, again, of the low value of x^*/λ, resulting in many PMs not allocating more than 1 or 2 VMs. In fact, steady state for *VMA1*, *VMA2*, *WC* and even *LBVMA* is reached between the 4000 and the 5000 PMs while in the previous cases was reached before the 2000 PMs. Again, this behavior is not replicated with Trace B due to its lower average task load.

Note also that the non power aware algorithms pay a higher power bill due to the use of a larger amount of PMs (in general) with a smaller amount of load per PM, resulting in a very inefficient usage of the available resources. This behavior is consistent for both Traces A and B.

Finally, observe how the larger the ratio x^*/λ, the larger the gap between our proposed algorithms, *VMA1* and *VMA2*, and the other ones. This would be the case when we have tasks that consume a very small amount of CPU in the system.

Let us now analyze the results of the last scenario. Here we keep $\lambda = 30$ GCPS and $\alpha = 2$ constant while we vary the value of x^*. These results are shown in Figure 4.6 for Trace A and in Figure 4.7 for Trace B.

The results are similar for both traces. We can see how for the smallest value, $x^* = 10$ GCPS ($x^*/\lambda = 1/3$) all algorithms achieve a similar result. As the ratio x^*/λ increases, the results obtained by *VMA1* and *VMA2* become better than the ones achieved by *WC*, *LLF*, *NF*, *Striping*, and *RF*, that increase with x^* and, therefore, lead to a larger power consumption. These results are in line with the results from Scenario 1 and are is motivated by the fact that the state-of-the-art algorithms tend to use a large amount of PMs keeping its average load low and, hence, paying a high price because of the b parameter. This, however, is not the case of *WC*, which, on the other hand, obtains a more packed partition, loading PMs beyond x^* and paying an extra cost due to the superlinearity of the power consumption with respect to the load.

4.5.3 Experimental Results for (C, m)-VMA

As we did with $(\cdot, m)$-VMA in Subsection 4.5.2, we start by defining the set of PMs we are going to work with. While for $(\cdot, m)$-VMA we assumed that PMs had infinite capacity, in (C, m)-VMA the PMs capacity is bounded. We denote the capacity as C. We define 3 sets of instances that we name after 3 real PMs from our laboratory, `Nemesis`, `Euler`, and `Erdos`. We use, as a reference, their maximum capacity, 11.2, 41.6, and 153.6 GCPS; and approximated idle cost b, 80, 100, and 200 Watts.

Jointly with C and b, we need a value of α and x^* to compute each value of μ. We will use $\alpha = \{1.5, 2, 2.5, 3\}$, and $x^* = \{0.5, 0.65, 0.75, 0.9, 1, 1.1\} \cdot C$. We can now compute the value of μ for each combination of these 4 parameters fully defining, then, our set of PMs. The combination of values for each one of these PM instances can be found in Tables 4.5, 4.6 and 4.4. We base the performance analysis of our proposed algorithms and the other state of the art assignment algorithms on these sets of PM instances.

Like for the $(\cdot, m)$-VMA case, we also consider three different scenarios. In the first one we

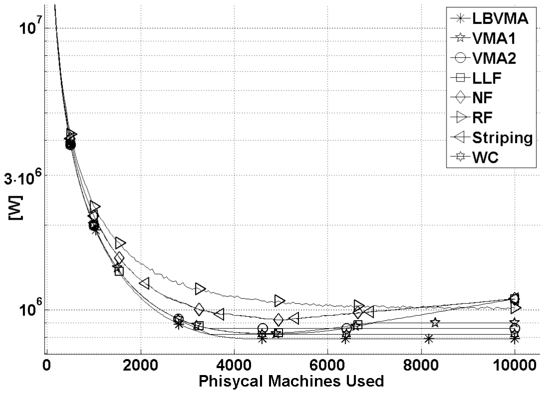

(a) $x^* = 10$ GCPS

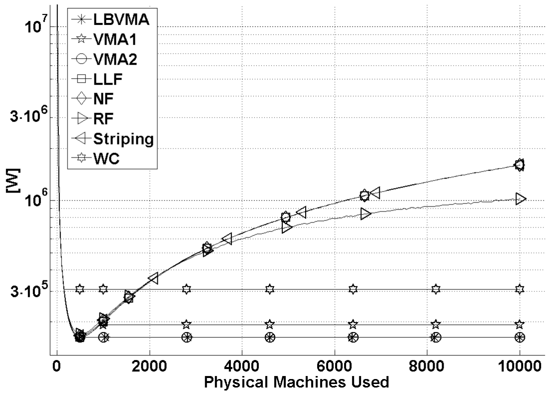

(b) $x^* = 100$ GCPS

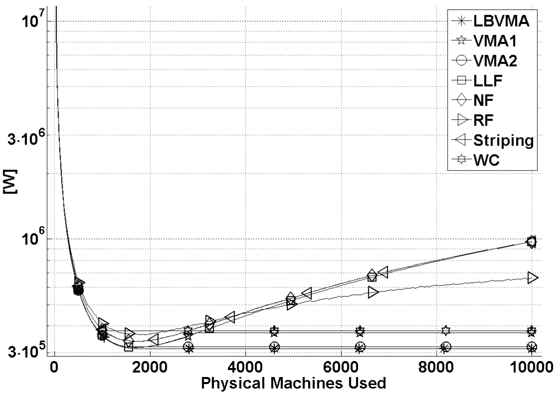

(c) $x^* = 30$ GCPS

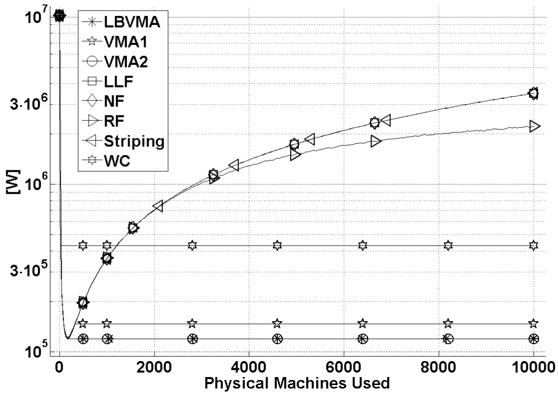

(d) $x^* = 300$ GCPS

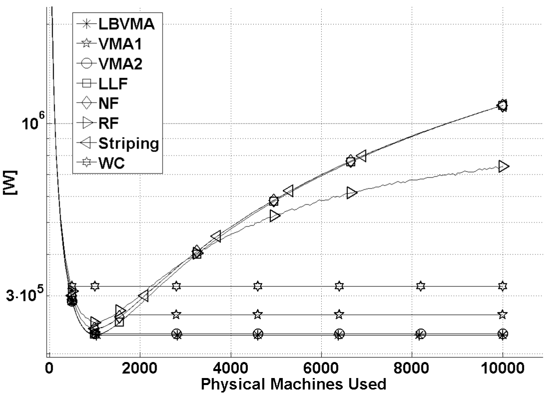

(e) $x^* = 50$ GCPS

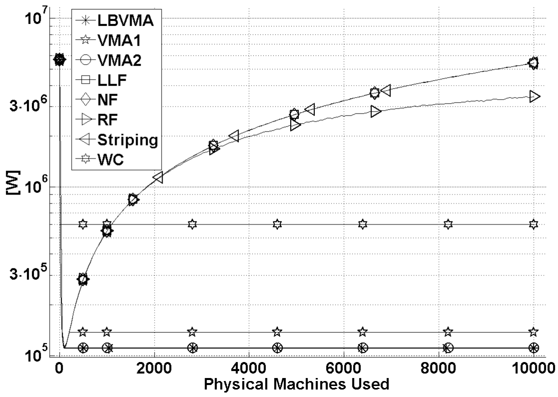

(f) $x^* = 500$ GCPS

Figure 4.6: $(\cdot, m)$-VMA: Comparing the power consumed by the different assignment algorithms for $\lambda = 30$ GCPS, $\alpha = 2$ and increasing values of x^* for Trace A (Synthetic traces).

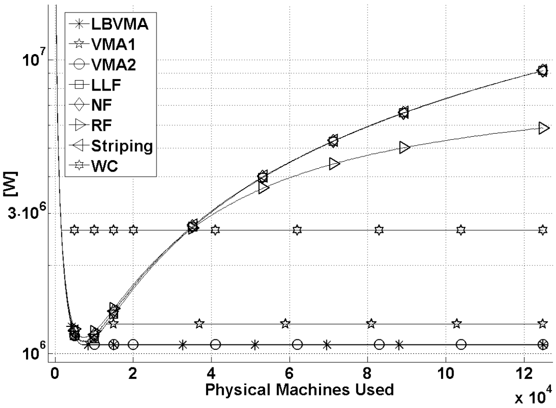

(a) $x^* = 10$ GCPS

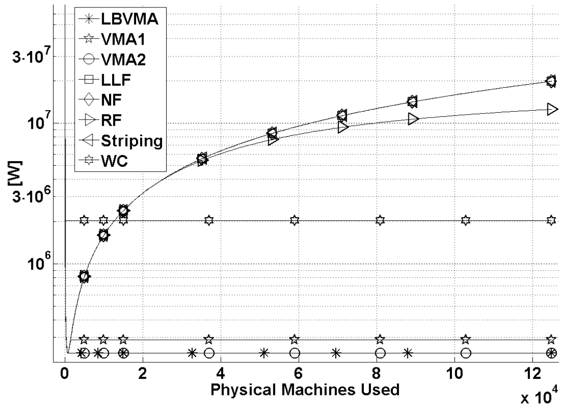

(b) $x^* = 100$ GCPS

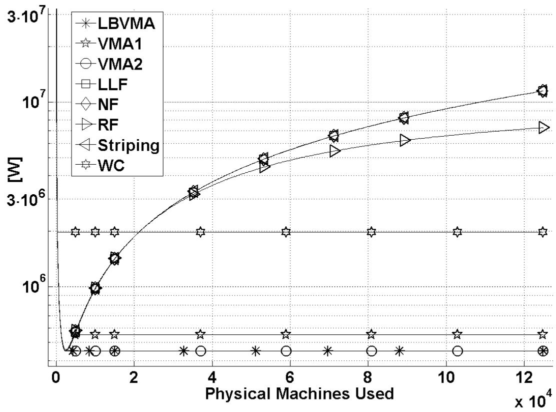

(c) $x^* = 30$ GCPS

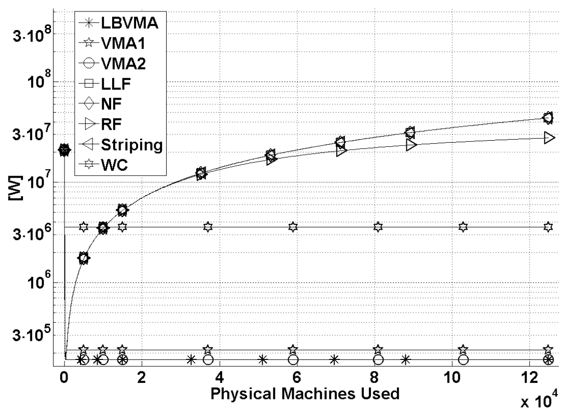

(d) $x^* = 300$ GCPS

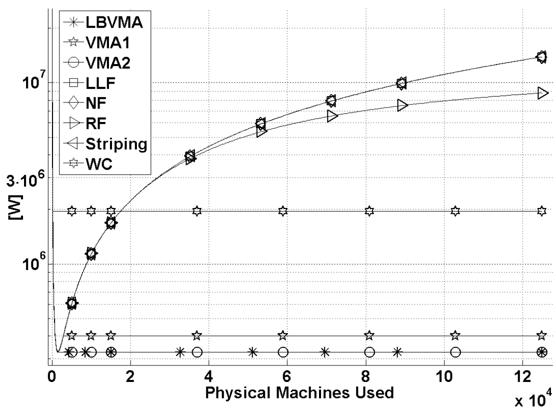

(e) $x^* = 50$ GCPS

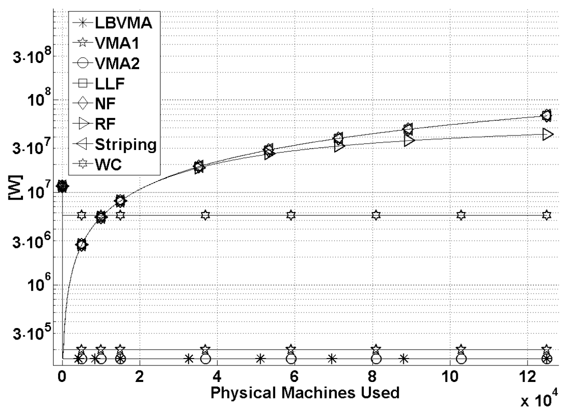

(f) $x^* = 500$ GCPS

Figure 4.7: $(\cdot, m)$-VMA: Comparing the power consumed by the different assignment algorithms for $\lambda = 30$ GCPS, $\alpha = 2$ and increasing values of x^* for Trace B (Google traces).

Table 4.4: Simulation parameters for a set of machines with b and C similar to `Erdos`.

`Erdos`-like Family		Max. Capacity $C = 153.6$ GCPS			
		$b = 200$ W			
x^* [GCPS]	$x^*[\%C]$	$\alpha = 1.5$	$\alpha = 2$	$\alpha = 2.5$	$\alpha = 3$
76.8	50	1.88E-14	3.39E-20	8.16E-26	2.21E-31
99.84	65	1.27E-14	2.01E-20	4.23E-26	1.00E-31
115.2	75	1.02E-14	1.51E-20	2.96E-26	6.54E-32
138.24	90	7.78E-15	1.05E-20	1.88E-26	3.79E-32
153.6	100	6.64E-15	8.48E-21	1.44E-26	2.76E-32
168.96	110	5.76E-15	7.01E-21	1.14E-26	2.07E-32

Table 4.5: Simulation parameters for a set of machines with b and C similar to `Nemesis`.

`Nemesis`-like Family		Max. Capacity $C = 11.2$ GCPS			
		$b = 80$ W			
x^* [GCPS]	$x^*[\%C]$	$\alpha = 1.5$	$\alpha = 2$	$\alpha = 2.5$	$\alpha = 3$
5.6	50	3.82E-13	2.55E-18	2.27E-23	2.28E-28
7.28	65	2.58E-13	1.51E-18	1.18E-23	1.04E-28
8.4	75	2.08E-13	1.13E-18	8.25E-24	6.75E-29
10.08	90	1.58E-13	7.87E-19	5.23E-24	3.91E-29
11.2	100	1.35E-13	6.38E-19	4.02E-24	2.85E-29
12.32	110	1.17E-13	5.27E-19	3.17E-24	2.14E-29

Table 4.6: Simulation parameters for a set of machines with b and C similar to `Euler`.

`Euler`-like Family		Max. Capacity $C = 41.6$ GCPS			
$b = 80$ W					
x^* [GCPS]	$x^*[\%C]$	$\alpha = 1.5$	$\alpha = 2$	$\alpha = 2.5$	$\alpha = 3$
20.8	50	5.33E-14	1.85E-19	8.55E-25	4.44E-30
27.04	65	3.60E-14	1.09E-19	4.44E-25	2.02E-30
31.2	75	2.90E-14	8.22E-20	3.10E-25	1.32E-30
37.44	90	2.21E-14	5.71E-20	1.97E-25	7.62E-31
41.6	100	1.89E-14	4.62E-20	1.51E-25	5.56E-31
45.76	110	1.63E-14	3.82E-20	1.19E-25	4.17E-31
$b = 100$ W					
20.8	50	6.67E-14	2.31E-19	1.07E-24	5.56E-30
27.04	65	4.50E-14	1.37E-19	5.54E-25	2.53E-30
31.2	75	3.63E-14	1.03E-19	3.88E-25	1.65E-30
37.44	90	2.76E-14	7.13E-20	2.46E-25	9.53E-31
41.6	100	2.36E-14	5.78E-20	1.89E-25	6.95E-31
45.76	110	2.04E-14	4.78E-20	1.49E-25	5.22E-31
$b = 120$ W					
20.8	50	8.00E-14	2.77E-19	1.28E-24	6.67E-30
27.04	65	5.40E-14	1.64E-19	6.65E-25	3.03E-30
31.2	75	4.35E-14	1.23E-19	4.65E-25	1.98E-30
37.44	90	3.31E-14	8.56E-20	2.95E-25	1.14E-30
41.6	100	2.83E-14	6.93E-20	2.27E-25	8.33E-31
45.76	110	2.45E-14	5.73E-20	1.79E-25	6.26E-31

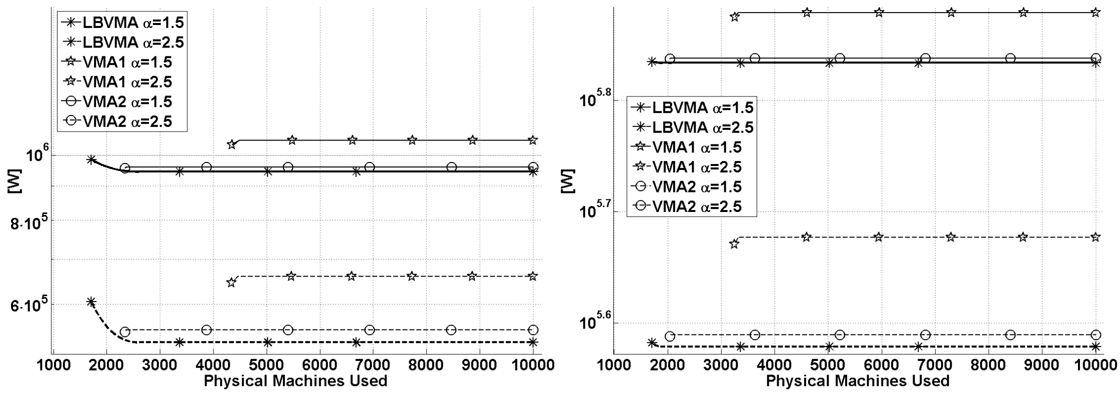

(a) $C = 153.6$ GCPS, $x^* = 0.65 \cdot C$ (b) $C = 153.6$ GCPS, $x^* = 0.9 \cdot C$

Figure 4.8: (C, m)-VMA: Comparing the power consumed by *VMA1* and *VMA2* with the lower bound *LBVMA* for $x^* = \{0.65, 0.9\} \cdot C$, $C = 153.6$ GCPS, $\alpha = \{1.5, 2.5\}$ for Trace A (synthetic traces).

compare again *VMA1* and *VMA2* with *LBVMA* for different values of α to evaluate its influence. In the second scenario, we evaluate the performance of *VMA1*, *VMA2*, and *LBVMA* jointly with *FF*, *MLF*, *Packing*, *RR*, *RF*, *LFF*, *NF*, *Striping* and *WC* when $alpha$, b and C are fixed and x^* varies. Finally, in the third scenario, we evaluate the influence of b on the performance of the different algorithms when α and C are fixed and different values of x^* and b are considered. Note that all experiments are run for both Trace A and Trace B. Similarly, observe that results are presented, again, as graphs in which the power consumed is represented as a function of the number of PMs used but, this time, these results do not start from 1 PM. Basically, each one of the results of the different algorithms starts from the number of PMs for which it obtained a valid solution and PMs where PMs do not have to be loaded beyond their capacity C.

The results for the first scenario, shown in Figures 4.8 and 4.9, throw very similar results to the ones obtained for $(\cdot, m)$-VMA. As for $(\cdot, m)$-VMA, *VMA2* performance is very close to *LBVMA* when it does not match it. Similarly, *VMA1* is again worse than *VMA2* due to the fact that it tends to pack less VMs per PM than *VMA1*.

The results for the second scenario are presented in Figure 4.10 and Figure 4.11 for Traces A and B, respectively. In them we compare the performance of *VMA1*, *VMA2*, *LBVMA*, *FF*, *MLF*, *Packing*, *RR*, *RF*, *LFF*, *NF*, *Striping* and *WC* for PMs such as the ones defined in Table 4.5. We can easily observe, independently of the trace used, that *RP*, *LFF*, *NF* and *Striping* consistently obtain partitions which result in a higher cost in Watts than the other algorithms. The nature of this set of algorithms implies using a large number of PMs with a small amount of load per PM, resulting, always, in a higher power consumption. For this reason, and for the sake of clarity when plotting and comparing the other algorithms, we only show *RP*, *LFF*, *NF* and *Striping* in subfigures 4.10(a), 4.10(b), 4.11(a) and 4.11(b) as an example. The rest of the subfigures from Figure 4.10 and Figure 4.11 zoom in the results of the other algorithms.

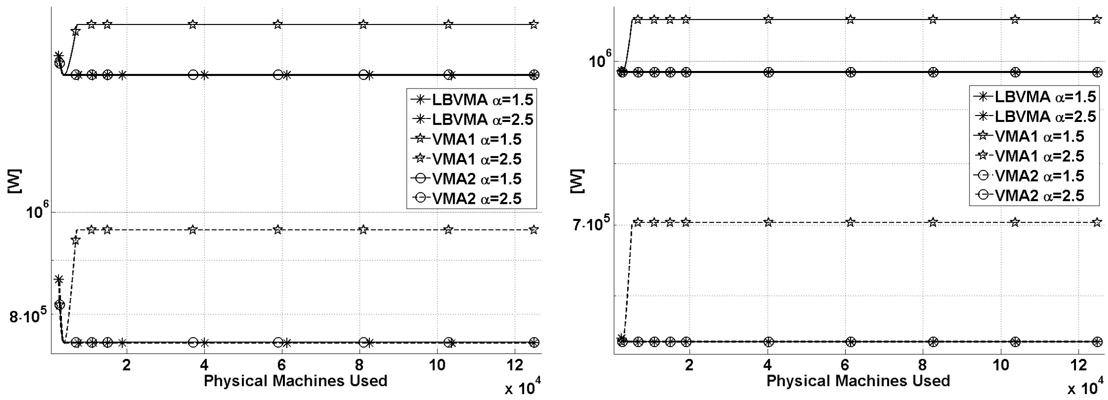

(a) $C = 153.6$ GCPS, $x^* = 0.65 \cdot C$ (b) $C = 153.6$ GCPS, $x^* = 0.9 \cdot C$

Figure 4.9: (C, m)-VMA: Comparing the power consumed by *VMA1* and *VMA2* with the lower bound *LBVMA* for $x^* = \{0.65, 0.9\} \cdot C$, $C = 153.6$ GCPS, $\alpha = \{1.5, 2.5\}$ for Trace B (Google traces).

Oppositely to the $(\cdot, m)$-VMA case, *WC* exhibits a better performance than *VMA1*, that also loses in the comparison with the group formed by *FF*, *MLF*, *Packing* and *RR* in every case except for Trace A and $x^* = 0.5 \cdot C$, shown in Figure 4.10(f)[4]. This worse performance of *VMA1* is a consequence of the ability of the other algorithms to obtain more packed solutions whose load is, additionally, closer to x^* than *VMA1*'s solution. On the other hand, *VMA2* results are similar or slightly worse than *FF*, *MLF*, *Packing* and *RR* for Trace A when when $x^* > 0.75C$ (Figures 4.10(a), 4.10(b) and 4.10(c)). However, *VMA2* still outperforms the other algorithms when $x^* \leq 0.75C$ (Figures 4.10(d),4.10(e) and 4.10(f)) or when the average VM task load is smaller, i.e., all the cases of Trace B. Once again, *VMA2* stays close to the optimal and, in general, performs better than the other algorithms.

We finally evaluate the influence of b when fixing other parameters. These results are shown in Figures 4.12 and 4.13 for Traces A and B, respectively. We compare, again, the performance of *VMA1*, *VMA2*, *LBVMA*, *FF*, *MLF*, *Packing*, *RR*, *RF*, *LFF*, *NF*, *Striping* and *WC* when $\alpha = 2$, $C = 41.6$ GCPS, $x^* = \{0.65, 0.9\} \cdot C$ and $b = \{80, 100, 120\}$ W. Similarly to what happened in Scenario 2 (Figures 4.10 and 4.11), the results obtained by *RP*, *LFF*, *NF* and *Striping* are consistently worse than the ones obtained by the other algorithms. Therefore, as in the previous case, we only show them in the two first figures for each trace, namely Figures 4.12(a), 4.12(b), 4.13(a) and 4.13(b). The rest of the subfigures from Figures 4.12 and Figure 4.13 zoom in the results of the other algorithms.

Observing these results we can barely appreciate any difference in behavior/shape between the different algorithms when we x^* remains constant. Of course there are differences in the power consumed for each value of b, but this is mainly because of the contribution of b. The

[4]This behaviour is not replicated for Trace B because of the smaller average load of Trace B VMs.

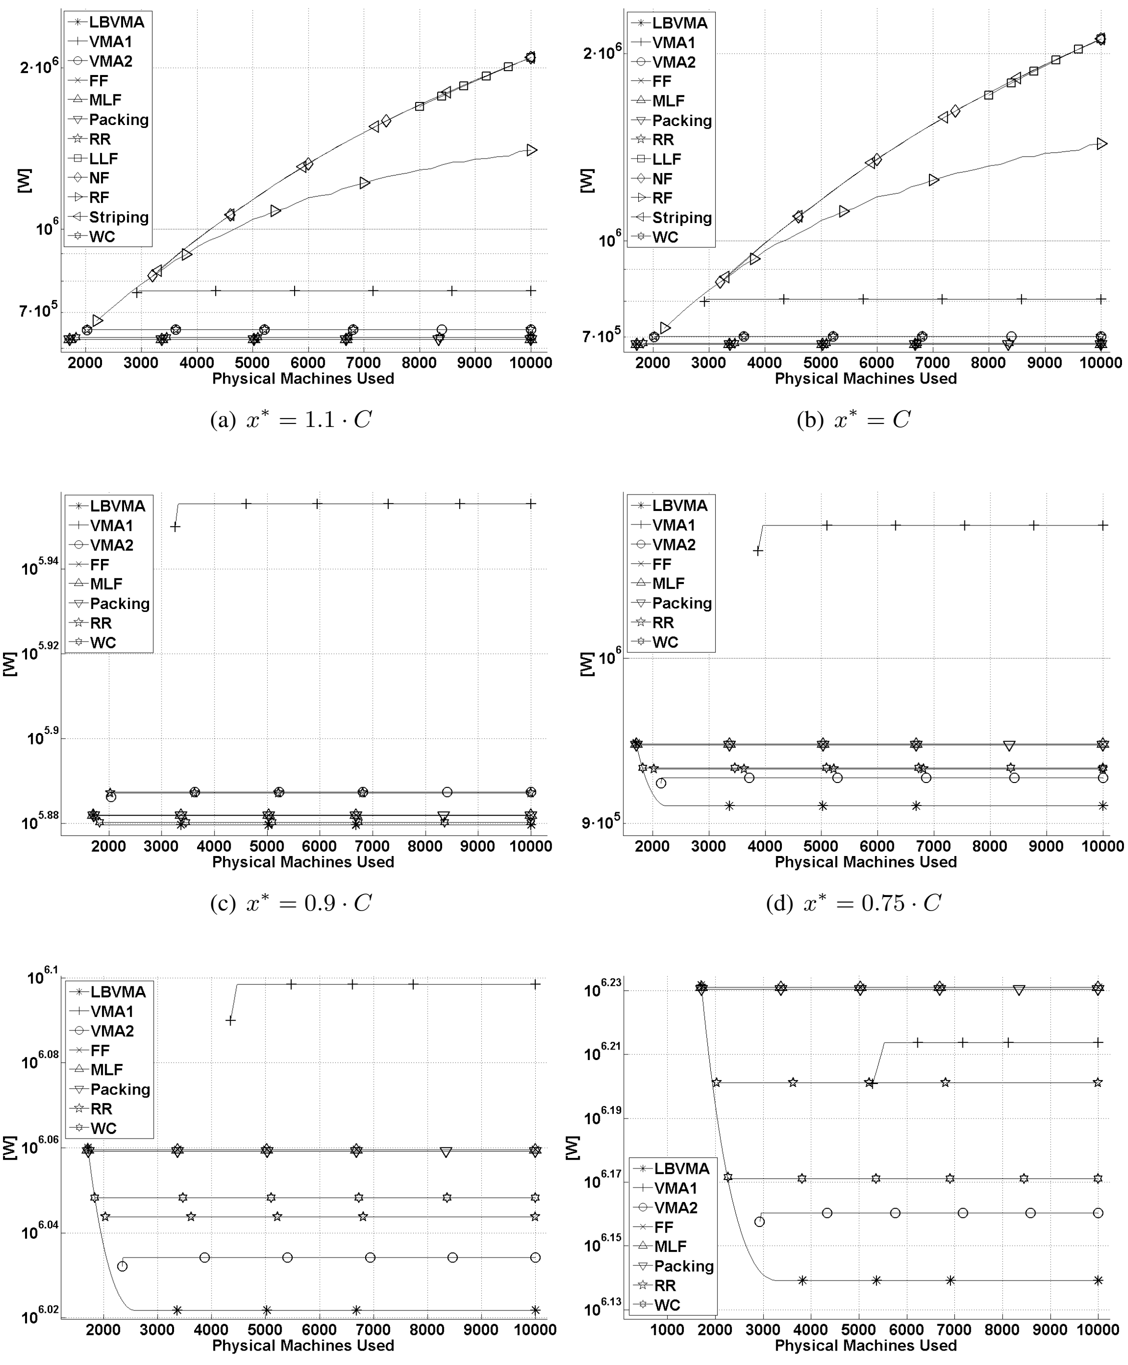

(a) $x^* = 1.1 \cdot C$

(b) $x^* = C$

(c) $x^* = 0.9 \cdot C$

(d) $x^* = 0.75 \cdot C$

(e) $x^* = 0.65 \cdot C$

(f) $x^* = 0.5 \cdot C$

Figure 4.10: (C, m)-VMA: Comparing the power consumed by the different assignment algorithms for $C = 11.2$ GCPS, $\alpha = 2$ and different values of x^* for Trace A (Synthetic traces).

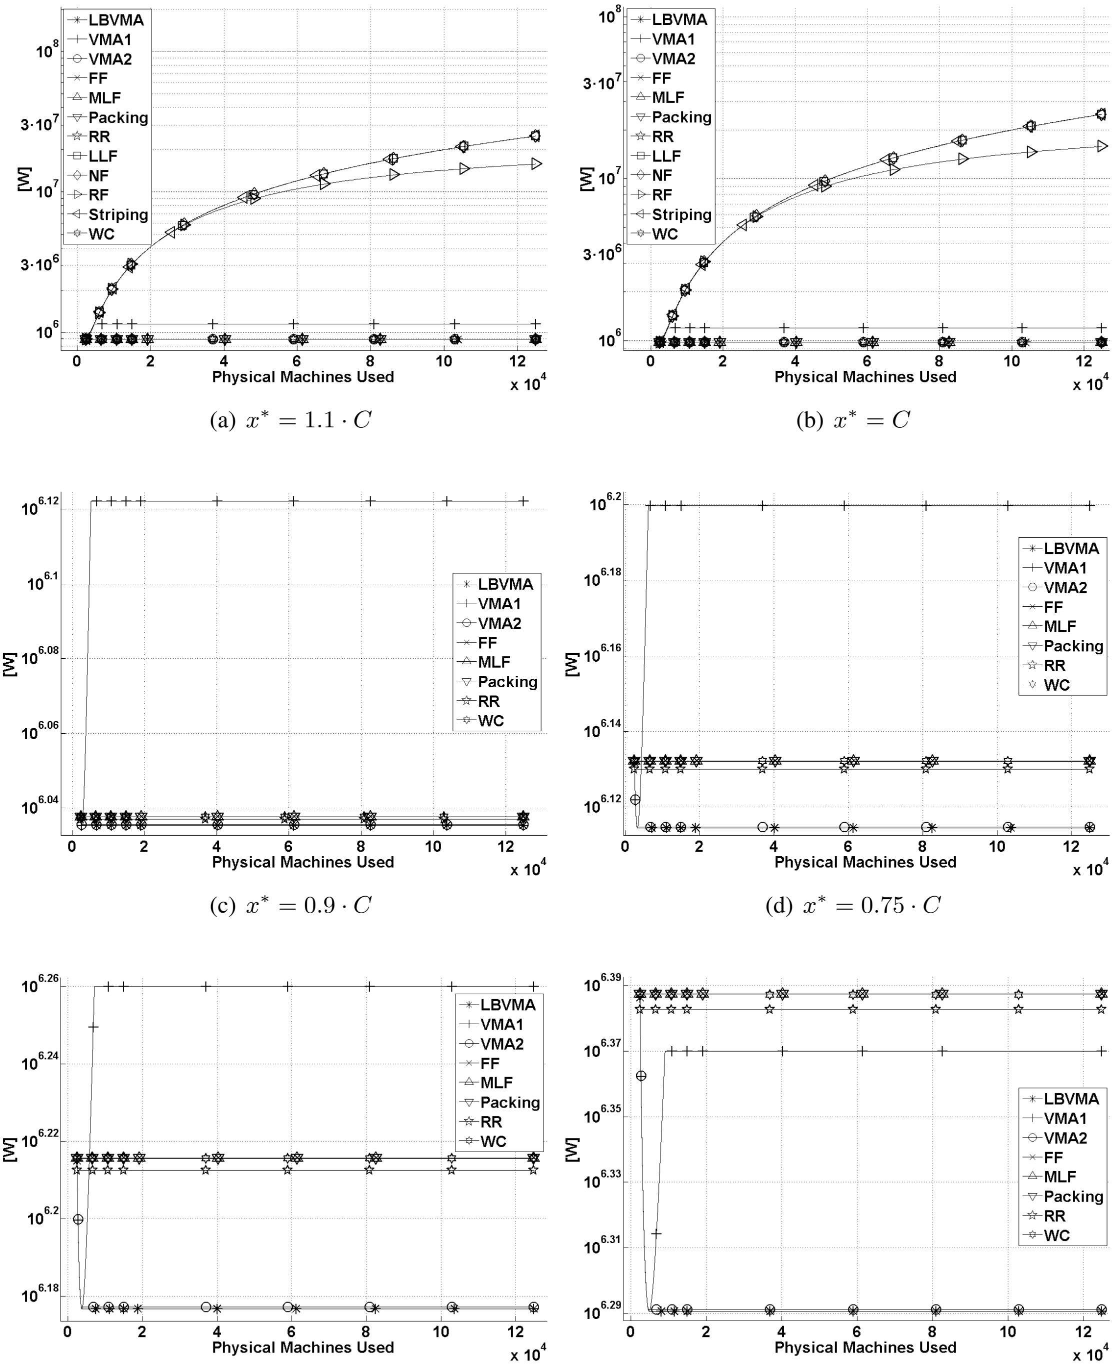

(a) $x^* = 1.1 \cdot C$

(b) $x^* = C$

(c) $x^* = 0.9 \cdot C$

(d) $x^* = 0.75 \cdot C$

(e) $x^* = 0.65 \cdot C$

(f) $x^* = 0.5 \cdot C$

Figure 4.11: (C, m)-VMA: Comparing the power consumed by the different assignment algorithms for $C = 11.2$ GCPS, $\alpha = 2$ and different values of x^* for Trace B (Google traces).

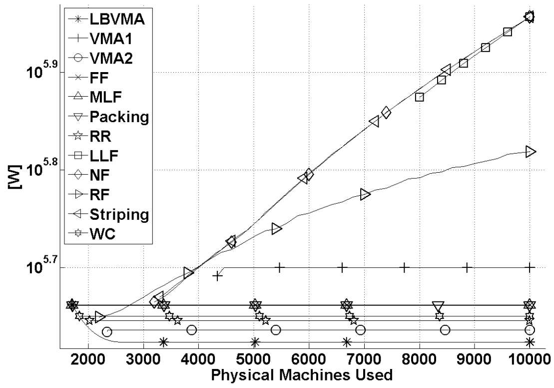

(a) $b = 80$ W, $x^* = 0.65 \cdot C$

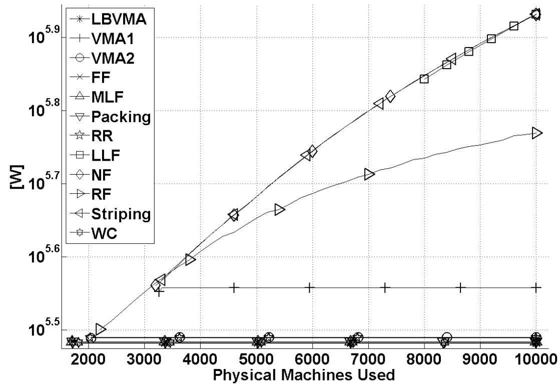

(b) $b = 80$ W, $x^* = 0.9 \cdot C$

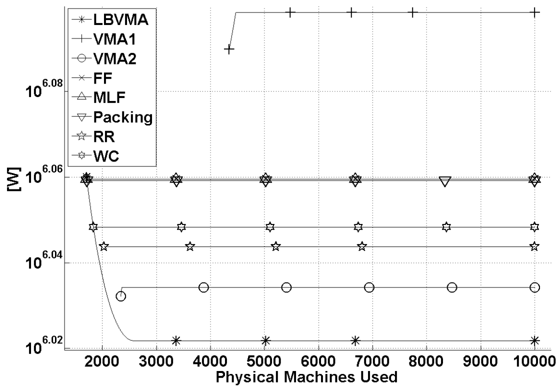

(c) $b = 100$ W, $x^* = 0.65 \cdot C$

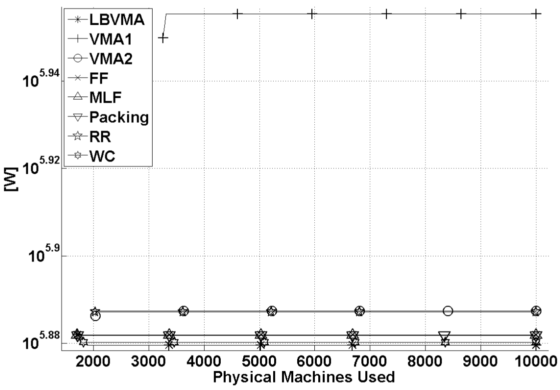

(d) $b = 100$ W, $x^* = 0.9 \cdot C$

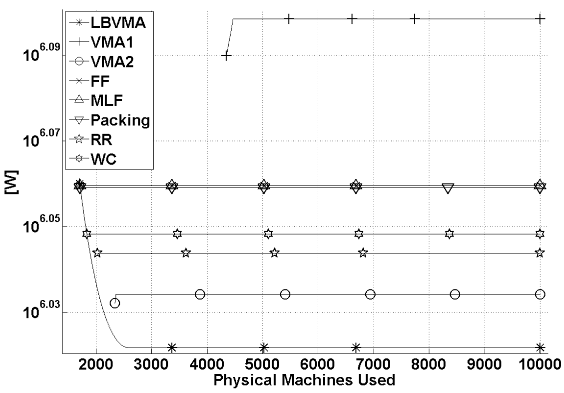

(e) $b = 120$ W, $x^* = 0.65 \cdot C$

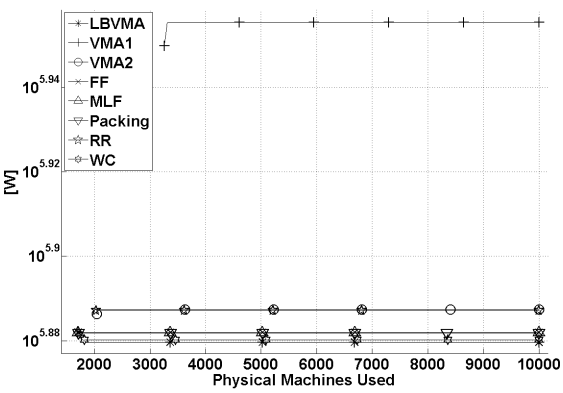

(f) $b = 120$ W, $x^* = 0.9 \cdot C$

Figure 4.12: (C, m)-VMA: Comparing the power consumed by the different assignment algorithms fixing for $C = 41.6$ GCPS, $\alpha = 2$ and 2 different values of x^* while varying b for Trace A (Synthetic traces).

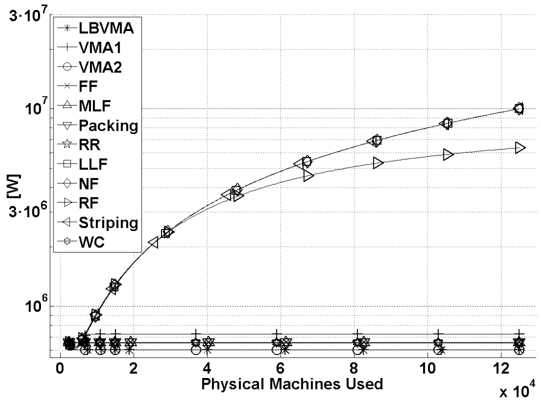

(a) $b = 80$ W, $x^* = 0.65 \cdot C$

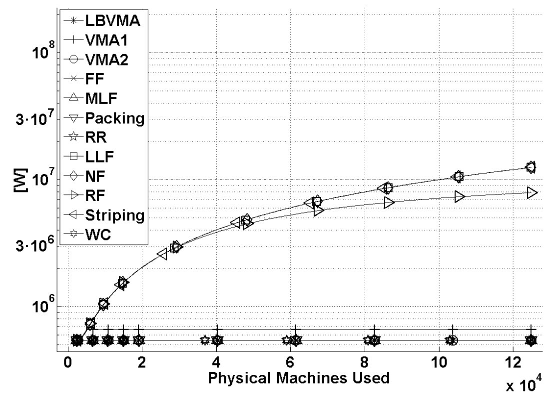

(b) $b = 80$ W, $x^* = 0.9 \cdot C$

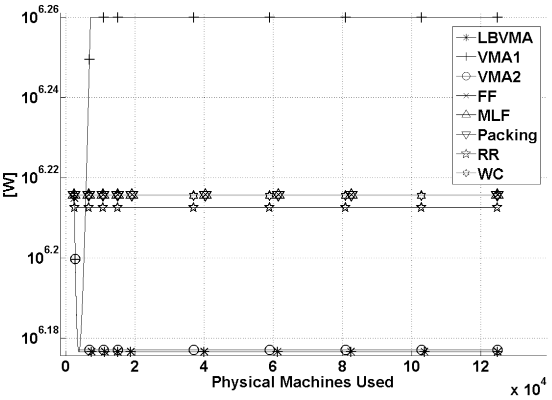

(c) $b = 100$ W, $x^* = 0.65 \cdot C$

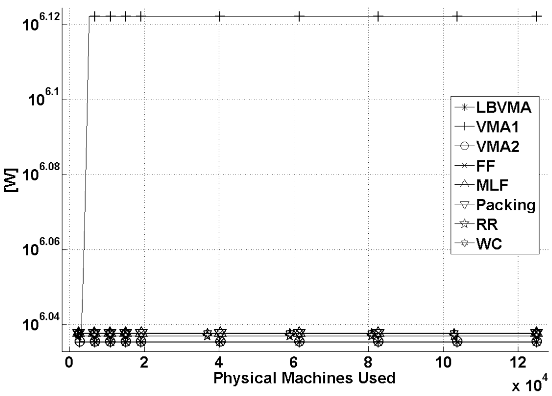

(d) $b = 100$ W, $x^* = 0.9 \cdot C$

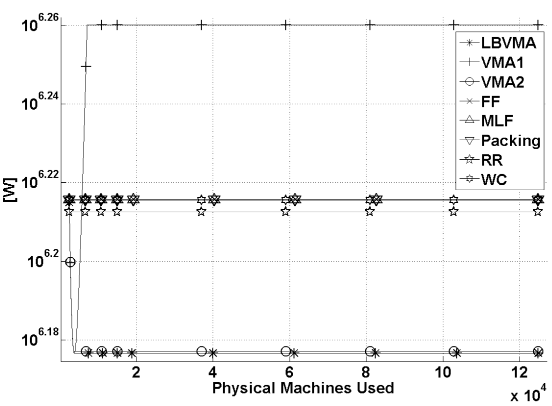

(e) $b = 120$ W, $x^* = 0.65 \cdot C$

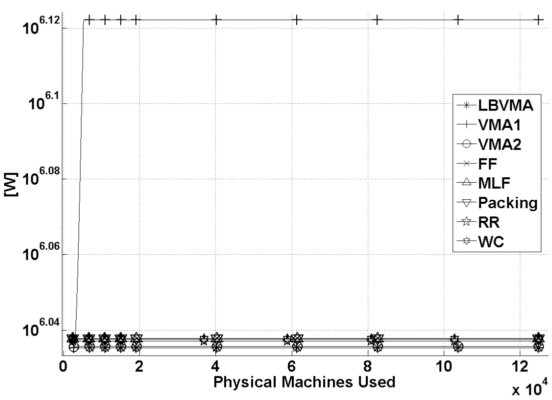

(f) $b = 120$ W, $x^* = 0.9 \cdot C$

Figure 4.13: (C, m)-VMA: Comparing the power consumed by the different assignment algorithms for $C = 41.6$ GCPS, $\alpha = 2$ and 2 different values of x^* for Trace B (Google traces).

reason for these results is basically a limitation of our model when trying to emulate different PMs changing only the value of b. In this case, as μ depends also on b absorbs the changes on its value, resulting in an identical behavior when x^* is kept constant. The results are, therefore, similar to the ones achieved in the second scenario.

4.6 Discussion

We discuss in this section practical issues that must be addressed to apply our results to production environments.

Heterogeneity of Servers. For the sake of simplicity, we assume in our model that all servers in a data center are identical. We believe this is reasonable, considering that modern data centers are usually built with homogeneous commodity hardware. Nevertheless, the proposed model and derived results are also amenable to heterogeneous data center environments. In a heterogeneous data center, servers can be categorized into several groups with identical servers in each group. Then, different types of applications can be assigned to server groups according to their resource requirements. The VMA model presented here can be applied to the assignment problem of allocating tasks from the designated types of applications (especially CPU-intensive ones) to each group of servers. The approximation results we derive in this paper can be then combined with server-group assignment approximation bounds (out of the scope of our work) for energy-efficient task assignment in real data centers, regardless of the homogeneity of servers.

Consolidation. Traditionally, consolidation has been understood as a bin packing problem [75, 95], where VMs are assigned to PMs attempting to minimize the number of active PMs. However, the results we derived in this chapter, as well as the results shown in Chapter 3, show that such approach is not energy-efficient. Indeed, we showed that PM's should be loaded up to x^* to reduce energy consumption, even if this requires having more active PMs.

VM arrival and departure. When a new VM arrives to the system, or an assigned VM departs, adjustments to the assignment may improve energy efficiency. Given that the cost of VM migration is nowadays decreasing dramatically, our offline positive results can also be accommodated by reassigning VMs whenever the set of VM demands changes. Should the cost of migration be high to reassign after each VM arrival or departure, time could be divided in epochs buffering newly arrived VM demands until the beginning of the next epoch, when all (new and old) VMs would be reassigned (if necessary) running our offline approximation algorithm.

Multi-resource scheduling. This work focuses on CPU-intensive jobs (VMs) such as MapReduce-like tasks [39] which are representative in production datacenters. As the CPU is generally the dominant energy consumer in a server, assigning VMs according to CPU workloads entails energy efficiency. However, there exist types of jobs demanding heavily other computational resources, such as memory and/or storage. Although these resources have limited impact on a server's energy consumption, VMs performance may be degraded if they become the bottleneck resource in the system. In this case, a joint optimization of multiple resources (out of the scope of our work) is necessary for VMA.

Implementation on real systems. Most of the allocation algorithms that we have tested in Section 4.5 are already available in popular cloud platforms like OpenNebula [84] or Eucalyptus [44]. Including another allocation policy, such as our algorithms, in the cloud controllers of these and other platforms (e.g. Apache Mesos [12]) is feasible. Introducing our algorithms would make those platforms power efficient, providing power-aware allocation policies. This feature is not found on any of the algorithms tested in Section 4.5 except for the *Watts per Core* algorithm from [63]. We leave such integration for future work.

4.7 Conclusions

In this chapter, we have studied a particular case of the generalized assignment problem with applications to Cloud Computing. We have considered the problem of assigning virtual machines (VMs) to physical machines (PMs) so that the power consumption is minimized, a problem that we call virtual machine assignment (VMA). In our theoretical analysis, we have shown that the decision version of (C, m)-VMA problem is strongly NP-complete. We have shown as well that the $(C, \cdot)$-VMA, $(\cdot, m)$-VMA and $(\cdot, \cdot)$-VMA problems are strongly NP-hard. Hence, there is no FPTAS for these optimization problems. We have shown the existence of a PTAS that solves the $(\cdot, \cdot)$-VMA and $(\cdot, m)$-VMA offline problems. On the other hand, we have proved lower bounds on the approximation ratio of the $(C, \cdot)$-VMA and (C, m)-VMA problems. With respect to the online version of these problems, we have proved upper and lower bounds on the competitive ratio of the $(\cdot, \cdot)$-VMA, $(C, \cdot)$-VMA, $(\cdot, m)$-VMA, and (C, m)-VMA problems.

Chapter 5

Conclusions

Cloud computing is alive and probably no more than a toddler yet. However, even in this early age we are already able to see many of the advantages, surely not all, that cloud computing can offer us. Despite of all these advantages we must not forget about its drawbacks and this was one of the targets of this master thesis.

The main objective of this master thesis was to face one of the main problems of data centers, energy consumption, which was introduced in Part I. Reducing the costs associated to this issue is not only an economical target but also an environmental problem. Power is generated at a cost and, indefinitely increasing our power consumption will have a prejudicial effect on our world. It is our duty, as researchers, to care about sustainability and ensure that future times will be better.

For these reasons we tried to provide of tools to directly or indirectly tackle these problems. This master thesis presents our contributions to the field of energy consumption in data centers. We studied how to reduce power consumption by optimizing how virtual machines are placed in physical machines, the virtual machine assignment problem; and also studied how servers consume power and energy in a data center, characterizing the contribution of each component and confirming the superlinearity on the load that we assumed in the virtual machine assignment problem.

Starting with servers, in Chapter 3 we characterized how different components of a data center server consumes power. One of the most interesting aspects, and also differentiator from previous works, is how we studied the effect of having multiple cores and how varying their frequency affects, not only to their energy consumption, but also to other components. The main idea to be extracted here is that a server is a puzzle, where each piece is a component, and some pieces needed others to perform a task, thus, being affected by how those other pieces are being operated. This study throw very interesting results as clearly confirming the superlinearity of power consumption on the load in the server as well as showing that the Active Cycles per Second are a proper unit to measure the load in a server. We concluded this part of the study showing how the characterization of these components can be used to predict the energy consumption of an application from its profiling.

Then, taking advantage of some of the conclusions of the previous chapter, as the superlinearity of the power consumption on the load, we studied the virtual machine assignment problem and how such a superlinear power consumption model affects it. We thoroughly analyzed four different cases, depending on whether the capacity and the number of the servers where bounded or not. For each one of them, when possible, we provided upper and lower bounds on the approximation ratio, as well as upper and lower bounds on the competitive ratio of the algorithms we proposed for them. We also proved, by simulation, that the algorithms we proposed can obtain substantially cheaper solutions, from the point of view of power consumption of the system, than other algorithms proposed in the literature.

In general we proved that huge amounts of energy can still be saved in data centers with no need of updating the already deployed hardware, in some cases, and that new solutions are waiting for the new generation of servers and network devices, ready to optimize the way energy is used in data centers.

5.1 Future Work and Open Problems

As we said at the beginning of this section, Cloud Computing is still a toddler. This means that the amount of open problems is practically unlimited and that, usually, an answer bring up two more questions. Allow us, then, to focus only in the particular issues that we have worked at during this master thesis.

We are aware that the characterization of power consumption we presented in Chapter 3 is neither complete nor perfect. However, this is a key problem because the more accurately can we predict the power required by a server in a particular situation or the energy it will consume when a certain application is run, the better we will be able to assign tasks to servers or maneuver in case moving virtual machines across our servers is needed in order to increase the efficiency of our data center. In any case, we believe that some of the aspects we have proposed here will be helpful for future works. The most obvious open problems in this case are, first, to properly characterize the power consumption due to the RAM and, second, improve power characterizations so the accuracy is increased.

Finally, regarding the virtual machine assignment problem presented in Chapter 4, our future work will consider the possibility that the load incurred by a VM changes over time or that the assignment of VMs to PMs is not final (and VMs can migrate, maybe at a cost). In fact, if the migration of VMs is available for free, our offline positive results can also be used in these new models, since an offline approximation algorithm can be run each time a load changes or a new VM arrives. Then, the VMs can be redistributed accordingly. Another future extension of the model will consider that the power consumption of a feasible solution of the VMA problem depends on several parameters simultaneously (e.g., memory space or communication bandwidth, in addition to processing load). Finally, as stated in Section 3.5, we plan to deploy our algorithm in a cloud platform, probably OpenNebula, and compare the performance of our proposed algorithms against other state-of-the-art allocation algorithms such as the ones analyzed in Section 4.5.

References

[1] Amazon web services. `http://aws.amazon.com`. Accessed August 27, 2012.

[2] Citrix. `http://www.citrix.com`. Accessed August 27, 2012.

[3] Infiniband. `http://www.infinibandta.org`.

[4] Rackspace. `http://www.rackspace.com`. Accessed August 27, 2012.

[5] Gerald Aigner and William H Whitted. Modular data center, October 9 2007. US Patent 7,278,273.

[6] Mohammad Al-Fares, Alexander Loukissas, and Amin Vahdat. A scalable, commodity data center network architecture. In *ACM SIGCOMM Computer Communication Review*, volume 38, pages 63–74. ACM, 2008.

[7] Noga Alon, Yossi Azar, Gerhard J. Woeginger, and Tal Yadid. Approximation schemes for scheduling. In Michael E. Saks, editor, *SODA*, pages 493–500. ACM/SIAM, 1997.

[8] Noga Alon, Yossi Azar, Gerhard J Woeginger, and Tal Yadid. Approximation schemes for scheduling on parallel machines. *Journal of Scheduling*, 1(1):55–66, 1998.

[9] Matthew Andrews, Antonio Fernández Anta, Lisa Zhang, and Wenbo Zhao. Routing for power minimization in the speed scaling model. *IEEE/ACM Trans. Netw.*, 20(1):285–294, 2012.

[10] Matthew Andrews, Spyridon Antonakopoulos, and Lisa Zhang. Minimum-cost network design with (dis)economies of scale. In *Proc. of 51-st Annual IEEE Symposium on Foundations of Computer Science*, pages 585–592, 2010.

[11] Antonio Antoniadis, Sungjin Im, Ravishankar Krishnaswamy, Benjamin Moseley, Viswanath Nagarajan, Krik Pruhs, and Cliff Stein. Hallucination helps: Energy efficient circuit routing. In Erlebach and Persiano [43].

[12] Apache. Apache mesos. `http://http://mesos.apache.org/`, 2014. Accessed December 11th, 2014.

[13] M Bailey, M Eastwood, T Grieser, L Borovick, V Turner, and RC Gray. Special study: Data center of the future, 2007.

[14] Jayant Baliga, Robert WA Ayre, Kerry Hinton, and Rodney S Tucker. Green cloud computing: Balancing energy in processing, storage, and transport. *Proceedings of the IEEE*, 99(1):149–167, 2011.

[15] Nikhil Bansal, Ho-Leung Chan, and Kirk Pruhs. Speed scaling with an arbitrary power function. In *Proc. of 20-th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 693–701, 2009.

[16] Nikhil Bansal, Anupam Gupta, Ravishankar Krishnaswamy, Viswanath Nagarajan, Kirk Pruhs, and Cliff Stein. Multicast routing for energy minimization using speed scaling. In *MedAlg*, pages 37–51, 2012.

[17] Luiz André Barroso and Urs Hölzle. The case for energy-proportional computing. *IEEE Computer*, 40(12):33–37, 2007.

[18] Luiz André Barroso and Urs Hölzle. The datacenter as a computer: An introduction to the design of warehouse-scale machines. *Synthesis lectures on computer architecture*, 2009.

[19] Luiz André Barroso and Urs Hölzle. The datacenter as a computer: An introduction to the design of warehouse-scale machines, 2nd edition. *Synthesis lectures on computer architecture*, 2013.

[20] Robert Basmadjian, Nasir Ali, Florian Niedermeier, Hermann de Meer, and Giovanni Giuliani. A methodology to predict the power consumption of servers in data centres. In *ACM e-Energy*, pages 1–10, 2011.

[21] Robert Basmadjian and Hermann de Meer. Evaluating and modeling power consumption of multi-core processors. In *Third International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet (e-Energy)*, pages 1–10. IEEE, 2012.

[22] CL Belady. Projecting annual new datacenter construction market size. *Microsoft, Global Foundation Services Report*, 2011.

[23] Umesh Bellur, Chetan S. Rao, and Madhu Kumar SD. Optimal placement algorithms for virtual machines. *CoRR*, abs/1011.5064, 2010.

[24] Anton Beloglazov, Jemal Abawajy, and Rajkumar Buyya. Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing. *Future Generation Computer Systems*, 28(5):755–768, 2012.

[25] Bruno Biais and Paul Woolley. High frequency trading. *Manuscript, Toulouse University, IDEI*, 2011.

[26] Juan Felipe Botero, Xavier Hesselbach, Michael Duelli, Daniel Schlosser, Andreas Fischer, and Hermann de Meer. Energy efficient virtual network embedding. *IEEE Communications Letters*, 16(5):756–759, 2012.

[27] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.

[28] Alaa Brihi and Waltenegus Dargie. Dynamic voltage and frequency scaling in multimedia servers. In *IEEE AINA*, 2013.

[29] M. Cardosa, A. Singh, H. Pucha, and A. Chandra. Exploiting spatio-temporal tradeoffs for energy-aware mapreduce in the cloud. In *Cloud Computing (CLOUD), 2011 IEEE International Conference on*, pages 251 –258, 2011.

[30] P Castagna. Having fun with pagerank and mapreduce. *Hadoop User Group UK talk. Available: http://static. last. fm/johan/huguk-20090414/paolo_castagna-pagerank. pdf*.

[31] Deeparnab Chakrabarty, Chandra Chekuri, Sanjeev Khanna, and Nitish Korula. Approximability of capacitated network design. In *IPCO*, pages 78–91, 2011.

[32] Ashok K. Chandra and C. K. Wong. Worst-case analysis of a placement algorithm related to storage allocation. *SIAM J. Comput.*, 4(3):249–263, 1975.

[33] Jeffrey S Chase, Darrell C Anderson, Prachi N Thakar, Amin M Vahdat, and Ronald P Doyle. Managing energy and server resources in hosting centers. In *ACM SIGOPS Operating Systems Review*, volume 35, pages 103–116, 2001.

[34] Shih-Chang Chen, Chih-Chun Lee, Hsi-Ya Chang, Kuan-Chou Lai, Kuan-Ching Li, and Chunming Rong. Energy-aware task consolidation technique for cloud computing. In *Proceedings of the IEEE Third International Conference on Cloud Computing Technology and Science*, pages 115–121, 2011.

[35] Susanta Nanda Tzi-cker Chiueh and Stony Brook. A survey on virtualization technologies. *RPE Report*, pages 1–42, 2005.

[36] Michael Chlistalla, Bernhard Speyer, Sabine Kaiser, and Thomas Mayer. High-frequency trading. *Deutsche Bank Research*, pages 1–19, 2011.

[37] R. A. Cody and Edward G. Coffman Jr. Record allocation for minimizing expected retrieval costs on drum-like storage devices. *J. ACM*, 23(1):103–115, 1976.

[38] Standard Performance Evaluation Corporation. Spec power benchmark, 2007.

[39] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, 2008.

[40] Dimitris Economou, Suzanne Rivoire, Christos Kozyrakis, and Partha Ranganathan. Full-system power analysis and modeling for server environments. In *Proceedings of Workshop on Modeling, Benchmarking, and Simulation*, pages 70–77, 2006.

[41] Kemal Efe and Antonio Fernández. Mesh-connected trees: A bridge between grids and meshes of trees. *IEEE Trans. Parallel Distrib. Syst.*, 7(12):1281–1291, 1996.

[42] Leah Epstein and Jiri Sgall. Approximation schemes for scheduling on uniformly related and identical parallel machines. *Algorithmica*, 39(1):43–57, 2004.

[43] Thomas Erlebach and Giuseppe Persiano, editors. *Approximation and Online Algorithms - 10th International Workshop, WAOA 2012, Ljubljana, Slovenia, September 13-14, 2012, Revised Selected Papers*, volume 7846 of *Lecture Notes in Computer Science*. Springer, 2013.

[44] Eucalyptus. Eucalyptus. `http://www.eucalyptus.com/`. Accessed January 20th, 2013.

[45] W. Fernandez de la Vega and G.S. Lueker. Bin packing can be solved within $1 + \epsilon$ in linear time. *Combinatorica*, 1(4):349–355, 1981.

[46] Anshul Gandhi, Mor Harchol-Balter, Rajarshi Das, and Charles Lefurgy. Optimal power allocation in server farms. In *ACM SIGMETRICS*, pages 157–168, 2009.

[47] Andres Garcia-Saavedra, Pablo Serrano, Albert Banchs, and Giuseppe Bianchi. Energy consumption anatomy of 802.11 devices and its implication on modeling and design. In *ACM CoNEXT*, pages 169–180, 2012.

[48] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1979.

[49] Rong Ge, Xizhou Feng, Shuaiwen Song, Hung-Ching Chang, Dong Li, and Kirk W Cameron. Powerpack: Energy profiling and analysis of high-performance systems and applications. *Parallel and Distributed Systems, IEEE Transactions on*, 21(5):658–671, 2010.

[50] Albert Greenberg, James R. Hamilton, Navendu Jain, Srikanth Kandula, Changhoon Kim, Parantap Lahiri, David A. Maltz, Parveen Patel, and Sudipta Sengupta. Vl2: a scalable and flexible data center network. *SIGCOMM Comput. Commun. Rev.*, 39:51–62, August 2009.

[51] Chuanxiong Guo, Guohan Lu, Dan Li, Haitao Wu, Xuan Zhang, Yunfeng Shi, Chen Tian, Yongguang Zhang, and Songwu Lu. Bcube: a high performance, server-centric network architecture for modular data centers. *SIGCOMM Comput. Commun. Rev.*, 39:63–74, August 2009.

[52] Chuanxiong Guo, Haitao Wu, Kun Tan, Lei Shi, Yongguang Zhang, and Songwu Lu. Dcell: a scalable and fault-tolerant network structure for data centers. *SIGCOMM Comput. Commun. Rev.*, 38:75–86, August 2008.

[53] Anupam Gupta, Ravishankar Krishnaswamy, and Kirk Pruhs. Online primal-dual for non-linear optimization with applications to speed scaling. In Erlebach and Persiano [43], pages 173–186.

[54] Mark D. Hansen. Approximation algorithms for geometric embeddings in the plane with applications to parallel processing problems. In *FOCS, 1989., 30th Annual Symposium on*, pages 604–609. IEEE, 1989.

[55] Ward Van Heddeghem, Sofie Lambert, Bart Lannoo, Didier Colle, Mario Pickavet, and Piet Demeester. Trends in worldwide ICT electricity consumption from 2007 to 2012. *Computer Communications*, Submitted.

[56] Brandon Heller, Srini Seetharaman, Priya Mahadevan, Yiannis Yiakoumis, Puneet Sharma, Sujata Banerjee, and Nick McKeown. Elastictree: saving energy in data center networks. In *Proceedings of the 7th USENIX conference on Networked systems design and implementation*, NSDI'10, pages 249–264, Berkeley, CA, USA, 2010. USENIX Association.

[57] Joseph L. Hellerstein. Google cluster data. Google research blog, January 2010. Posted at `http://googleresearch.blogspot.com/2010/01/google-cluster-data.html`.

[58] Lei Huang, Qin Jia, Xin Wang, Shuang Yang, and Baochun Li. Pcube: Improving power efficiency in data center networks. In *Cloud Computing (CLOUD), 2011 IEEE International Conference on*, pages 65–72. IEEE, 2011.

[59] IBM. Ibm's integrated server room. `http://www-935.ibm.com/services/in/en/it-services/integrated-server-room.html`.

[60] IEEE Std. 802.3az. Energy Efficient Ethernet, 2010.

[61] Sungjin Im, Benjamin Moseley, and Kirk Pruhs. Online scheduling with general cost functions. In *Proc. of 23-rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1254–1265, 2012.

[62] Sandy Irani, Sandeep Shukla, and Rajesh Gupta. Algorithms for power savings. *ACM TALG*, 3(4):41, 2007.

[63] R. Jansen and P.R. Brenner. Energy efficient virtual machine allocation in the cloud. In *Green Computing Conference and Workshops (IGCC), 2011 International*, pages 1–8, 2011.

[64] Bhavani Krishnan, Hrishikesh Amur, Ada Gavrilovska, and Karsten Schwan. VM power metering: feasibility and challenges. *ACM SIGMETRICS Performance Evaluation Review*, 38(3):56–60, 2011.

[65] Dara Kusic, Jeffrey O Kephart, James E Hanson, Nagarajan Kandasamy, and Guofei Jiang. Power and performance management of virtualized computing environments via lookahead control. *Cluster computing*, 12(1):1–15, 2009.

[66] Dara Kusic, Jeffrey O. Kephart, James E. Hanson, Nagarajan Kandasamy, and Guofei Jiang. Power and performance management of virtualized computing environments via lookahead control. *Cluster Computing*, 12(1):1–15, 2009.

[67] Etienne Le Sueur and Gernot Heiser. Dynamic voltage and frequency scaling: The laws of diminishing returns. In *Proceedings of the 2010 international conference on Power aware computing and systems*, pages 1–8. USENIX Association, 2010.

[68] F. Thomson Leighton. *Introduction to parallel algorithms and architectures: array, trees, hypercubes*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1992.

[69] Adam Wade Lewis, Soumik Ghosh, and Nian-Feng Tzeng. Run-time energy consumption estimation based on workload in server systems. *HotPower'08*, pages 17–21, 2008.

[70] Chenguang Liu, Jianzhong Huang, Qiang Cao, Shenggang Wan, and Changsheng Xie. Evaluating energy and performance for server-class hardware configurations. In *IEEE NAS*, pages 339–347, 2011.

[71] Ning Liu, Ziqian Dong, and Roberto Rojas-Cessa. Task and server assignment for reduction of energy consumption in datacenters. In *Proceedings of the IEEE 11-th International Symposium on Network Computing and Applications*, pages 171–174, 2012.

[72] F. Machida, M. Kawato, and Y. Maeno. Redundant virtual machine placement for fault-tolerant consolidated server clusters. In *Network Operations and Management Symposium (NOMS), 2010 IEEE*, pages 32 –39, 2010.

[73] C.C.T. Mark, D. Niyato, and Tham Chen-Khong. Evolutionary optimal virtual machine placement and demand forecaster for cloud computing. In *Advanced Information Networking and Applications (AINA), 2011 IEEE International Conference on*, pages 348 –355, 2011.

[74] K. Mills, J. Filliben, and C. Dabrowski. Comparing vm-placement algorithms for on-demand clouds. In *Proceedings of the IEEE Third International Conference on Cloud Computing Technology and Science*, pages 91–98, 2011.

[75] M. Mishra and A. Sahoo. On theory of vm placement: Anomalies in existing methodologies and their mitigation using a novel vector based approach. In *Cloud Computing (CLOUD), 2011 IEEE International Conference on*, pages 275 –282, 2011.

[76] Mayank Mishra and Anirudha Sahoo. On theory of vm placement: Anomalies in existing methodologies and their mitigation using a novel vector based approach. In *IEEE CLOUD*, pages 275–282, 2011.

[77] Akihiko Miyoshi, Charles Lefurgy, Eric Van Hensbergen, Ram Rajamony, and Raj Rajkumar. Critical power slope: understanding the runtime effects of frequency scaling. In *ACM ICS'02*, pages 35–44, 2002.

[78] Justin D Moore, Jeffrey S Chase, Parthasarathy Ranganathan, and Ratnesh K Sharma. Making scheduling "cool": Temperature-aware workload placement in data centers. In *USENIX annual technical conference, General Track*, pages 61–75, 2005.

[79] Ripal Nathuji and Karsten Schwan. Virtualpower: coordinated power management in virtualized enterprise systems. In *SOSP*, pages 265–278, 2007.

[80] Hien Nguyen Van, Frederic Dang Tran, and Jean-Marc Menaud. Autonomic virtual resource management for service hosting platforms. In *Proceedings of the 2009 ICSE Workshop on Software Engineering Challenges of Cloud Computing*, CLOUD '09, pages 1–8. IEEE Computer Society, 2009.

[81] Radhika Niranjan Mysore, Andreas Pamboris, Nathan Farrington, Nelson Huang, Pardis Miri, Sivasankar Radhakrishnan, Vikram Subramanya, and Amin Vahdat. Portland: a scalable fault-tolerant layer 2 data center network fabric. *SIGCOMM Comput. Commun. Rev.*, 39:39–50, August 2009.

[82] Leonard Nonde, Taisir EH El-Gorashi, and Jaafar MH Elmirghani. Energy efficient virtual network embedding for cloud networks. *Journal on Lightwave Technology*, 33(9):1828–1849, 2015.

[83] Bruce Nordman and K Christensen. Reducing the energy consumption of network devices. *IEEE 802.3 tutorial*, pages 1–30, 2005.

[84] OpenNebula. Opennebula. `http://opennebula.org/`. Accessed January 20th, 2013.

[85] Ramya Raghavendra, Parthasarathy Ranganathan, Vanish Talwar, Zhikui Wang, and Xiaoyun Zhu. No power struggles: Coordinated multi-level power management for the data center. In *ACM SIGARCH Computer Architecture News*, volume 36, pages 48–59. ACM, 2008.

[86] Steve Scott, Dennis Abts, John Kim, and William J. Dally. The blackwidow high-radix clos network. In *ISCA*, pages 16–28, 2006.

[87] Shekhar Srikantaiah, Aman Kansal, and Feng Zhao. Energy aware consolidation for cloud computing. In *Proceedings of the 2008 conference on Power aware computing and systems*, volume 10 of *HotPower'08*. USENIX Association, 2008.

[88] Matt Stansberry. Important to recognize the dramatic improvement in data center efficiency, 2012.

[89] Wendy Torell. Tco analysis of a traditional data center vs. a scalable, containerized data center. Technical Report 164, Schneider Electric, 2012.

[90] R. Van den Bossche, K. Vanmechelen, and J. Broeckhove. Cost-efficient scheduling heuristics for deadline constrained workloads on hybrid clouds. In *Proceedings of the IEEE Third International Conference on Cloud Computing Technology and Science*, pages 320–327, 2011.

[91] Arunchandar Vasan, Anand Sivasubramaniam, Vikrant Shimpi, T Sivabalan, and Rajesh Subbiah. Worth their Watts? - An empirical study of datacenter servers. In *IEEE HPCA*, pages 1–10, 2010.

[92] Vijay V. Vazirani. *Approximation Algorithms*. Springer, March 2004.

[93] H. Viswanathan, E.K. Lee, I. Rodero, D. Pompili, M. Parashar, and M. Gamell. Energy-aware application-centric vm allocation for hpc workloads. In *Parallel and Distributed Processing Workshops and Phd Forum (IPDPSW), 2011 IEEE International Symposium on*, pages 890 –897, 2011.

[94] M Mitchell Waldrop. Data center in a box. *Scientific American*, 297(2):90–93, 2007.

[95] Meng Wang, Xiaoqiao Meng, and Li Zhang. Consolidating virtual machines with dynamic bandwidth demand in data centers. In *IEEE INFOCOM*, pages 71–75, 2011.

[96] Mark Weiser, Brent Welch, Alan Demers, and Scott Shenker. Scheduling for reduced CPU energy. In *Mobile Computing*, pages 449–471. Springer, 1996.

[97] Jing Xu and José Fortes. A multi-objective approach to virtual machine management in datacenters. In *Proceedings of the 8th ACM international conference on Autonomic computing*, ICAC '11, pages 225–234. ACM, 2011.