POLYTECHNICS OF TURIN

MasterDegree in ICT 4 Smart Societies

MasterThesis

Load managment in green data center

Supervisors

Prof. Michela Meo

Dr. Daniela Renga

Candidate

Damiano Scantamburlo

July 2020

Acknowledgements

To my family which always back me and gave me all the possibilities by sustaining me economically and personally. To my friends and my girlfriend which have always been by my side. To my supervisors which patiently tutor and supervise me in this thesis path.

Table of Contents

List of Figures

Abstract

Data centres usage and impact is a sensible matter, even more nowadays. Data centres energy impact is not negligible, driven by continuous technologies improvement and the research of higher performances. Among several methods to improve efficiency, this thesis work main goal is to study energy and cost impact of a proper and more tailored load management approach. The thesis aims to implement and extend a previous work born by a collaboration between University of Calabria and Polytechnics of Turin. The case study remains the same and it is represented by 4 hypothetical data centres placed in different locations. This spatial dislocation implies different energy cost and different renewable production among the 4 structure at the same hours of the day. This work, and this thesis as well, studies the efficiency gain brought by performing a load migration between these 4 data centres in order to match better renewables production and hourly energy cost. The load moved consisted in a set of virtual machines referring to real data from a Telecom data centre collected by researchers of University of Calabria. For this purpose, a java program was implemented to properly simulate the migration of virtual machines over this 4 data centre giving back energy parameters results. This thesis continues to analyse these performances making further considerations and implementations focusing on the influence of a proper virtual machine characterization. The virtual machine features considered are principally cpu, ram and disk usage. The analysis use these metrics as constraints to migration, this permits to calculate eventual benefits based on the main characteristic of each virtual machine which results evaluated individually. The output is an overall consideration on performances, sum of the 4 data centres results. The final goal is to understand and contextualize these results to find relations between the virtual machine characteristics and an optimal migration policy.

Acronyms

DC

Data Center

VM

virtual machine

PUE

Power Usage Effectiveness

DR

Demand Response

PV

Photo-Voltaic

TH

threshold

Chapter 1

Introduction

1.1 Green Data Center

Green data Centers (DC) are facilities where data are stored, managed and elaborated and in which are adopted energy efficiency technologies. Indeed, DC are huge structures, full of an high number of hardware computing components, which summed together, have a very high power consumption. So the ideal aim of green data centers is to reduce energy cost and footprint without compromising the performance by investing and implementing in new hardware and software technologies.

The continuous research for higher performance, the world' seek for IT resources and the technology improvement brought data center employment to grow exponentially during the past years. Moreover, data centers are directly connected to data that are one of the main resources of this technological era where data importance and magnitude grows dramatically over last years[1] and now a good architecture to manage them is an essential requirement.

Data center are indeed complex architectures which permit to perform many different tasks and they evolve to adapt to modern user in a society more and more IT compliant.

Data centers buildings themselves are full of modern electronic components and

---

[1]citing one of the latest Cisco report [1]:

- "Globally, Internet traffic will grow 3.7-fold from 2017 to 2022, a compound annual growth rate of 30%";
- "Globally, Internet traffic will reach 350.8 EB per month by 2022, up from 96.0 EB per month in 2017";

This statements help to contextualize better the actual and future importance of data and the magnitude of the data management problem.

each of them have an energy consumption. These buildings importance make them a necessary asset and the relative world electricity consumption grow accordingly and now energy demand is not negligible and brought the consequent necessity of improve efficiency and reduce maintenance cost.

Power Usage Effectiveness (PUE) 1.1 is one of the main metric to indicate the energy efficiency of a data center.

$$PUE = \frac{P_{DC}}{P_{comp}} \tag{1.1}$$

The numerator is the total power supplied to the data center's structure (P_{DC}), while the denominator is the fraction used to effectively run all data center computational task (P_{comp}). Therefore, ideal PUE is 1, indicating a data center with no energy waste where all the electricity demand is used exclusively for the DC operation.

Figure 1.1 is an histogram representing PUE distribution based on a JRC report [2]. The data refers to a sample set of 268 data centers principally located in Europe in 2016 and the average PUE is 1.8. There are a small subset of DCs with high efficiency and, on the contrary, many others with very poor performance. A value of PUE near 2 means that only 50% of the energy supplied is effectively used for the DC computational purpose. Now, in 2020 the average global average PUE drops to 1.59, this value is still improvable and the DCs in 2016 with PUE lower than 1.5 proves that.

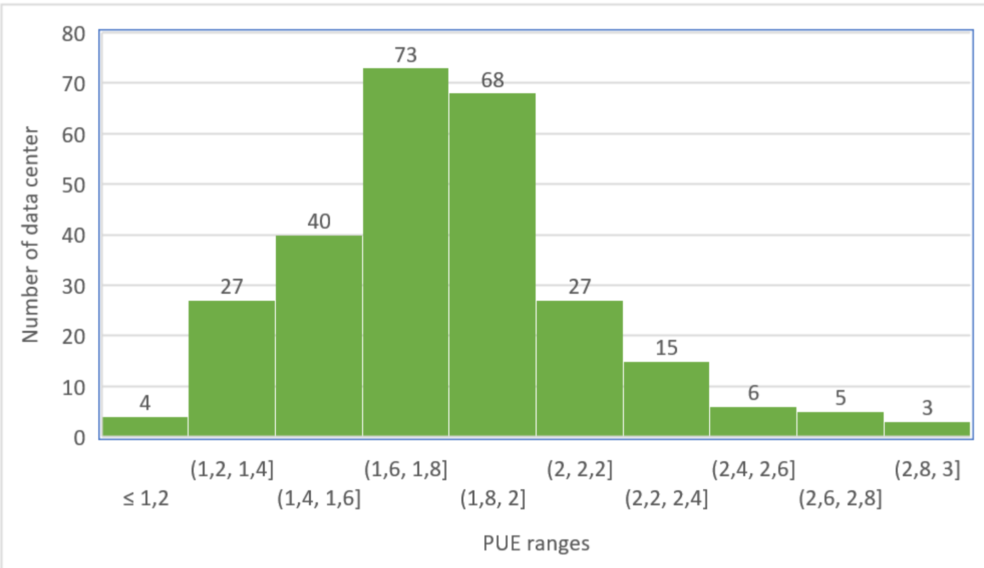

Figure 1.1: PUE distribution [2]

The histogram in figure 1.2 , brought by a study of the synergy research group [3], shows an exponential growth in cloud infrastructure services investments (one of many DC main feature); along with a massive and constant value of DC's hardware and software investments. This second aspect seems to be stable but it should be considered the difference in the last 10 years is about 40 billion dollars, (about 40% of the total) and the money invested are already 60 billion in 2009.

We can also make further considerations: for example, the 2019 data are predicted, and indicate a peak beneath 100 billion of dollar. According to one of the last Canalys research, the real value reached in 2019 is 107 billion dollars and it is expected to reach 284 billion by 2024 [4]. Moreover, the prediction made previously does not consider the huge implications relative to the last changes and the pandemic problem. In fact, many recent articles[2]describe the rebound of smart working approach and necessity to rely on many data center services more than ever and with a reliability even higher than before.

[2]Covid pandemic, lockdown, distance learning, smart working, these are all words people becomes familiar with. Many companies and services change and adapt their approaches to new needs to match better the world situation. All these changes increase the burden the IT infrastructures have to bear. Moreover, also in a post pandemic period, there are high probability that some companies chose to maintain some of the adopted changes. There are many other considerations worth to be done about this topic but here are some articles (a subset) which highlight the importance of data center in a post covid world and the reliability needed to bear the higher data traffic demand:

- "Coronavirus Is Bad and Good News for the Data Center Switch Market" [5]
- "'You can't just stop.' How data centers are dealing with the coronavirus crisis" [6]
- "The future of data centers in a post Covid-19 world" [7]
- "Le sfide per i Data Center al tempo del Coronavirus" [8]

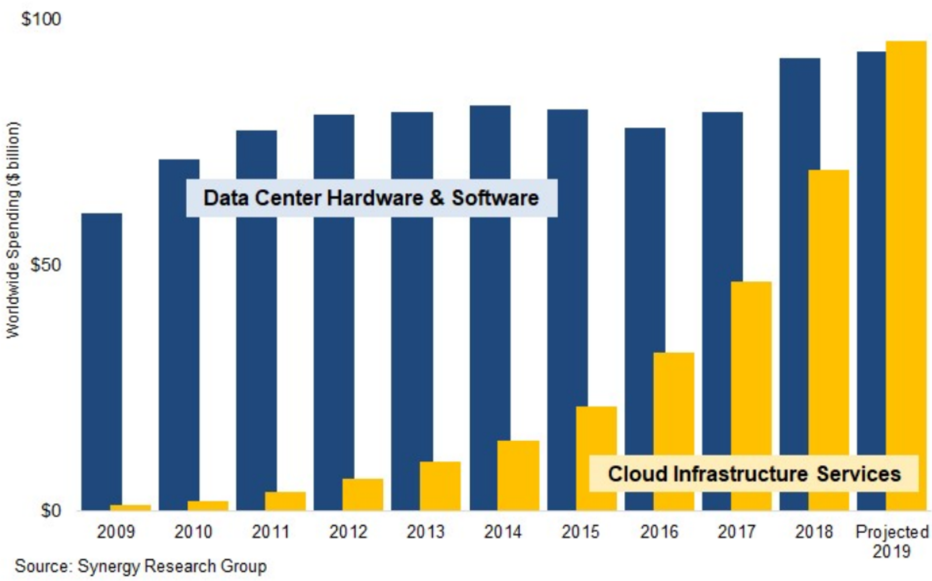

Figure 1.2: Enterprise spending on cloud and data centers [3]

1.2 Data center footprint and energy consideration

According to World Health Organization (WHO), carbon footprint is "a measure of the impact your activities have on the amount of carbon dioxide (CO_2) produced through the burning of fossil fuels and is expressed as a weight of CO_2 emissions produced in tonnes" [9]. CO_2 emission is only one of the many metrics useful to evaluate world pollution and energy impact. It still reamains a usefull term of comparison because well-known and used in many different study areas.

As stated in one of the latest international agency energy (IEA) reports [10] global data center electricity demand in 2019 was around 200 TWh corresponding to 0,8% of the total electricity demand [11]. Considering $1\frac{KgCO_2}{KWh}$ for coal propelled power plants (coal is the worst fuel regarding the CO_2 emission, slightly worse than petroleum) [12] we can estimate a worst case with roughly 200 mega tonnes of CO_2 (for comparison in 2016, Italy, emitted about 358 metric mega tonnes [13]).

So, it is clear data centers have an energy impact and a value of CO_2 emissions not negligible. Although to grasp even better the data center energy situation, some considerations should be added:

- The emissions estimated in the previous rows are about a worst case but coal is slowly decreasing its employment and many companies start to rely on renewable or in general more green resources than coal. Figure 1.3 shows a huge decrease in coal-fired power capacity subject to FID[3], confirmed also from the chart in figure 1.4 which displays a slow but steady increase in the usage percentage of renewables (this last data should be evaluated together with the historical and geopolitical worlds background[4]).

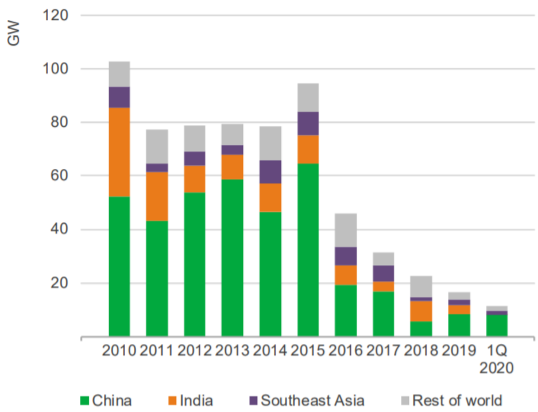

Figure 1.3: Coal-fired power generation capacity subject to an FID [10]

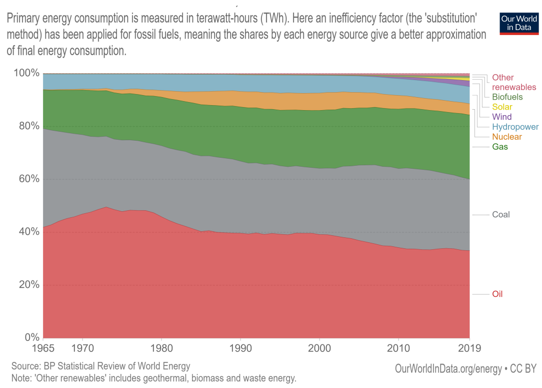

Figure 1.4: Energy world consumption by source [14]

- As mentioned in the previous section PUE is one of the common metrics use to evaluate energy efficiency and the PUE actual average is 1.59. Many companies achieved excellent PUE values over the years. Anyway it should be also considered the increasing effort in lowering PUE as it approaches the optimal value of 1. The plot in Fig.1.5 shows the excellent google 's data centers performances but it also display how many years were needed to reach the actual value; without investments and effort is impossible for world's data centers achieve optimal PUE values ,especially in a short amount of time.

[3]FID stands for Final Investment Decision and represents the final step of a long term assets managing project, this metric can be very useful to understand the economical implication and above all the future direction a company intends to follow.

[4]The figure 1.4 demonstrates that the world comes from a long history of carbon fossil usage. With the only exception of hydropower, the renewables start to be considered only in the last decades and clearly with a different timing among countries with a different development degree.

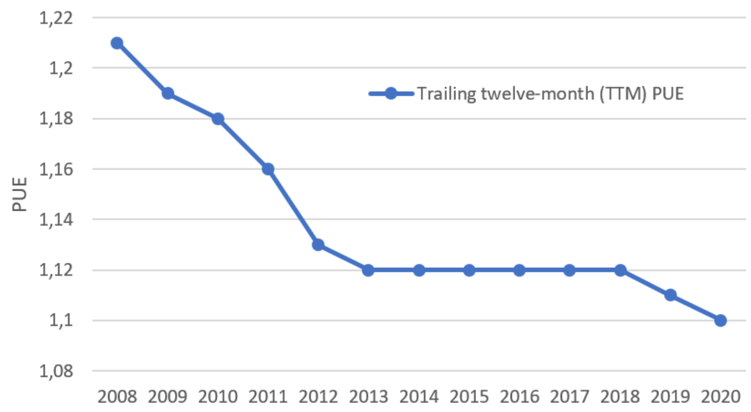

Figure 1.5: Google Data Center PUE performance [15]

- According to international energy agency (iea), the world increasing demand for data center's services is biased by a shift in higher efficiency hyper-scale structures [10]. Figure 1.6 shows this hyper-scale shift trend while, on the right, (figure 1.7) it is possible to observe the flat energy demand trend (the charts are elaborations based on iea data [10]).

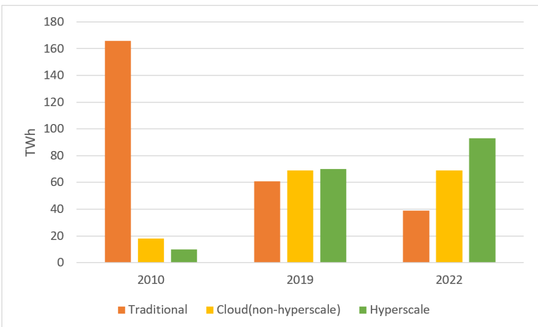

Figure 1.6: Global data center energy demand by data center type [16]

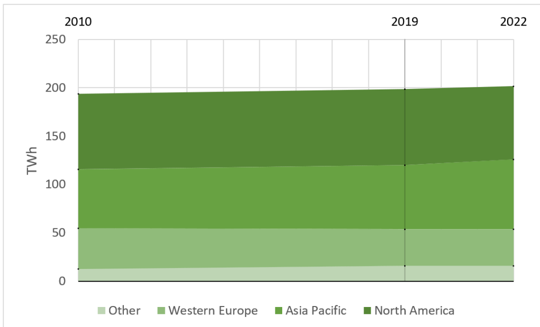

Figure 1.7: Global data center energy demand by region [17]

Hyper-scale data centers permit an improved efficiency and give the possibility to cushion better investments in high performance components. At the same time many distributed data centers are more difficult to track and can represent an unpredictable coefficient in the energy management considerations. Indeed these two charts show, from a certain point of view, the advantages of a single high efficiency architecture in comparison with a multi-distributed one.

Although Hyper-scale data centers manage larger capacity but at the same time need huge investments because their main feature may be represented from the ability to scale fast with high level of redundancy and reliability. Figure 1.7 shows quite a static trend in energy demand (especially in Europe), but also in this case the recent world's changes will surely imply some differences within the predicted values (all the values after 2019). The pandemic changes may question many researches made before 2019 with results projected after 2020, in particular at world wide scale.

Data centers performances and footprints are a sensible matter, so not every big company shares data about it. This means difficulties to gather consistent data from reliable sources and consequentially make truthful considerations at world level.

The backfire of this particular aspect, as explained in an interesting article written by Eric Masanet and Nuoa Lei (of the McCormick School of Engineering and Applied Science) [18], is the production of a considerable amount of white papers and reports with inaccurate considerations and a consequent wrong perception of the real situation. The article indeed highlights discrepancies between ambiguous and sometimes overestimated data and real ones. In addition make proper argumentation about energy efficiency, world pollution, footprint and renewables may be difficult because it is a very hot topic and the variables to keep in consideration are various.

To overcame data 's problem just mentioned, Green Peace developed a percentage value called the green energy index which also may fulfill the role of giving an alternative, different metrics. This index aim to measure effectively how much clean energy is used to provide the data center services [19]; It is calculated as average among the difference facilities and its computation is based on the percentage of different fuels used.

These results (written in a series of reports over different years [19] [20] [21] [22] [23]) was used to build the histogram in figure 1.8 . It shows the energy green index trend of a subset of companies:

The results show a different point of view on the green energy performance of a company. The green energy index indeed is a completely different metric with respect to PUE. Apple hits 100% but it does not declare its PUE making the energy index one useful term of comparison common with others. All the companies with high energy index generally have also optimal low PUE values (Microsft declares value under 1.1 and Yahoo and Facebook too).

Anyway, one metric does not exclude the other and it is better to consider them together in order to have a wider and more complete idea about data center footprint. The histograms show data before 2017 but for example Google data centers are now 100% powered by renewable source, Facebook declares to achieve

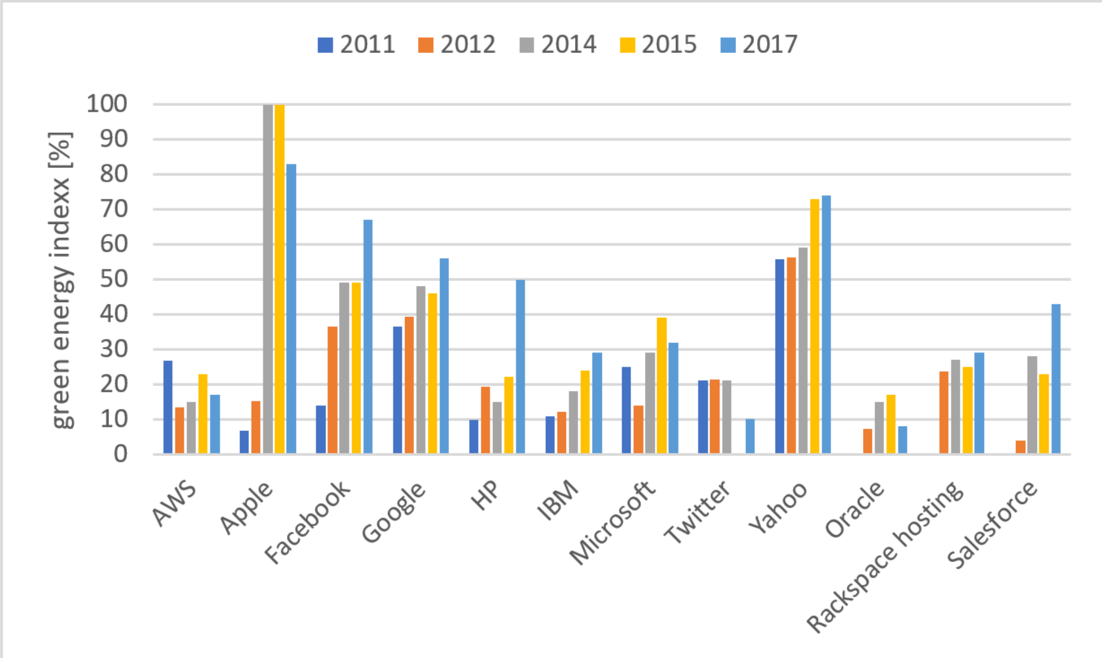

Figure 1.8: Green peace computed green energy index

the same goal by 2020 and Microsoft by 2025 [5].

Another important but less known footprint needed to be mentioned is the water one. Water is an aspect sometimes underrated especially talking about energy use. Water footprint is not as discussed and known as carbon footprint in particular for DC case. On this purpose, water is clearly largely used to produce energy in first place, and then it also used to cool data centers.

According to a research by National Renewable Energy Laboratory (nrel) with 2003 techonologies the solely amount of water evaporated for $1KWh$, used to be 7.6 liters [24]. "Considering the whole water used in the energy production process Tamis Younos [6] affirms this number reaches the value of $\frac{95L}{kWh}$."

[5] All the major IT companies are developing dedicated web pages which talk about their improvements in terms of ecological footprint and all the data mentioned about Microsoft Google and Facebook are retrieved from there:

- `https://sustainability.fb.com/`
- `https://sustainability.google/`
- `https://www.microsoft.com/en-us/corporate-responsibility/sustainability`

[6] Tamis Younos is the associate director of the Virginia Water Resources Research Center (in

Despite technology improvements in the last 17 years, these numbers remain considerable, especially if multiplied for the average data center electricity demand.

The water used for cooling purposes should be added as well and it represents a more difficult data to collect because linked to many aspects (first of which the cooling mechanism). So, in order to evaluate better the data center water footprint Green Grid developed in 2011 a new metric called Water Usage Effectiveness (WUE) which is the "Annual Water Usage" over the "IT Equipment Energy" expressed in $\left[\frac{L}{KWh}\right]$.WUE is not an a-dimensional metric like PUE and can be very useful to evaluate the water usage exceeding the power purpose.

Figure 1.9 is an open-source dashboard made by Facebook which shows the company efficiency. Facebook is one of the few companies which highlights also this pretty new metric. A lower value of WUE means a DC less water-intensive. the value of 0.35 should be multiplied for the DC power consumption as well to obtain the total water usage.

Indeed, figure 1.10 can help in understanding better the volume of water used by data centers.

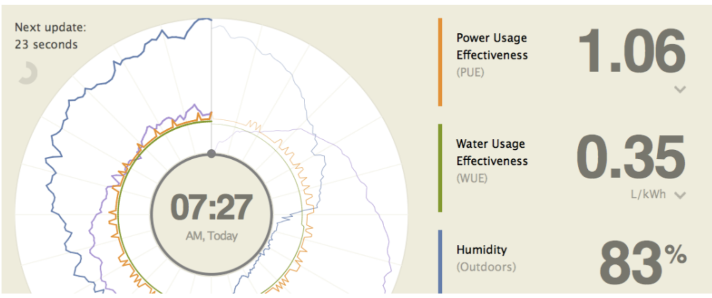

Figure 1.9: PUE/WUE Facebook dashboard [26]

This last figure is self-commentary, between the other subject choose to match the DC's water usage there is also a forest tree and that represents yet another consideration useful to understand data centers energy impact.

Blacksburg) and a professor of water resources at Virginia Tech. This statement is reported in an interesting article written in 2008 by Willie D. Jones for IEEE Spectrum [25]

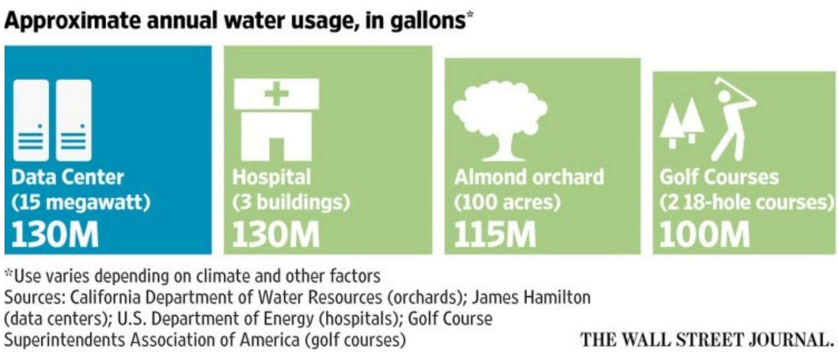

Figure 1.10: Water usage comparison [27]

1.3 Virtualization

Virtualization means the creation of an abstract, virtual resource of non physical entities like application, server, storage etc... Virtualization has a key role to enhance cloud computing possibilities and usability, providing several benefits e.g:

- improve scalability
- high avaliability options
- greater workload achievement
- improve flexibility and adaptability
- greater workload portability

Data centers are an important technology to achieve a more efficient and dynamical allocation of hardware resources. Virtualization implementation growth very fast in the previous years finding also new and more field of applicability as shown in Fig.1.11, it refers to a SpiceWorks research [28] which obtain thess results as consequence of a interview to 530 IT decision makers in 2019.

The higher bar in the figure is an undeniable value which shows the importance of virtual machine in server cases and a trend still increasing. The rest of the plot highlights in particular how virtualization technology is becoming an important reality for many other applications and higher percentage of future growth.

Server virtualization market was valued around 6 Billion dollars in 2019 and it is expected to surpass 9 billion dollar by 2026. Virtual machines are the main components of the virtualization technology able to move such a huge amount of money and they are basically a software implementation of real machine which permits the user to run code on it like on physical hardware.

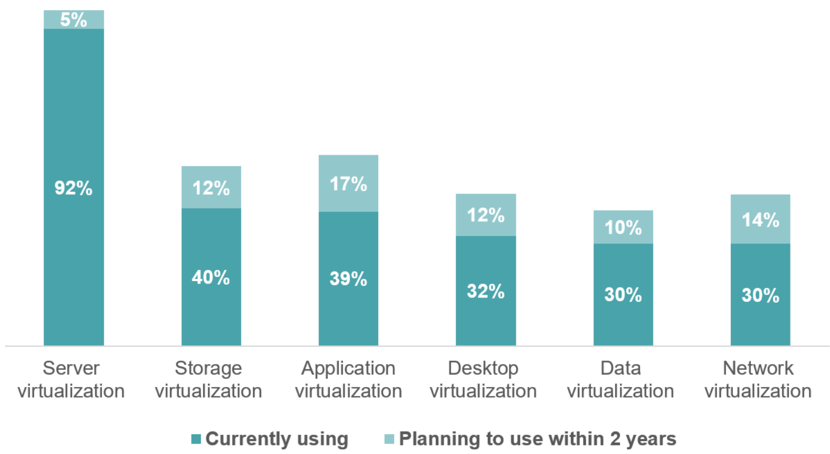

Figure 1.11: Business Adoption of Virtualization Technology [28]

Another important player in this technology is the hypervisor that have the role to supervise virtual machines and interface them with the hardware components which effectively make the computing and represent the effective computational power. The hypervisor itself represents the virtualization level with the role to to manage and distribute the hardware resources over the virtual machines.

There are several server virtualization architectures and in figure 1.12 is shown the simplest one (b) highlighting the difference with a non-virtualized (a)

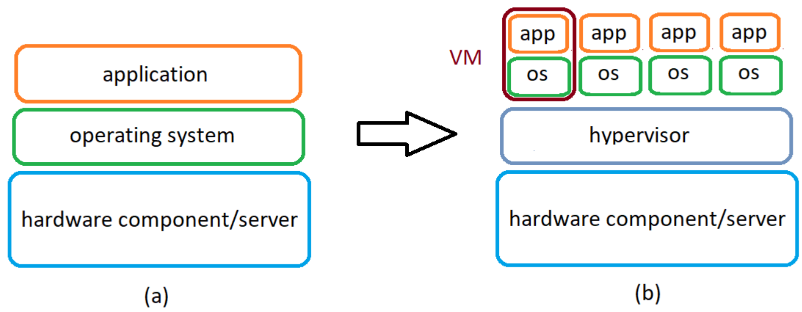

Figure 1.12: standard architecture(a) and virtualization implementation (b)

This figure is also useful to understand better part of the virtualization advantages mentioned before. In particular VMs are the main elements fulfilling the end user necessity in a cloud computing environment. In particular we can notice how hypervisor and virtual machine permit to exploit better the hardware capacity. At the same time end user have access to personalized environment easier to customize for fulfilling his/her own purposes.

The VM based architecture gives advantages also from energy and financial point of view because the main facility with hardware components that costs and are accountable for the majority of power consumption remain the same, but with the ability to service higher number of users which at the same time may pay less money for a more tailored service.

Virtual machine plays a key role in this thesis work and their characteristics will be explained better in the next chapters.

1.4 Aims and contributions

Starting point of this thesis is a partnership between University of Calabria and Polytechnics which developed and use a simulator environment named ”eco4cloud” in order to study the performances of virtual machines migrations. These simulator, thanks to values of 414 real VM collected by a Telecom data center, compute load migration performances among 4 different hypothetical DC around the world with various energy characteristics over the day.

Previous study outlined a behaviour of the migration related to energy parameter changes over time and consider also influence of renewables and batteries capacity. This thesis focuses mainly on implementing this simulator to study further migration behaviour and in particular the influence of the VM characteristics over the migration performances.

Focusing better on virtual machine object itself and its main features is the main purpose of this project. The features chosen are CPU, RAM and disk usage which may represent the core of each virtual machine. In order to evaluate them was developed a system of thresholds to constrain the migration and observe eventual variation of performances in function of these three characteristics.

These thresholds were used to perform e selective migration based on features and CPU is the one which influences more energy cost of the DC (accordingly with the simulator cost computational function). After studying the relation between CPU and other features, different simulations were made to observe if the outputs match the expectation for the different characteristics.

Two types of analyses were made: the first one considers the aggregated results and the total output of the simulator while the other one is more punctual and observe the migration performances at each iteration.

Other considerations, less related to the VM feature themselves, were also put under observation. In particular, it was contemplated the relation between the amount of load migrated and the energy costs' reduction and a comparison between performance of migration and local consolidation which is as well, a mechanism to increase energy efficiency based on host optimization.

Chapter 2

State of the art

The previous section is an overview to Data Center energy impact with some power consumption consideration. Now can it can be useful consider the main solution adopted to reduce or optimize DCs energy demand and data centers footprint.

2.1 Cooling mechanism

Among different approaches to reduce the energy demand of DC, one of the most important is the reduction of energy used for cooling purposes. This also implies the reduction of PUE and consequentially an higher efficiency.

Optimal temperature and humidity conditions improve also hardware performances because each electronic component has its own optimal working region and performs better if maintained inside of it.

Therefore, a proper cooling strategy means accomplishing the same computing task faster and better with a corresponding monetary and energy saving.

So cooling mechanisms are fundamental both for energy efficiency and hardware functionalities and for that reason many studies were done on this subject leading to different technologies and approaches. Despite slight minor differences in the implementations and the structure design, in general the cooling mechanism can be divided in 2 subgroups:

- air based;

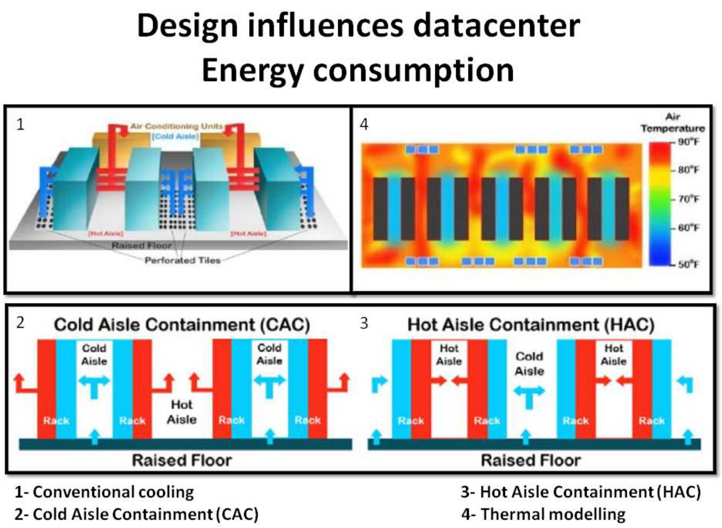

Figure 2.1: Air cooling mechanism [29]

- fluid based;

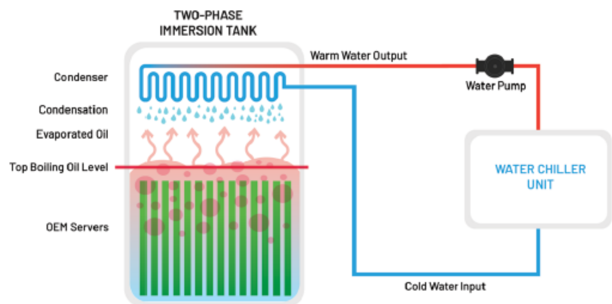

Figure 2.2: Advanced fluid cooling mechanism [30]

Figure 2.1 and 2.2 show two examples of cooling mechanisms. Each of them has different advantages and disadvantages, and the structural cost varies as well.

One important aspect worth considering is the DCs placing. Data center can be placed in many different parts of the world for many different reasons and choosing of the cooling mechanism should also take into account external environment.

Indeed, there is a lot of variables to be considered in order to achieve a reasonable efficiency.Cold Isle Containment(CAC) and Hot Isle Containment(HAC)[1] are two air based cooling mechanism and for example the first one has the advantage over the second to keep the surrounding cooler, keeping a better working environment for the DC building's employers. Although, this feature may represents a disadvantage at high latitude regions where the cold air can be reused to warm DC facility. The humidity and the water usage can be discriminant factors as well. It is important to consider data center as a whole, a building or a structure in general related to many other different activities with its own maintenance needs.

Figure 2.3: Microsoft DC Natick project [31]

Project Natick is an interesting Microsoft attempt with a peculiar choice for the DC location. In fact Microsoft placed a DC on the bottom of the sea for 2 years, choosing to trade data center structure accessibility (at least from physical handling point of view), with the possibility to make the most of the environmental characteristics.
The data center (in figure 2.3) was retrieved from its location on July 2020 and now is being analysed. The first data declared on the project's official page ([31]) tells about a data center 8 times more reliable and very prone to energy efficiency.
An underwater DC has clearly a series of constraints but this project represents an important proof of concept and contributes to demonstrate the importance of a proper DC placing and cooling.

[1]CAC and HAC are two of the air based cooling mechanism shown in figure 2.1 (respectively 2 and 3). As it is possible to see in the figure HAC traps the hot exhausted air which will be drained away while CAC approach traps the cold air while the hot air remains external and will be managed by ventilation system. So in one case the external environment results way cooler

2.2 Beyond PUE

Reducing PUE is very important to reduce energy waste but also an ideal value of 1 does not imply any reduction for the effective energy demand for computational purpose and so, the DC power consumption still remains not negligible.

We have already seen in the previous chapter the importance of some other aspects (e.g. green energy index, water footprint, energy and fuel considerations). PUE metric alone, cannot be the only metric to define the energy efficiency and in particular also an ideal PUE does not exclude the possibility of further improvements from power consumption point of view.

So, some considerations should be done about the possibility of reducing the overall energy demand with proper management policies. Next paragraphs will focus on different approaches for improving energy efficiency, based on smart load management. The final goal remains to reduce power consumption or CO_2 emission or generally increase efficiency, but the implementations are different and often less hardware dependent.

2.2.1 Server usage optimization

In 2007 Luiz André Barroso and Urs Hölzle published an important research regarding ratio between energy consumed and sever utilization [32]. It is common practice to spread data and resource among different server in order to improve system reliability and content accessibility, "As a result, all servers must be available, even during low-load periods. In addition, networked servers frequently perform many small background tasks that make it impossible for them to enter a sleep state" (citation the report [32]).

The key issue is the servers optimization at high loads without proper consideration for lower activity cases. Clearly a data center with all servers active 24/7 performs better because all the resources are always completely available but this condition has a huge drawback from the point of view of energy efficiency.

The ideal solution proposed in the paper is to change toward a configuration based on a proportionality between servers utilization and energy demand. This approach implies to find a trade-off where, slightly better performances in low load cases, are exchanged for higher efficiency. It means to introduce a proper sleep and power off policy together with an accurate load distribution between different hosts of the same data center.

Figure 2.4 shows that servers consume almost 50% of the peak energy demand also in virtually idle case, the consequence is a considerable waste of energy for the only advantage of maintain a server more reactive to sudden requests. On contrary figure 2.5 represents the behaviour of a more energy-proportional situation. The two red efficiency curves are very different and highlights the huge advantages, in terms of power consumption, of the second configuration in respect to the first one.

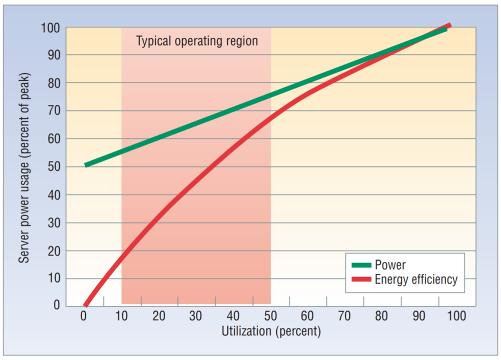

Figure 2.4: Power and efficiency over utilization: non proportional case [32]

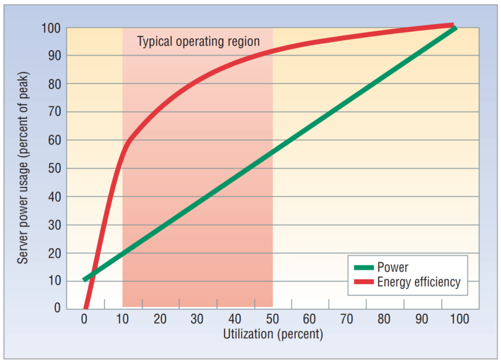

Figure 2.5: Power and efficiency over utilization: proportional case [32]

As expected the two figures show substantial differences in the low utilization region because the green line has indeed a starting point way lower in the second one. Figures show a typical operating region below 50%, because usually it is better to avoid the server saturation and prevent overheating and an excessive strain on electronic components.

The solution proposed implies to increase load on a set of servers, to save energy with proper sleep or power off strategies on some other ones. Also in the new resources allocation phase, it may be more convenient to maintain some of the possible receiving servers off, increasing load on others and consequentially their utilization regions, so as mentioned before a trade-off between this two condition should be found.

New hardware components can help in this task due to higher reliability and resilience but in general all technology improved considerably in the last decades considering also software development and new data center management policies. So, on august 2017 Rathijit Sen and David A. Wood published an article [33] aiming at updating the previous considerations of Barroso and Hölzle.

Figure 2.6 reconsider and extend plot 2.4. The green line EP [2] represents an ideal behaviour also in the zero load case while the red line (dynamic EP) refers to a more realistic scenario where a low energy demand is present, also near 0% load. The second one is nearest to old behaviour of EP (green line, figure 2.5) and consider minimal power demand due to background activities.

In this second report considerations about the performances of new servers were added, which reach peaks levels at different loads as consequence of the topologies

[2]EP stands for energy proportional

and configurations used. The dotted light blue lines instead, represent the behaviour of a system with different configurations aiming at better explaining the modern servers' functionalities.

Sen and Wood consider the innovations on modern DC and their possibility to adopt a more re-configurable system. The dotted light blue lines represent the nonlinear behaviours of different configurations and the black dotted line is the consequential pareto frontier, which intersected the EP line and where this happens the super proportional threshold is crossed.

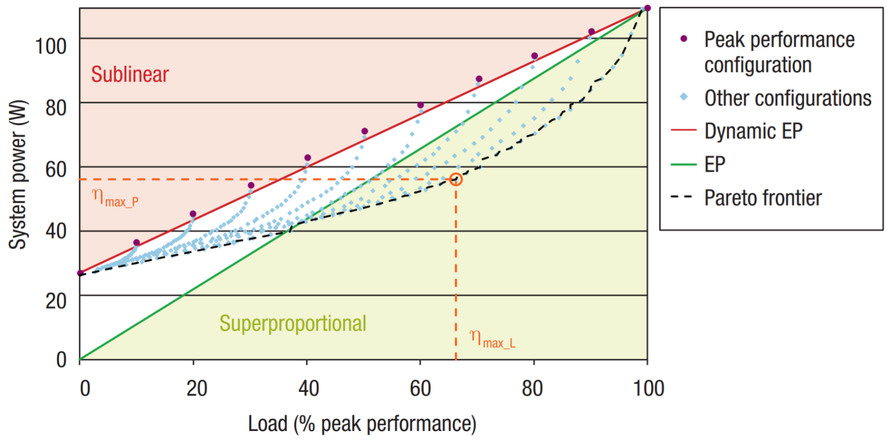

Figure 2.6: System Power at various loads with with multiple configurations [33]

The old definition of energy proportional computing asserts the max efficiency is reached in a 100% load case and this value represents an upper bound (all the saturation considerations made before are not considered, it is only an ideal upper bound). The possibility of a configurable system gives the opportunity to go beyond this limit and the key idea for the redefinition of an energy proportional computing is to achieves η_{max} for all loads.

Crossing the super proportional line means the hypothetical upper bound just mentioned is exceeded, this upper bound results from an ideal proportionality between load and power consumption, and the green area beyond it represents a situation where the system is beyond this proportionality.

Indeed, in the EP ideal case, the ratio between load and power is 1 for all loads while the pareto frontier shows better performance after the intersection (loads higher than 40%).

Also in this case, it could be useful to analyse the efficiency curve and figure 2.7 compares it with the pareto frontier line, highlighting η_{max} which results 29% higher than the upper bound imposed from the previous definition of energy proportional computing.

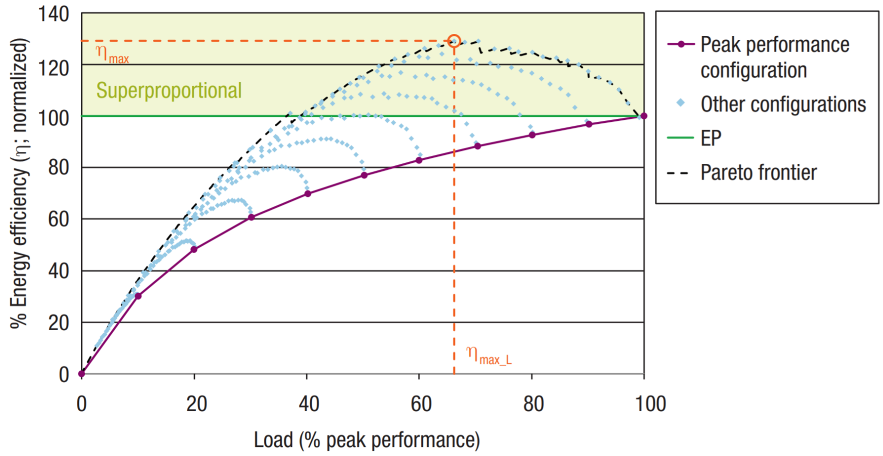

Figure 2.7: System efficiency at various loads with multiple configurations [33]

So, in the end the key idea is to exploit the re-configurable potentials of new data center architectures in order to find an optimal work region to maximise performances with the current technology development degree, indeed this approach permits to achieve higher levels of energy efficiency.

2.2.2 Data Center Demand and Response

Demand Side Management (DSM) is a mechanism where the consumers' demand is changed or influenced to match better the energy supply. Demand and response (DR) is a class of DSM where the consumer attain a series of benefits from the rescheduling of its energy profile. The key idea in DSM and more specific in DR is to move partially consumers demand accordingly to power distribution provider and the grid necessities. According to international energy agency (iea) demand response programs are becoming a valid solution in different world area and for example "In Italy, a total 280 MW of capacity were commissioned by the system operator across the country, while in Ireland 415 MW of demand response capacity was awarded in a T-4 (four-year-ahead) auction." This is a textual citation form iea report [34] and it is useful to contextualize and understand better the magnitude of energy implied in DR policies, a huge amount of energy dedicated to DR policies is a good indicator of this approach possibilities.

Figure 2.8 is a useful summary diagram (brought by an interesting overview on DR published on IEEE Explore in 2007 [35]) which shows the demand response advantages from different points of view.

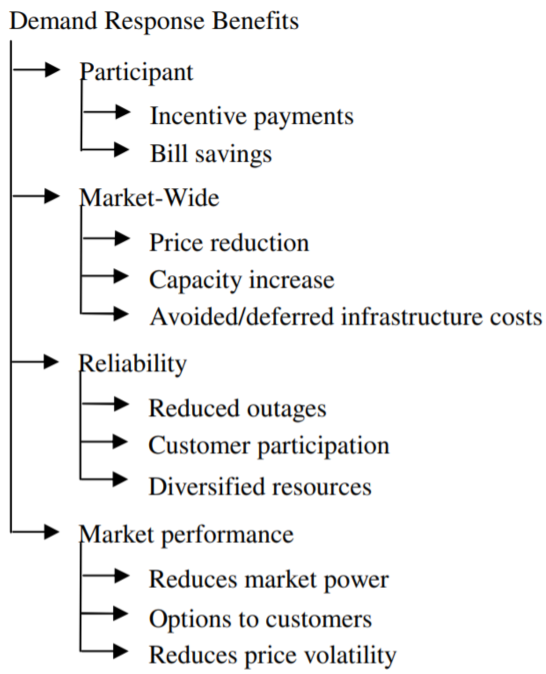

Figure 2.8: Classification of DR benefits [35]

First of all, DR helps in the important achievement of reducing burden on energy grid, this also means higher reliability and flexibility, a reduced strain on power supply infrastructure and the possibility to reduce overloads and handle better emergency situations.

Moreover, DR helps to match better renewables production and account for their intermittent behaviour especially considering the possibility to introduce energy accumulation solutions like batteries. In fact, demand response is not only about reducing energy peaks but more precisely this approach aims to move them accordingly with the grid possibilities.

So, if the power grid does not include intermittent power source in general the DR policy can still be useful to reduce energy peaks while, on the other hand, in presence of renewable sources, may be advantageous to match them properly.

The final considered scenario should be a smart grid infrastructure with renewable power sources and batteries combined with a properly tailored demand response policy. Clearly, one of the main difficulties of achieving such scenario is the intrinsic necessity of DR approach to find a compromise with end consumer. Indeed demand response gives back remarkable results if the load considered is large but for a consumer's subset, this approach may impose conditions too strict to be adopted, e.g. average workers used the electricity at home only in early morning and in the evening after work, at the same time a company may presents higher peaks in the morning and after lunch breaks.

So, DR performs better if adopted by a large number of consumers which cooperate to attain a series of benefits like money saving. There are also situations where different macro categories of consumers can potentially match their necessity demand each other, consequence of using energy in different moments of the day (e.g. the just mentioned huge set of employees and the facility in which they work). However a proper management is needed and more variables to be considered may increase the its complexity, also considering the issues related to have a flexible and not homogeneous consumers set which relies on the same energy provider finding an agreement.

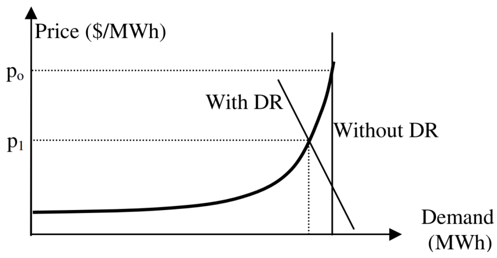

Figure 2.9: Simplified Demand Response effect [35]

In figure 2.9 there is a qualitative simplified explanations of DR effects, the key feature of this approach is to impose a negative slope at high demand and so, the higher the demand, the higher are the benefits obtained. Despite many different consumers may constitute this high load and high demand requested, they should also be willing to accepts limitations to their power usage in order to satisfy DR requirements.

A data center facility has high power consumption but at the same time counts as one single entity, so it is easier to manage than many different consumers with an equal summed total demand. Indeed, from many points of view, DCs may be the ideal consumers in this case and they have many characteristics which match perfectly DR polices:

- very large load and power demand but single entity
- active 24/7
- scalability
- flexibility
- necessity to reduce energy cost
- renewables and battery compliant
- different facility, one owner

Indeed, DC main features make it more inclined to load shifts and considering a data centers set of the same company instead of one single facility, benefits obtained may increase even more due to DR potentialities, especially if merged with important investments in renewables.

To evaluate demand response benefits in data centers case, it could be useful to cite for example a publication of a group of researchers from Chicago University, which investigate the potential saving in an energy bill for a small to medium data center situations [36]. The final conclusions considering a high performance computing scenario is, citing textually: "curtailing 1MW, 8 times per year, will lead to an annual electricity savings of approximately $100K. For our facility, this corresponds to a total cost savings of roughly 7% annually." This represents quite a remarkable result especially considering the study case and projecting it to higher load cases.

To conclude this argumentation about demand response can be useful to look at figures 2.10 and 2.11, by an article on official google 's blog page written by "Ana Radovanovic" [37].

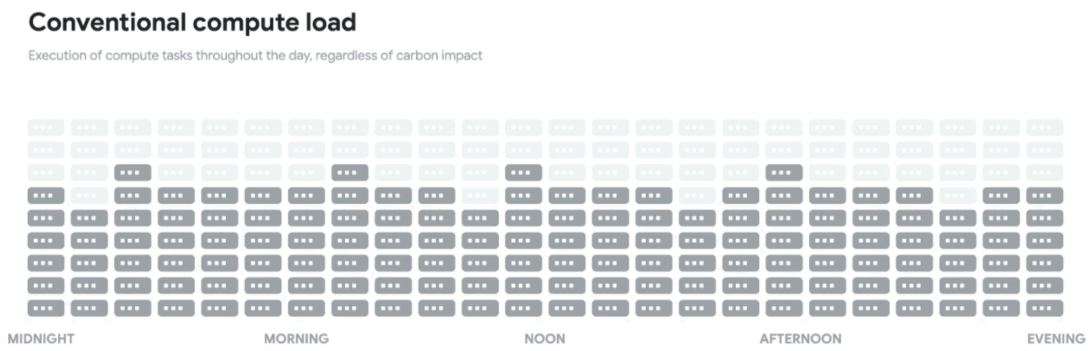

Figure 2.10: Conventional compute load

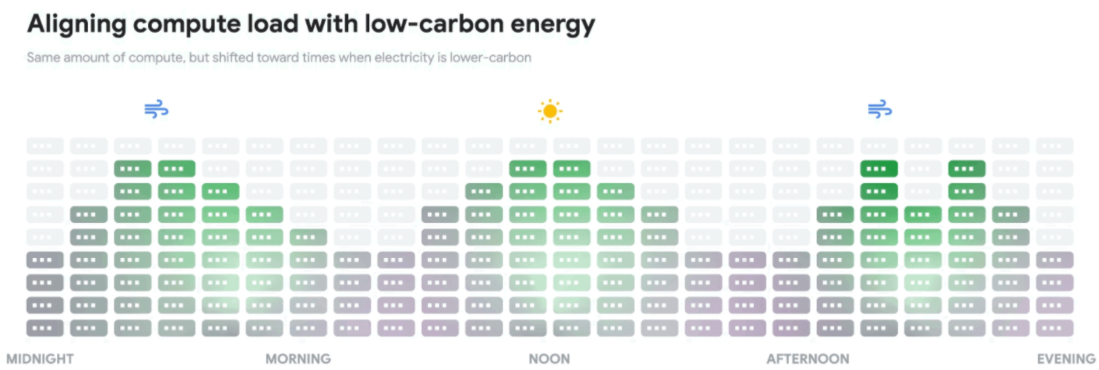

Figure 2.11: Aligning compute load with low-carbon energy

Figures 2.10 and 2.11 represents one of Google solutions to reduce CO_2 emission. The idea is to move computing load to match energy production from low carbon emission sources. This approach is clearly similar to demand response (and helps to understand it better) but it is CO_2 based and consider only one big company. However it still represents a load shift approach particularly interesting for its green oriented characteristics where there are remarkable results due to the high magnitude of Google 's data centers load.

2.2.3 Load management and virtual machine migration

Load management is a general concept, that was already mentioned a lot in the previous sections and indicates the handle of data centers load for many different purposes but generally with the aim of accomplishing some kind of optimization.

In chapter one is mentioned the importance of virtualization in modern data centers architectures where, in fact, virtual machines have a main role and represents the principal load for cloud DC computational resources.

Demand response key idea is to perform a temporal load shift in order to obtains benefit, but could be also interesting considering a spatial load shift for similar reasons. There are many factors which may influence DC performances and energy efficiency like DC placing, environmental conditions, renewables production, DC characteristics, energy cost, scope and usage.

All these factors (which clearly differ for each data center around the world) contribute to give a sort of characterization that may be used to outline a DC performance and convenience profile. One of the main approaches to load management is to balance this load among different servers and data centers but, as we have seen so far, there are many alternatives which outperform this method at least from the energy efficiency point of view. VM migration is another policy which may provide benefits by migrating virtual machine (and consequentially DC load) to the most convenient DC among the series at disposal.

Clearly, these observations should be contextualized considering the number and type of DCs, the type and magnitude of benefits desired and the constraints imposed:

Obtainable benefits:

- reducing energy expenditure;
- reducing carbon emission;
- increasing energy efficiency;
- exploiting of renewables;
- improving system flexibility;
- emergency resilience

Drawbacks and constraints:

- increasing distance between resources and end user;
- cost of migration;
- possible latency;
- increase data center workload at certain timestamp;

Figures 2.12, 2.13 and 2.14 (taken from a Google sustainability report [38]) are three heat maps representing the daily percentage of DC energy demand matching a free CO_2 emission power generations.

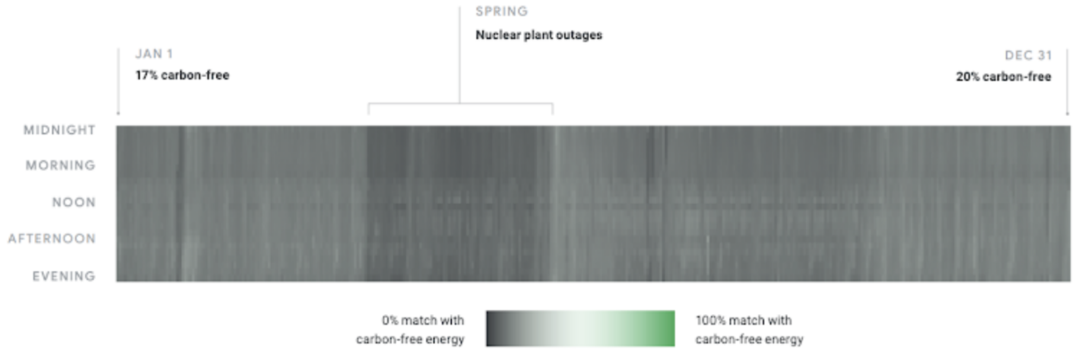

Figure 2.12: Google DC in Changhua, Taiwan

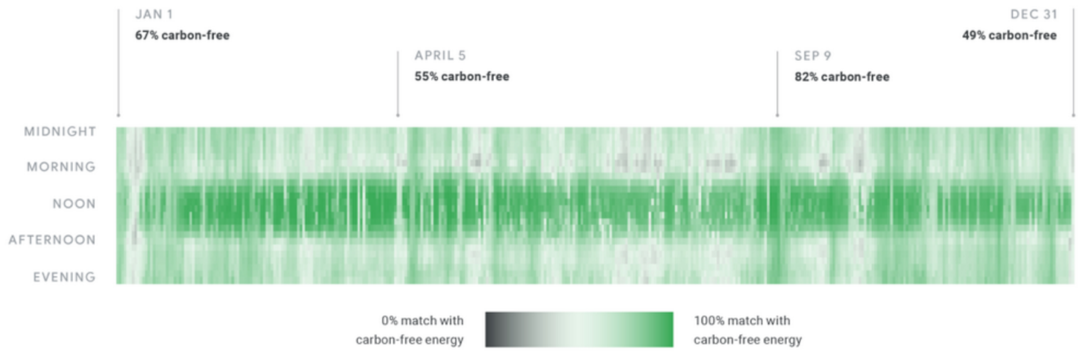

Figure 2.13: Google DC sin North Carolina

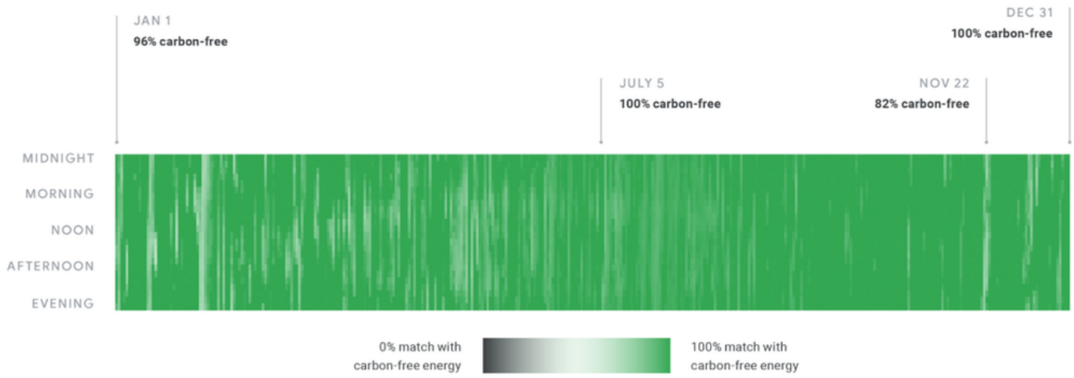

Figure 2.14: Google DC in Finland

There are huge differences among this three figures consequence of three different locations with different environmental and economic scenarios, in particular DC in Taiwan relies more on fossil power sources while North Carolina 's one use a lot solar energy and finally Finland, with the best performance (from the carbon emission point of view), are mostly powered by wind turbines. These power sources are peculiar of the related regions and may not be usable in the other places (like solar in Finland), due to environmental and geographical constraints. Moreover each location should be evaluated also form the geopolitical and economical point of view since each country have different energy cost and taxes and regulatory policies (Finland gasoline cost is on average almost double whit respect to Taiwan, without considering the different taxation).

This thesis uses a previous research [39] as starting point. The related conclusions will be treated again in the next chapters but could be useful, now in the state of the art, mentions it and look at figure 2.15 that is one of the results obtained. In particular this image shows the effect on migration versus the variation of photo voltaic panel size (PV) for the case studies considered. As we can notice with a well sized solar panel the migration benefit becomes not negligible aslo without any particular energy saving oriented policy.

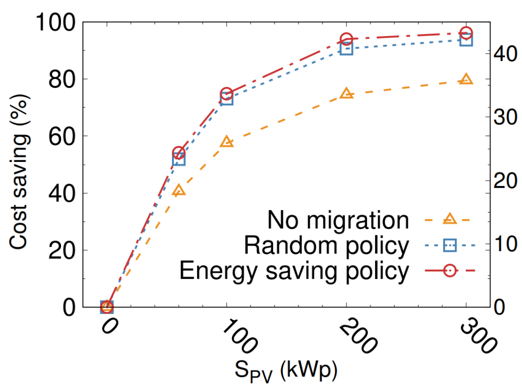

Figure 2.15: Cost-saving versus PV panel size for different migration policies and battery=200kWh [39]

One last consideration about VM migrations worth to be done is relative to real time migration. Indeed, with current technologies is possible to move virtual machines without interrupting workflows. This is possible thanks to a huge implementation of hypervisors functionalities which are actually able to reallocate and copy, all the VM virtual resources amoving the executions itself only in the end, without necessarily abort it. This achievement makes easier to adopt load migration policies because cancel the huge downside effect of interrupting execution of activities on the migrated virtual machines.

Chapter 3

eco4cloud simulator

As mentioned, this thesis work uses as starting point a previous research which study load migration. In this section it is explained the case study and the simulator environment and its functionalities.

3.1 Case study and Database

First of all can be useful to have an overview of the case study and the database which is the base of the analysis argued in this paper. Load migration aims to exploit the difference in geographical location of group of different DCs and figure 3.1 is a graphical representation of the case study considered.

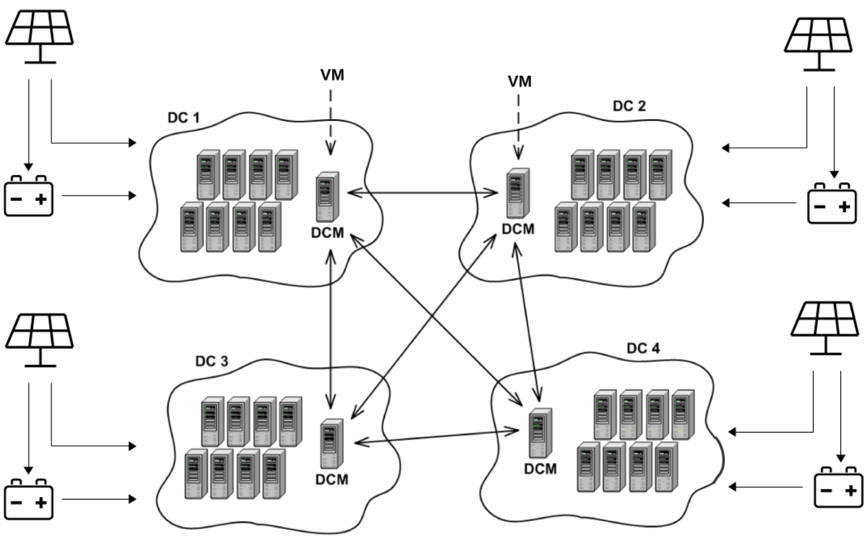

Figure 3.1: Case study [39]

There are four data centers located among different geographical positions (different coordinates). Each of them is provided with PV panels which provide a certain amount of renewables in different time of the day and batteries which store the unused energy to supply it when needed.

Each data center may have different energy characteristics related to locations. In the state of the art it has been already argued about what a different geographical location may imply. The four data centers features simulate this scenario with different values of energy prices, PUE, renewables production, carbon emission. Moreover also the variation of this parameters is not synced to consider the different time zones.

Migration aims to move virtual machines among different data centers to exploit as much as possible these differences. Indeed these variations implies that each data center may result more convenient in different timestamps The final aim of migration is to periodically move load to this most convenient data center to have related benefit (e.g. reduction of carbon emissions or reduction of energy expenditure).

Virtual machines are initially distributed over the four data centers and after that, a periodical migration was performed. DCM is the data center manager which represent the administrator level of the data center facility which receives the VM migrated and reallocates them among the hosts.

The VM data used refers to real virtual machines of a Telecom data center. It has been created a database with different tables containing the information about each virtual machine and their corresponding servers. In particular the tables of the database mainly used are:

- "server": this table contains all the data about the 28 server of the DC. These data are composed by the number of core assigned, CPU and RAM and the computational resources assigned in general to the servers.

- "v_m" and "v_m_performances": In this two tables are stored all the virtual machines info and the server id to which they belongs to. There are a total of 496 virtual machines but only a subset of 414 are actually active (they have RAM and CPU usage equal to zero). CPU and RAM are expressed in percentage and are multiplied for the resources of the corresponding host to compute the effective usage in MHZ and MB.

The four DC instead are fictitious entities created for the simulation purposes. They have been configured to represents four hypothetical data centers and are then populated with the virtual machines and hosts real data attained from database queries. The simulator uses the VMs described as sample pool to create the used virtual machines.

3.2 Simulator environment

The simulator environment is developed in Java language and it is based on the utilization of three important different objects: DC, host, and VM. All the simulator considers a timestamp as time base unit and one timestamps corresponds to one hour, so all the operation happening between two different timestamp are considered ideal with zero execution time.

Although the database is populated by only the virtual machines of one data center, the simulator creates 4 instances of the same object DC and treat them separately. In particular there are various parameters which changes for each data center and can be useful to cite some of them:

- some of the variables are set differently for each data center and does not change over timestamps:
  - PUE;
  - carbon emission (CE) expressed in grams;
- there are other variables which change over time, in particular the configurations has been set over 24 timestamp (one day), so for longer simulation the pattern is repeated but clearly the system has different starting point and the output may vary as well over different days.

  These time dependent variables are set to simulate the behaviour and condition of different data centers around the world:
  - Energy prices: differ between data centers and between different timestamps. Also the average price is not the same for all DCs in order to simulate the variation of price for energy sources in different location.
  - PV production: this variable varies as well and it is related to battery capacity and PV dimension which on contrary are fixed for each DC.
  - Arrival rate μ: it indicates the number of virtual machines incoming for each timestamp and the mean changes over timestamps but not among different DCs. This value is then multiplied for variable named "ls_factor" which does not change during the simulation and over DC.

Without going into too much details there are a bunch of functions and variables utilized from the simulator environment worth to be explained because are fundamental to understand its behaviour and the results obtained:

- "score": it is computed whenever in the simulator there is a necessity to outline the most convenient data center for a specific timestamp. The score is computed for each data center as function of the estimated load, PUE, energy costs and the renewables contribution.

So, score computation is affected only by the DC setting for the timestamp and by an expected load whic is the same for all DCs. In fact it does not consider in any way the real load effectively present on data centers.

- "dc.costi": in the code this variable stores the amount of energy costs for each DC object without renewables and batteries consideration. This variable has two contributions:
  - the first one is fixed for the timestamps and regard the DC and host performances, so similarly to "score", it only depends from the parameters of the DC;
  - the other one sums all the contributions of each virtual machine.Considering ,in particular, the CPU in MHz of each VM. This parameter is then scaled by considering features of DC and host to which the VM is assigned of.
- "dc.energyRen": this variable keeps track of the variation in renewables at disposal for each data center, considering also the batteries which are charged only if the PV production exceed the consumption.
- "dc.updateTotalEnergyCost(timestamp)": this important function updates, at each timestamp, the data centers condition by updating the variables regarding costs, load, renewables etc.. The drawback of this function is to have a cumulative approach. For example, if this function is called one times after another without any other action in the middle, the outputs will change anyway.

  This function behaviour is being created to be used only once for each timestamp to update the energy system situation before the next iteration.

The flow chart in Fig. 3.2 is a simplification of the simulator steps which may help to understand better the overall behaviour.

The diagram shows the main steps of the simulation:

1. Initial setting: the initial settings are mainly composed by the creation of the 4 DCs objects and the from external files import of all the simulation parameters, including the ones mentioned before.
2. Initial assignment: the 4 DCs created are populated with the virtual machines taken from the database. It is important to understand that the 414 active virtual machines of the database are "cloned" and used among the system as different entities in each DC so the simulator is not bounded to consider only 414 of them.

   In this step for each virtual machine it is attained a list composed by the 4 data centers ordered by the most to less convenient one. The aim is to obtain

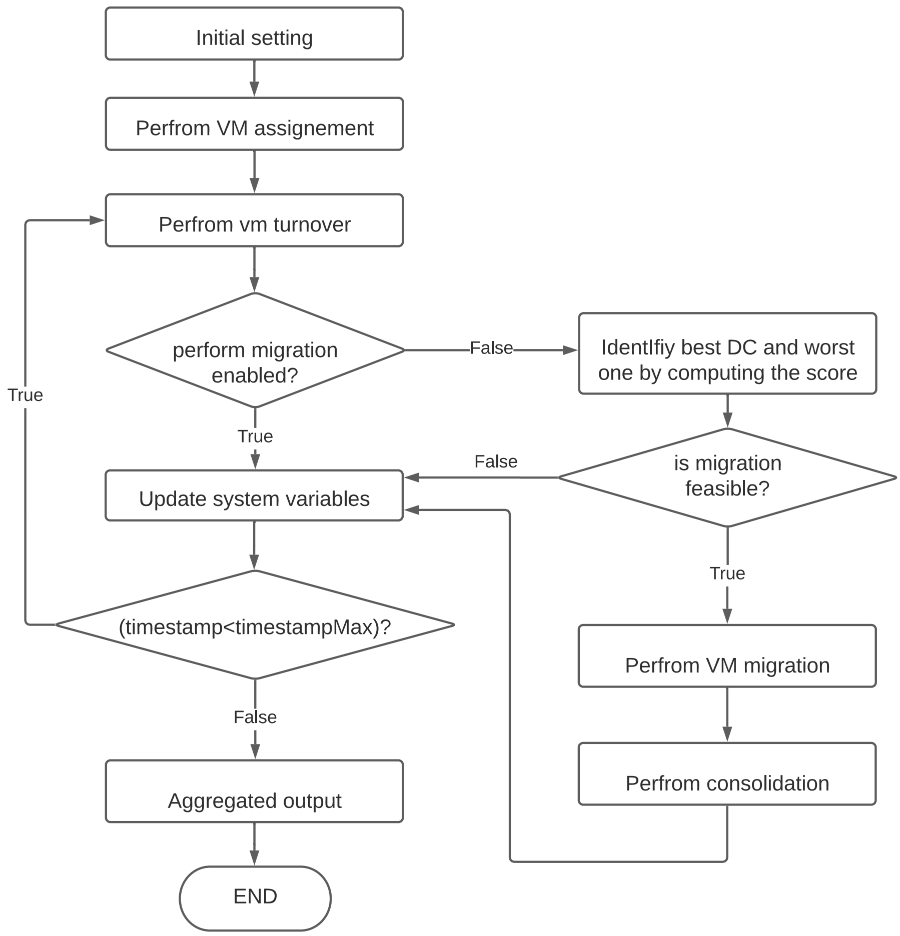

Figure 3.2: Simulator main steps

an assignment 's order by computing the already mentioned "score" which does not depend from the DC real load. So the score is evaluated after each single VM assignation.

3. VM turnover: new VM are introduced in the system while the ones already served are removed. It is a very important step which gives flexibility to the system and behave in according with the arrival and the service rate of each timestamp.

4. VM migration: if enabled the virtual machine migration is performed. Also in this case it is computed the score to determines the best and worst DC from the point of view of performances and energy costs.

   Unlike the assignment step, in this case the score is computed only once at the function beginning. Consequentially,worst and best DCs remains the same for all the migrations performed for the timestamp. Change this policy for the migration is not such a simple task, especially in a real environment. The total number of VMs migrated can also reach the tens of miles and a score computation after each single VM migration may be quite expensive in term of computational resources.

After having outlined the source and destination DCs, there is a control on the migration feasibility where three conditions are checked:

- bandwidth: it must be enough to perform migration;
- host saturation: if the receiving DC has saturated it cannot take any new VM;
- delta: this check is done to avoid unnecessary migrations and it indeed permits to avoid it when the scores computed between worst and best DC are too similar.

5. Perform local consolidation": in the state of the art it was already mentioned the importance of a proper optimization policy among different servers of the same DC. The objective of the consolidation is exactly to optimize the load distribution among hosts by performing a sort of internal migration between them.

6. Update system variables: in this step the function "updateTotalEnergyCosts" is used to update the system counters and variables before a new iteration begins.

7. output: the simulator environment gives back two types of results, one aggregated and one for each timestamp. In the first one the output is one row for all the simulation with a final results which is a combination of the contributions of the various timestamp. The other result instead is more punctual and considers only the performances of each timestamp.

3.3 Parameters choice

Having understood the simulator environment main functionalities, is now possible to talk about the parameters chosen for this thesis analyses. As seen there are a lot of variables that must be considered to understand better and contextualize properly the the results obtained from the simulator environment and could be useful to cite again some of them:

- batteries capacity;
- PV dimension;

- ls_factor;
- lifetime;
- μ service rate;
- λ arrival rate;
- bandwidth;

Clearly the majority of these variables are set accordingly with the previous results obtained, in order to exploit all the considerations and analyses already made and to provide a sort of continuity to the research.

Figures 2.15 shows percentage of cost saving over PV panel dimensions with 200KWh as battery capacity and it is possible to notice the migration performances became considerable over 100 kWp. For these reasons, the PV dimension used was set to 150 kWp while the battery capacity considered is 200 kWh.

Bandwidth represents an important constraint for the migration and in a first moment it was maintained the original setting of 2Gbps. However, this thesis focuses on virtual machine characterization during migration and for this purpose a constraint of this type can be neglected with proper contour considerations.

Indeed, the effects of migration are more evident with a huge amount of VM migrated but at the same time it is important to avoid system saturation in order to maintain the system dynamic and prone to migration.

λ is the arrival rate and it is set in an external configuration file where it varies over timestamps. "ls_factor" is a multiplicative factor for lambda which have the role to scale it, increasing the arrival rate. μ is function of the data centers load and the lifetime and together with λ is responsible for the amount of saturation of the system. In particular it is equal to the ratio between total number of VM present in the system and the lifetime.

So, some analyses focusing on migration are needed to provide proper values of "ls_factor" and lifetime. The final goal is to set them to create a system as prone to migration and as reactive as possible. So, "ls_factor" and lifetime are two critical variables to be set and to choose them properly it was considered the variation of VM migrated in relation to a variation of these two variables.

Figure 3.3 represents the output obtained after performing simulations with different lifetime values over 168 timestamps. The number of VM migrated drops after just one iteration while the differences for lifetimes greater than ten have lower magnitude. As expected, longer lifetimes imply more VM in the system also due to the decreasing of the service rate. So the system incurs in a saturated situation where there are less evident migration effect.

From another point of view a lifetime too low is unfeasible. Considering for example a value of 1 means that the virtual machine stay in the system (on average) an amount of time equal to the timestamp unit itself. In this condition almost zero virtual machines last enough to undergo a timestamp variation. The plot justifies and proves the correct choice of a values of lifetime equal to five, made from the precedent users of this simulation environment.

So, to maintain an environment prone to migration, a value of lifetime equal to 5 has been chosen for all the further considerations. On average this means after five timestamp a VM leaves the system or, from another point of view, at each timestamp $\frac{1}{5}$ of the total amount of virtual machines leaves the system.

Considerations about "ls_factor" are similar and aim to find an optimal value that also in this case probably cannot be too high in order to avoid system saturation. After fixating "$lifetime = 5$" the same type of simulation has been performed over ls_factor and figure 3.4 shows the number of VM migrated over a 168 timestamps simulation iterated for different "ls_factor" values.

In this case the differences are less dramatics, but it can be noticed a decreasing trend for value higher than ten due to the expected saturation of the system. Also in this case a trade-off must be considered because values of ls_factor too low may implies a system too empty and consequentially less informative. So, in this case the values chosen is 4 because shows the highest number of VM migrated and may represents a good compromise.

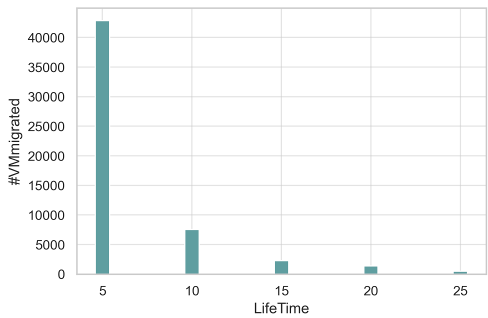

Figure 3.3: VMs migrated over lifetime: one week simulation

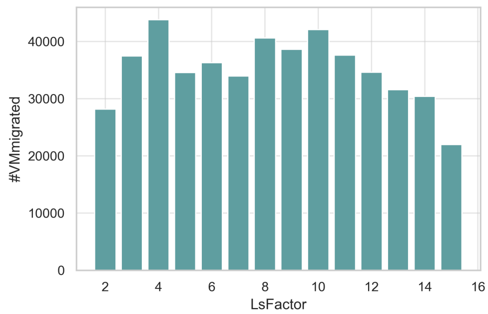

Figure 3.4: VMs migrated over ls_factor: one week simulation

Chapter 4

Virtual machine characterization

Section 3 is fundamental to understand better the following considerations. It explains "eco4cloud" functionalities and its main functions and purposes. This overview of the simulator environment is necessary because it represents a starting point for this thesis work and all the further consideration are based on implementation made over it.

4.1 VM characteristics

"eco4cloud" simulator already gives back a lot of information about VM migration performances and implications, focusing on understanding the behaviour in relation to batteries and PV panels dimension. Now, the main goal is to make further considerations, on the migration performances in general, and then analyse the impact of a proper VM characterization.

In particular three different features were chosen:

- CPU: this feature (expressed in percentage in the "v_m" table) is very important to define workload's contribution. Indeed, (as seen in the previous section) the energy cost considered in the simulator environment depends also from the amount of CPU usage in MHz sum of different VMs contributions.
- RAM: it is another important feature indicating (in % as well) VMs performances which strongly characterize a virtual machine and it may also represent an important load to be migrated because it consumes both resources and space allocation.
- Disk utilization: this last feature considered does not give directly information

on virtual machine performances but when a virtual machine is migrated its storage has to be migrated as well. So it represents an important indicator of the volume of data in MB to be migrated together with the related VM.

VM are in total 496 but only 414 of them are effectively active and are considered in our analysis. Table 4.1 reports maximum, minimum and mean of each features while Figures 4.2, 4.3 and 4.4 gives an idea of the VMs population in relation to this three characteristics chosen.

Features	Min	Max	Mean
CPU[%]	0.66	100	9.35949
RAM[%]	1.76	99.79	56.0425
Disk [MB]	10000	1.11256e+06	87601

Figure 4.1: features main indicators

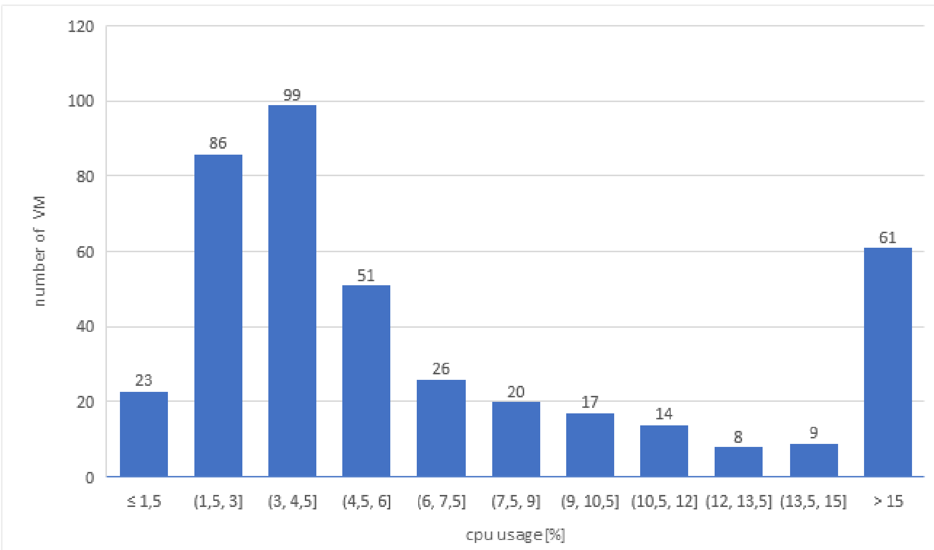

Figure 4.2: VM distribution with respect to CPU usage [%]

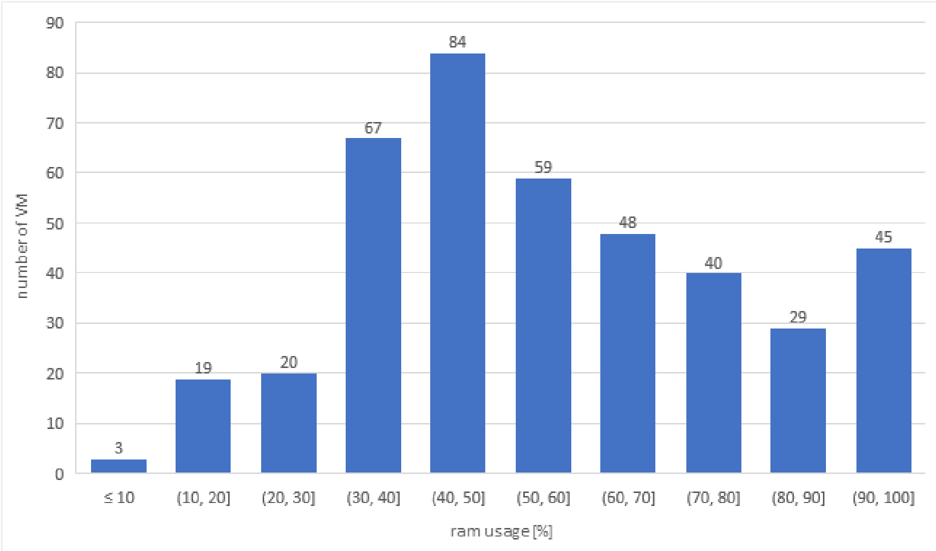

Figure 4.3: VM distribution with respect to RAM usage [%]

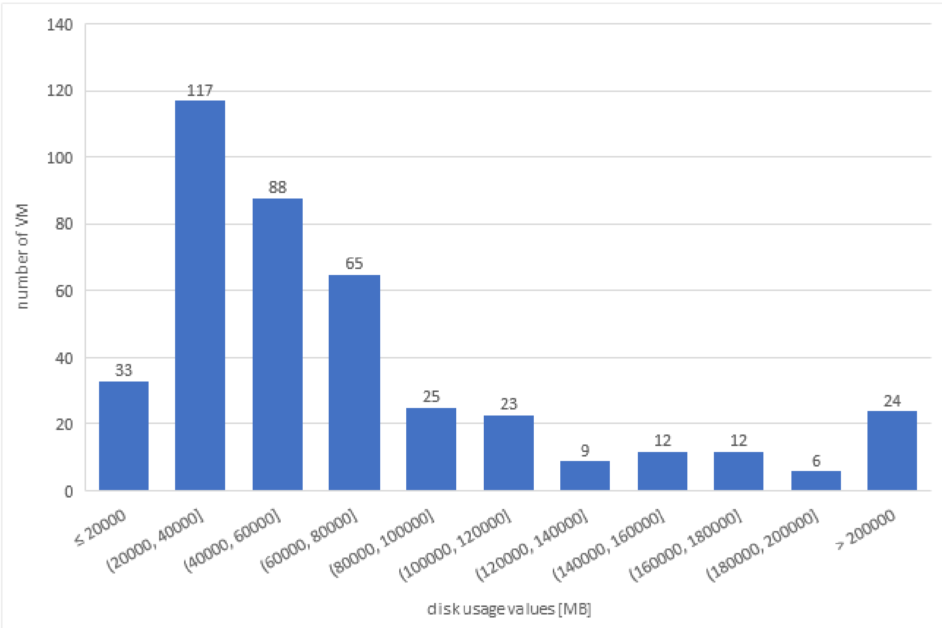

Figure 4.4: VM distribution with respect to disk usage [MB]

Disk usage and CPU distributions are very narrowed while ram distribution is more uniform. Indeed CPU usage is mostly narrowed below 15% values and disk usage values have low variance as well, on the contrary ram distribution is more spread between the whole set of possible values((0,100]). These considerations are double checked from the data reported in table4.1 where it is more evident the distances between mean values and the two external values of the distributions, indeed only RAM has a mean value almost centred between the maximum and minimum percentage.

The rounded covariance matrix [1] reported in table 4.5 is calculated over the 414 virtual machine considered after performing normalization in range [0-1] and figure 4.6 is basically a graphical not-normalized representation[2] of this covariance matrix.

	Disk	CPU	RAM
Disk	0.0135	0.0013	0.0009
CPU	0.0013	0.0203	0.0075
RAM	0.0009	0.0075	0.054

Figure 4.5: Rounded covariance matrix of normalized features

As expected in relation with the previous bar plots and table 4.1, RAM is the feature with higher variance and its distribution are more uniform over its minimum and maximum range. Disk usage and CPU present lower variance values and this is confirmed from the flattened distribution which also highlights the presence of outliers.

The three features do not present any strong correlation among them and this is deducible both from the off diagonal values of the covariance matrix and from the corresponding scatter plots. CPU and RAM anyway seems slightly correlated but otherwise the orange regression lines present very low angular coefficient.

A scatter plot representation helps to understand better the relation between feature but at the same time can be useful to partially exclude non linear correlation which are not possible to notice from the covariance matrix and in this case there is not any peculiar pattern that may be a signal of this situation.

Figure 4.7 is indeed a scatter plot where the 30 VM, corresponding to the more outliers values, are neglected for visualization purposes. This 3D visualization can

[1] Covariance matrix is a statistical tool used to describe correlation between features. On the main diagonal are computed the variances of the corresponding features, while other entries indicate the covariance among different features. So it is possible to understand better the distribution of variables considered and eventually find a statistical dependencies between each-other.

[2] the figures in the main diagonal are histogram representing distribution of the virtual machine based on the feature while the other squares are scatter plot with marked in orange the corresponding regression lines

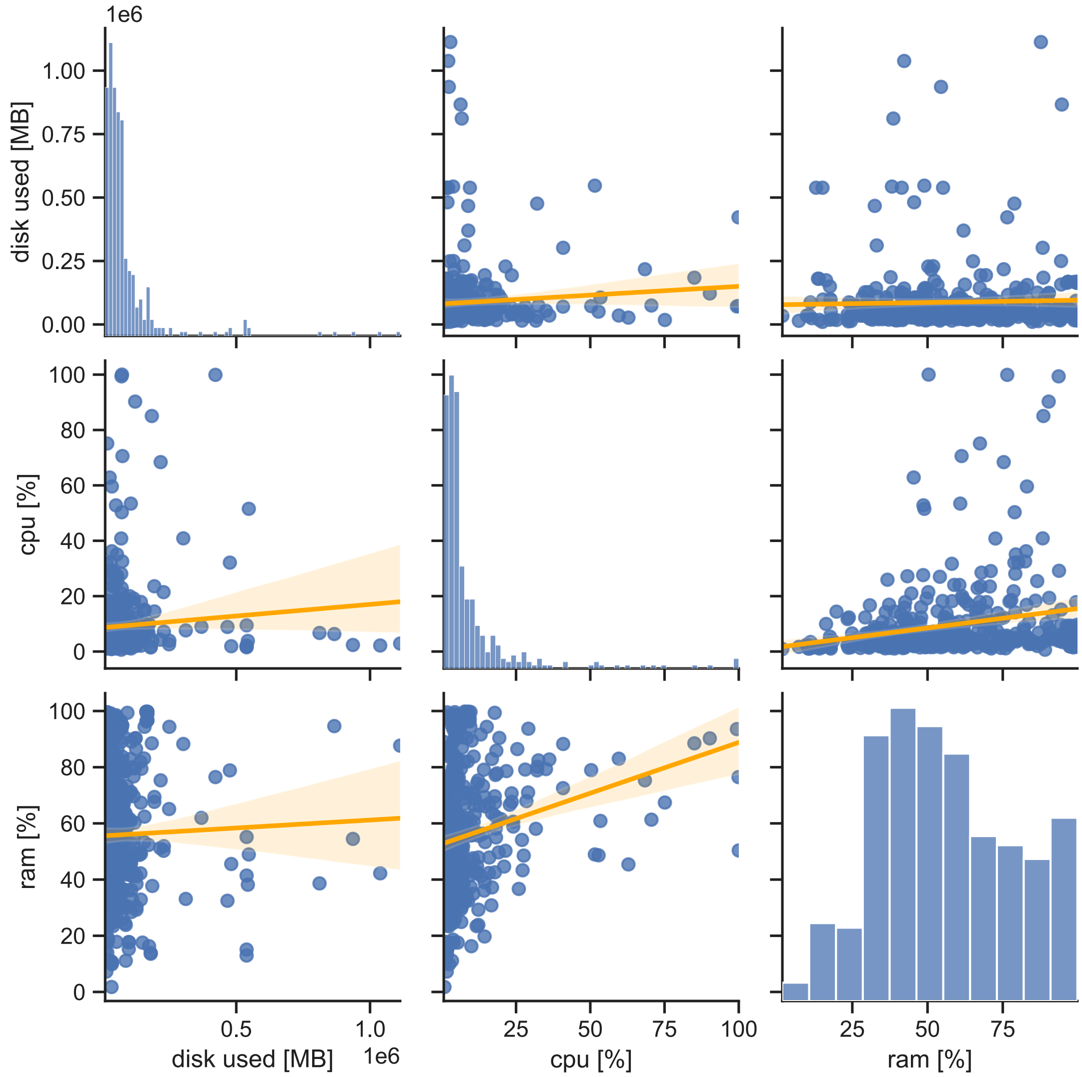

Figure 4.6: graphic representation of the covariance matrix

be useful to have a complete overview considering all the 3 features at the same time and it is basically a combination of the scatter plots represented in figure 4.6 which are anyway visible also separately in the appendix.

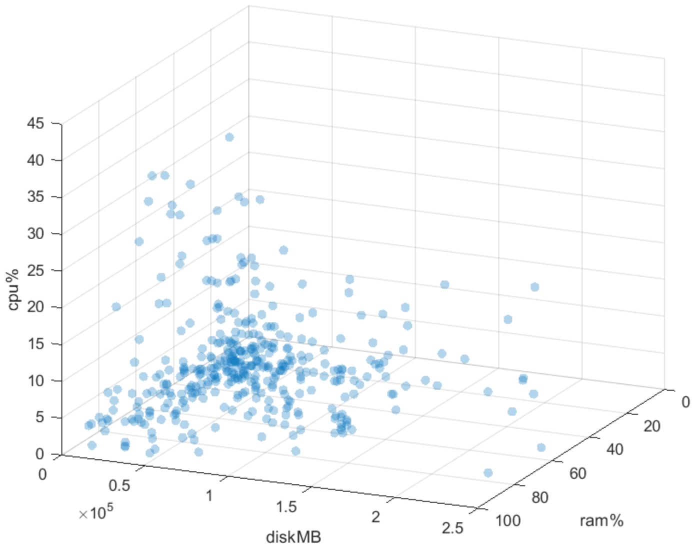

Figure 4.7: 3D scatter plot of VM based on the 3 main features

The various points in the this plot represent the virtual machines with theirs features on the three axes and we can notice how all the points are concentrated and flattened on the RAM axes which is the one with higher variance.

4.2 Methodology

The features used for characterization are disk usage in MB, CPU and RAM in percentage. Can now be useful outline the main steps used to derive results based on VM characterization:

1. find a way to introduce a control over virtual machines which permits to evaluate them based on the features mentioned;
2. implement this control policy over the "eco4cloud" simulator;
3. observe effects of this approach on the aggregated results which are natively returned by the simulator;
4. focus over migration itself and over the DCs state pre and post migration;
5. understand the results obtained taking into account the simulator functionalities ant the covariance matrix;

The first two points indicated will be discussed in the following subsections while the other three points will be argued in the results section.

4.2.1 Thresholds

To fulfil point 1 a virtual machine control policy was implemented. Before migration, a set of consideration, already mentioned in previous chapter(3), were done for each virtual machine but none of them consider the VM itself and its features.

So the key idea is to introduce in the code a routine where each virtual machine is evaluated also for its characteristics and eventually allowed to migrate.

To perform this control policy a set of thresholds (TH) were created with the purposes of constraint migration; the final goal is to allow each VM to move into another DC only if its disk usage, RAM or CPU value is under the TH imposed.

To choose these THs was used the median[3] function applied to each feature distribution. The choice of median is brought by the necessity of applying a control policy as even as possible and be aware of the influence of each TH implication during the simulation (TH corresponding to median implies a better control over the migration performances).

To identify more than one TH for each feature, the median was computed again over the two subset obtained from the first division and so is possible to get back nine values (three thresholds for each feature), which divide the virtual machine population in five subset on different bases.

Moreover, by considering also min and max values of each distribution, it was obtained a total of five different thresholds for each metric. Including min value as threshold is useful to have a double check on this control functionalities, the situation in which such a TH is used corresponds indeed with a simulation where the migration is completely avoided.

Table 4.8 reports these 5 values for CPU RAM and disk usage with an highlight on the main median which divides the population in half. As expected only RAM THs among them is near the mean value stated in table 4.1.

	TH1	TH2	TH3	TH4	TH5
CPU[%]	0	2.9	4.48	9.5	100
RAM[%]	0	39	53	73.2	100
Disk[MB]	0	34557	53451	87071	1113000

Figure 4.8: Thresholds used for each feature

So, for example by neglecting eventual lower boundary (imaging it equal to 0),

[3]Median of an array (or in general of a set of values) is a statistical indicator which identifies the value among them that divide the whole data set into two subset, with the same number of entities, composed respectively by the higher half and lower half of the distribution. So, for example in a vector of 100 random variables, the median is the value greater that 50% of them and lower than the other 50%. In a ideal Gaussian distribution median correspond to mean

changing the value of CPU threshold from 4.48% to 100% implies to give averagely %50 more possibilities to allow migration because the amount of virtual machine satisfying the constraints double.

Migration is granted only for Virtual machine which have a corresponding value, of the feature considered, included between an upper and a lower threshold. In particular thresholds were used in two different ways:

1. Cumulative case with increasing boundaries: at each iteration, the threshold considered as upper bound, changes. The lower one remains equal to zero.

   So in this case the simulator becomes less constraining at each iteration and as consequence also the number of virtual machines allowed to migrate increase accordingly.

2. Interval case with fixed interval: the lower and the upper bound changes together at each iteration. In particular, when the threshold set as upper bound changes the lower bound becomes equal to TH value used as upper bound at the previous iteration.

   In this case the possibility of migration for each virtual machine remains the same at each iteration despite the threshold considered but the kind of virtual machines allowed to migrate changes.

These two approaches main difference resides in the number of virtual machine migrated because the second one remains strict over the whole simulation while the first one increases the probability of migration at each iteration.

Clearly the number of TH used may vary as well and in a first moment the code implemented, consider only the external values of the distribution and the main median as TH values. Which basically creates 3 different situation that can be used as example to understand better the TH control policy:

1. TH=0 : no migration allowed

2. TH=median over the whole data set: only the virtual machine with a feature's value under this TH are allowed to migrate and this number is on average corresponding to half of the population.

3. TH=max value for the considered feature:this situation opens two scenarios:

   - cumulative case: all the virtual machines are allowed to migrate because the lower bound remains equal to zero while the upper one is at its maximum.
   - interval case: only the half of virtual machine's population presenting a feature's value higher than the lower threshold and smaller then MAX, are allowed to migrate. So, the probability to migrate remains equal to

case 1 but the type of VM migrated changes because now the considered VMs have corresponding features of higher magnitude.

The whole simulation is looped over the 3 kind of different features and over the various number of THs considered for each of them. Figure 4.9 gives an idea of the implementation of the thresholds policy over the flow chart in Fig. 3.2

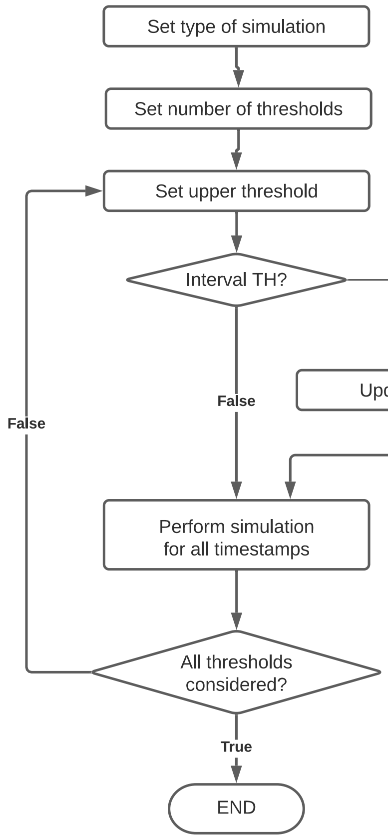

Figure 4.9: Thresholds implementation

As stated "eco4cloud" simulator iterates over different timestamps (up to one week corresponding to 168 timestamp) changing the parameter of each DC at each timestamp. Now to implement the control policy, the whole simulation is iterated multiple times using different thresholds which change both in type and value.

It is important to distinguish the loop over timestamp mentioned also in chapter 3, (which has the purpose of changing the energy parameters of the simulation for each timestamp) from the outer loop implemented after, which iterate all the simulation many times with different parameters and different THs.

So the number of obtained results changed as well. The aggregated result becomes one for each TH used and punctual results increase too.

Originally the simulation is performed considering for example 168 timestamp with a consequent punctual output composed by 168 rows while, with TH approach, the total numbers of rows returned are equal to:

$$rows = timestamps * features * THs \tag{4.1}$$

which means that considering a whole week simulation with all the features and thresholds argued so far, the final output has:

$$168(timestamps) * 3(features) * 5(thresholds) = 2520rows \tag{4.2}$$

The final goal is to use upper and lower boundaries to constraint migration and then see if the results obtained matches the expectation considering the simulator costs functions and the VM distributions just discussed.

4.2.2 Pre/Post migration

In order to focus better on migration performances, become important to isolate its contribution from the rest of the consideration and variables of the simulator environment.

The final aim is to compute all the system output parameters in two different moment, pre and post migration and then analyse the differences between the data obtained.

Can be useful to remember one more time the ".updateTotalEnergyCost()" function which have the important role to update the system condition at each timestamp. As said, it is a cumulative function which also keep track of the previous timestamp so at each iteration sums the new values obtained to the previous ones.

The use of this function is needed to efficiently update the system parameters but has the drawback that is a function born to be used only one times for timestamp. Indeed due to its cumulative behaviour whenever it is invoked after the first times alter the results for that timestamp because it sums more than once the costs and variables just computed.

So, in order to isolate properly the migration contribution it was necessary a series of steps to be performed after the VMs turnover but before the VMs migration, visible in the diagram in figure 4.10:

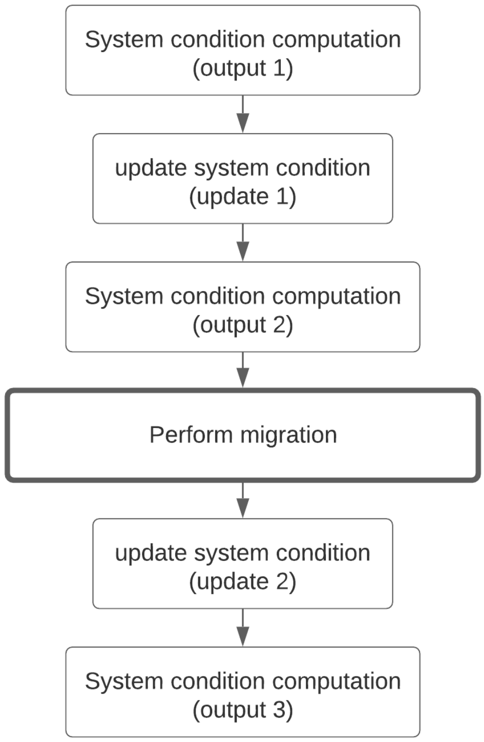

Figure 4.10: Pre/Post computation diagram

The different outputs obtained corresponds to three different moments during the same timestamp iteration and they were then subtracted by each other to obtain the final output needed, through the following computation:

1. $differenceA = "output2" - "output1"$

   $differenceA$ computation is needed to eliminate the influence of the previous timestamps of the simulation and indeed it is the equal to the raw contribution of the last update performed.

2. $differenceB = "output3" - "output2"$

   $differenceB$ computes effectively the difference between the migration impact on the general output but it is still biased by the previous updates made which have altered the values.

3. $differenceC = "differenceB" - "differenceA"$

   Finally, $differenceC$, is the real output desired which represents exclusively the difference between the system pre and post migration with all the bias and alteration avoided.

Thanks to this procedure was possible to use more than once the update function without incurring in the drawbacks. Indeed this fucntion is used to update the system status pre and post migration but without this algorithm it would have altered the result.

4.3 VM migration and DC consolidation: a comparison

Before considering the results obtained there is one more important consideration to be made regarding the consolidation impact on simulator environment.

In section 2.2.1 were already mentioned the concept of DC optimization and server consolidation and how fundamental can be to achieve optimal energy saving. Also the simulator environment considers this important approach and perform it once for each timestamp after the migration.

So, both VM migration and consolidations aims to obtain some benefits and one approach does not exclude the other but, on the contrary, it is convenient to perform them together to improve performances. Although for this thesis argumentation it is useful to distinguish their contribution in order to highlights only migration results.

In a first moment, consolidation contributions functions remain active but, proceeding with VM migration considerations, it becomes clear the importance of distinguish it. Consolidation mechanism gives a sort of bias on output parameters because it is always convenient to be performed and moreover it becomes even more important in case of VM migration.

In order to distinguish completely this two contributions, the code is modified and the consolidation is moved from after the migration to instead of it. Then the pre/post methodology just described was used in order to highlights the contributions and two different simulations are performed over 168 timestamps: one considering only migration and the other one considering only consolidation.

The plot in Fig. 4.11 is the merged results of this two simulations and shows the variation of costs in the simplest case without renewables, function of the timestamp and the energy efficiency approach used.

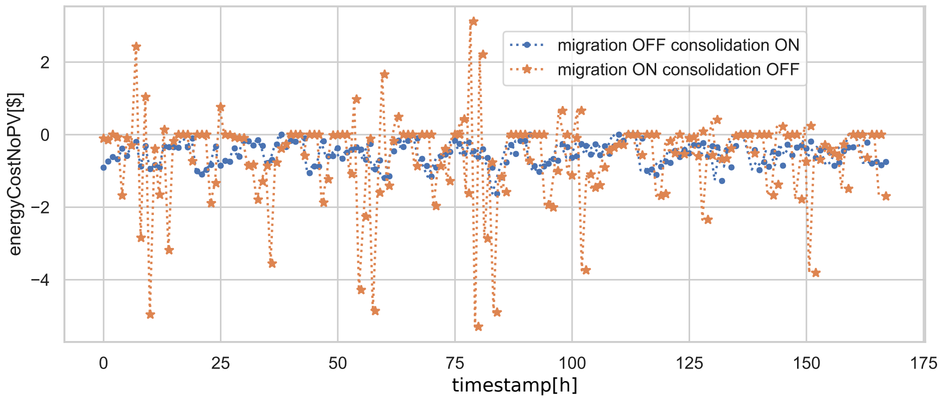

Figure 4.11: Comparison between consolidation and migration

Both migration and consolidation gives back, on average , a positive contributions to the reduction of costs but the behaviours differ. In particular consolidation has only negative values which means it is always convenient while migration curve fluctuates more and has higher peaks. So, the migration shows off a greater potentiality in reducing costs but has also a higher variability while the consolidation gives back a lower but steadier benefit.

This result is partially expected in particular because consolidation does not depends at all from the renewables at disposal and being performed on the same data center does not implicates any real change of environment for the virtual machine.

So, the simulation performed, which brought to the following section results, are done without performing any consolidation in order to show better the migration influence.

4.4 Results

Before starting to comment results, can be useful resume a set of considerations made so far:

- the score which determines best and worst DCs to perform migration does not consider load on the DC itself but on contrary it is function only of the DC energy properties like energy cost PUE etc.
- migration performances are better in a non-saturated system
- energy costs of the system are computed as function of a fixed component plus a variable part relative to the DCs' workload. This workload is represented

only by the amount of CPU usage in MHz, sum of the contributions of the several virtual machines presents in the DC.

- the consolidation is avoided in order to highlight only the migration effects

Moreover it is important to define the output variables considered:

1. "#VM migrated": it is always useful to monitor the number of VM migrated under different conditions.
2. "energyCostNoPV[$]":this it the energy consumption cost in dollars without considering renewables and solar contribution and it is computed at each iteration by considering also the load of each data center.
3. "finalEnergyCost[$]": this value is the total energy consumption cost in dollars considering renewables and batteries. It is equal to the difference between "energyCostNoPV" and the contributions of the renewables which it is also computed considering DCs load and battery capacity.

4.4.1 Aggregated results

Aggregated results are the final output of the "eco4cloud" simulator and indeed are obtained in a moment corresponding to the output box of the diagram in figure 3.2, in the next figures the plots' x axes are categorical in order to indicate the variation of thresholds for different features in the same image.

All the thresholds used are reported in table 4.8 but more important than the numerical value itself is to remember that using the median as TH implies an increment proportional to the distribution at each variation.

First of all, it is interesting to consider the energy results in the simplest and less noisy case represented by simulator output after the first single timestamp performed with different thresholds approaches. In this condition the system is far away from the saturation condition and the renewables resources are still completely at disposal in each data center.

However, at the same time, in the first timestamp the steady state condition is not reached yet because this is the timestamp right after initial assignment.

In the assignment the score of each DC is computed after each single VM placing while in the migration the score is computed only once at the beginning of the function and the origin and destination DCs remains the same for all the VMs. This means the two scores in this case can drive to particularly differences between each other's because based on two different computational functions.

Moreover it is important a last consideration about the ratio between thresholds and VMs migrated in this transient case. Indeed, with the exponential increase of the number of VMs at each timestamp the ratio between VMs migrated and thresholds is likely to increase in the cumulative case, despite the use of medians as THs, while at the first timestamp this condition is avoided.

Figure 4.12 shows the variation of the number of VM migrated in a scenario with cumulative thresholds so, where the boundaries becomes less strict at each iteration.

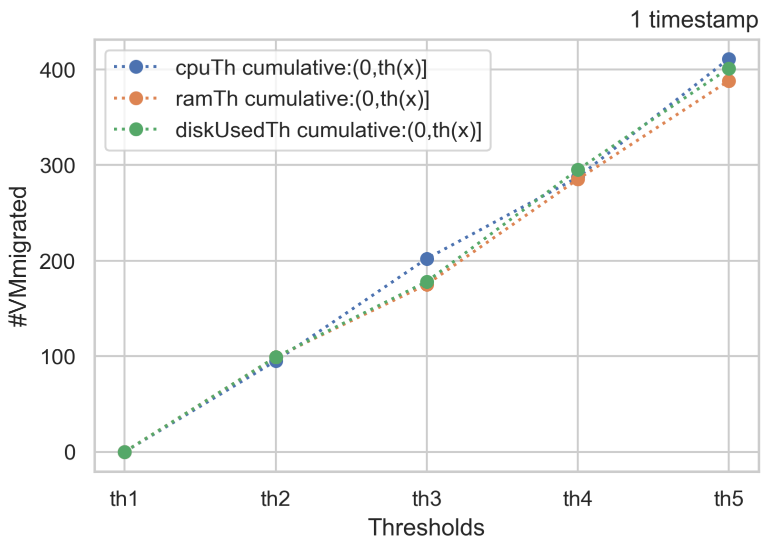

Figure 4.12: VM migrated over cumulative thresholds: one hour

As expected in this first case the increasing of the VM migrated is directly proportional to thresholds applied, with a ratio equal to 1 thanks to medians. Considering only one timestamps is like linearize the exponential by considering only a small region and so the exponential increase of VM migrated is neglected.

Figures 4.13 and 4.14 are two plots with "energyCostNoPV[$]" and "finalEnergy-Cost[$]" over THs in a cumulative case as well.

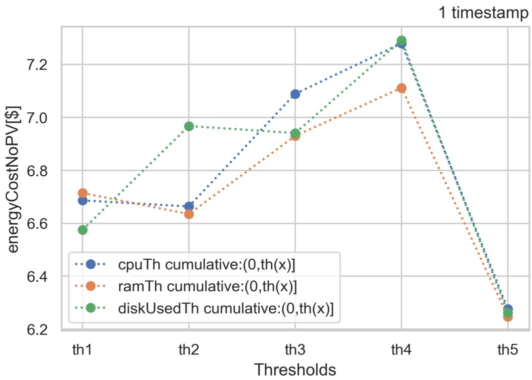

Figure 4.13: Energy costs no PV over cumulative thresholds: one hour

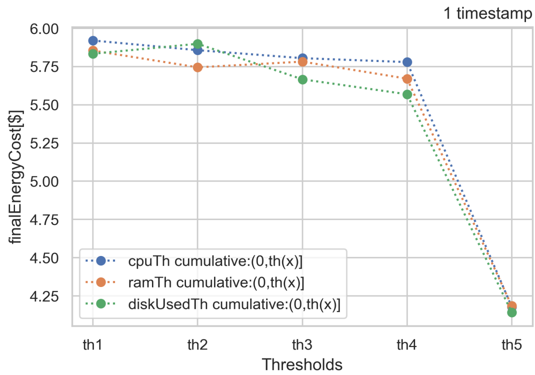

Figure 4.14: Energy final costs over cumulative thresholds: one hour

The first one shows an increasing trend until TH4 and then a sudden drop. This behaviour is partially predicted because influenced by the difference in VM allocation methodologies between migration and assignment. Moreover an higher amount of VM migrated helps to reach faster the steady state condition where the migration effects should be to lower the energy costs.

FinalEnergyCost[$] instead has a very smooth decreasing trend until the sudden drop with the highest TH. As expected, renewables and batteries exploit migration performances and in this particular case, cushions the negative slope of the energy costs without them.

With TH4 there is a sudden drop in energy costs despite the type of feature considered. This is an interesting results especially if linked to the increment of VM migrated. In fact, the migration increases by only a 25% but the behaviour of costs (in particular in the case without renewables) does not change accordingly.

The only features which is expected to have a considerable influence on cost is the CPU because each VM contribution is used to calculate the load of each DC. Despite that, the consideration made from the covariance between features (figure 4.6 and table 4.5)does not explain such a sudden change for all the three lines of the plot.

So, the reason may be searched elsewhere and in particularly can be related to the number of VM migrated. It is possible that there is a minimum amount of virtual machine migrated that helps the system to go beyond this transient condition and after that limit is surpassed the migration starts to become convenient anyway due to the increasing of the load migrated.

Figure 4.15 represents the number of VM migrated in the interval thresholds case with only one timestamp, while figures 4.16 and 4.17 are the two plots corresponding to figures 4.13 and 4.14 with the only difference to consider THs in the interval case.

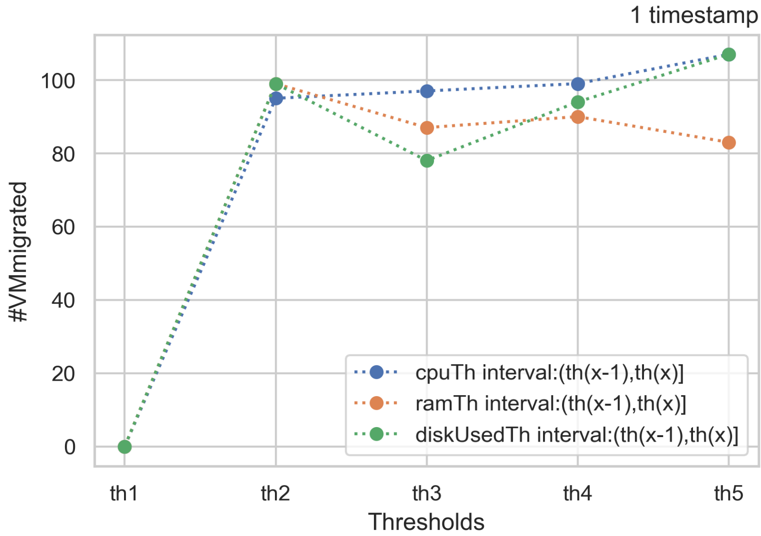

Figure 4.15: VM migrated over interval THs: one hour

In figure 4.15 is possible to see as the virtual machines migrated number remains quite stable. This happens because the possibility to migrate remains unchanged, due to the usage of median combined with a shift in threshold boundaries and the only exception occurs for th1 where both inferior and superior limits are zero and no migration is allowed.

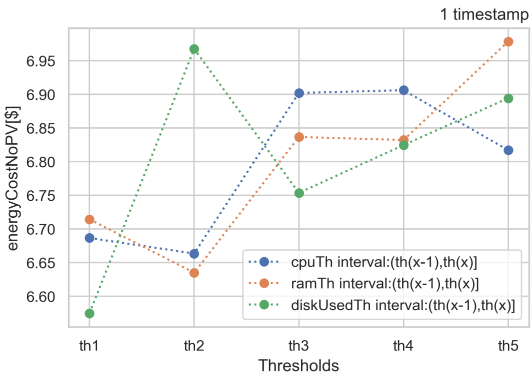

Figure 4.16: Energy costs no PV over interval THs: one hour

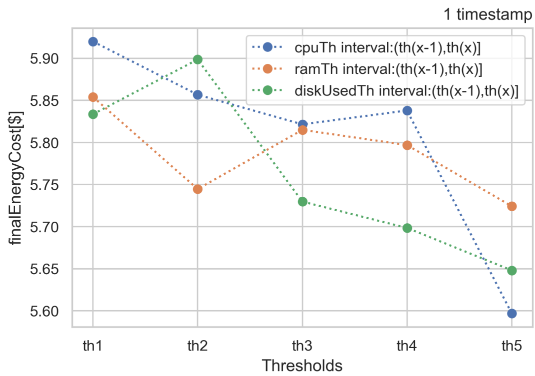

Figure 4.17: Final energy costs over interval THs: one hour

In figures 4.16 and 4.17 are evident the differences in respect to the cumulative case but remains an increasing trend regarding costs without renewables and e decreasing one for total costs where are also considered batteries and PV production. In this case anyway there is not present any drop after th4.

These are some first considerations for the aggregated results over one single timestamp right after performing the VM assignment. Now, can be useful to observe the same results attained with a simulation over a week (corresponding to 168 timestamps) with both thresholds approaches.

Figures 4.18, 4.19, 4.20 and 4.21 represents the final energy costs (considering also renewables and batteries) obtained in this case together with the related amount of virtual machines migrated.

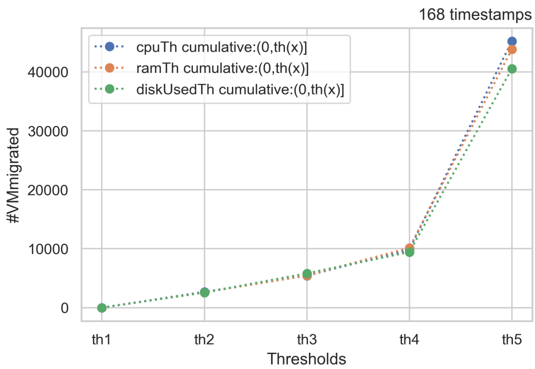

Figure 4.18: VM migrated over cumulative THs: one week

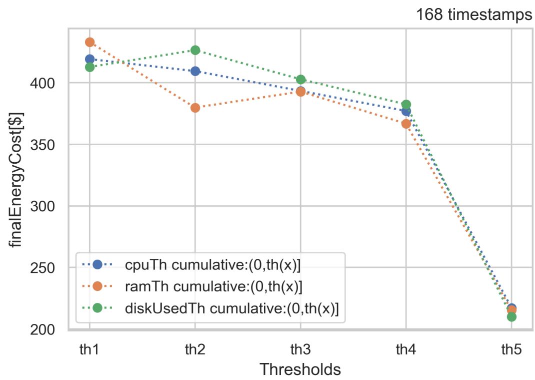

Figure 4.19: Final energy costs over cumulative THs: one week

First of all, it is evident the difference in the amount and behaviour of virtual machines migrated between figures 4.12 and 4.18. Indeed in this second case it is highlighted an exponential trend in the cumulative TH case that was not evident in the previous plot.

Increasing the threshold means increase the possibility of migration but the VMs which attempts the migration are not equally distributed. So, there is a cumulative migration effect driving to this exponential growth that is avoided in the case of only 1 timestamp because the difference in the number of VM allowed to migrate, is not summed for the 168 iteration.

Despite this difference, figure 4.19 shows a similar trend with respect to figure 4.12 with a sudden drop corresponding to the greatest thresholds.

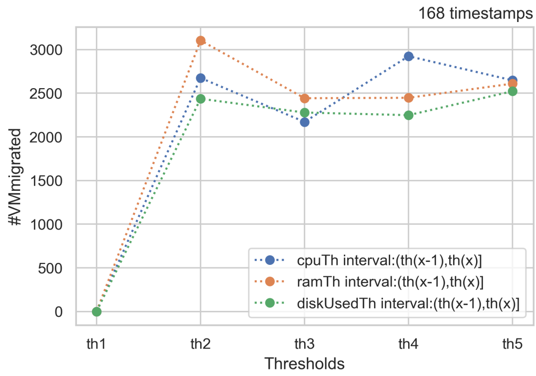

Figure 4.20: VM migrated over interval THs: one week

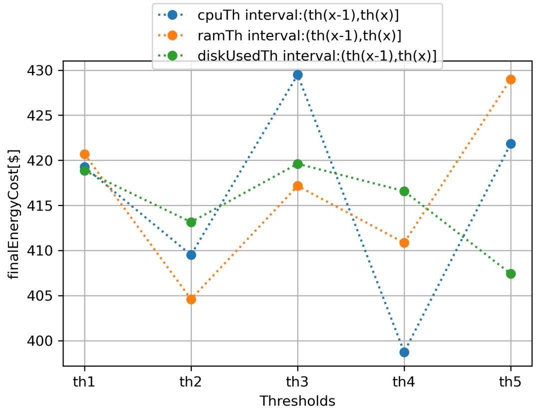

Figure 4.21: Final energy costs over interval THs: one week

In figure 4.20 is represented the amount of VM migrated with an interval TH approach and in fact there is not any qualitative difference within the plot in figure 4.15, proof of a correct functioning of this interval TH policy also in the week simulation case.

The main differences are presents in the plot in figure 4.21 especially with respect to figure 4.17 and figure 4.19. In this case the final energy costs do not display any trend or changes accordingly to TH and do not show either any difference with th1 case where no migrations are allowed. These fluctuations seem to be unrelated to the number of VM migrated and to the migrated VM feature values as well.

This discrepancy between the energy cost plots is a considerable result. Despite the application of different thresholds approaches all the other parameters of the system remains unchanged including the number of virtual machines in the system.

The only big difference between these two scenarios is the number of virtual machines migrated. The minimum amount of virtual machine migrated in the first case is indeed comparable with the average amount of VM migrated in the second case.

Moreover in figure 4.17 it was observed a decreasing trend despite the low amount of VM migrated. So this may implies that more important than the amount of VM migrated itself is the ratio between this number and the amount of VM of the system.

These two last considerations are very important because the output obtained may indicate that there is a minimum ratio which represent a condition to be matched, to obtain considerable performances.

It indeed does make sense that if the system has a huge amount of DC and a consequential higher load, migrate only a, too small, sub portion of it does not influence the performances. This situation comes out only in aggregated results with interval approaches where the number of VM migrated is low despite the huge number of VM in the system.

In particular the total amount of VM which populated the system, during these simulations, are on average:

- 33700 for simulation over 168 timestamps
- 5200 for simulation over 168 timestamps
- 700 for simulation over one single timestamp

These average values do not depend from the TH approach used but they are only function of the timestamps considered. By considering these values together with the related number of VM migrated and observing the consequential behaviour over the different threshold approaches, can be outlined a boundary. Indeed this consideration draw to the conclusion that if the ratio just mentioned is below 0.1 the migration performances are negligible. This means that for percentage of VM migrated less than 10% the migration does not impact enough on the system to observe real benefits.

Figure 4.23 represents the variation of costs without renewables over simulations with 24 timestamps with cumulative THs and figure 4.22 shows the VM migrated relative trend.

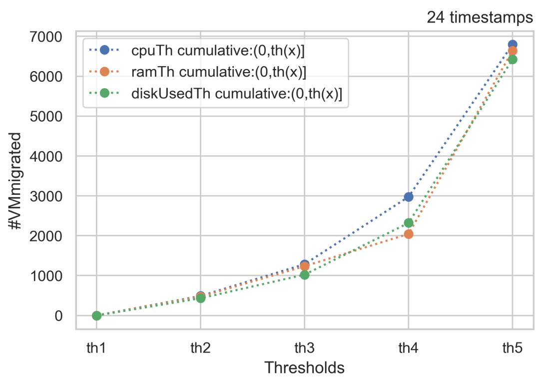

Figure 4.22: VM migrated no PV over cumulative THs: one day

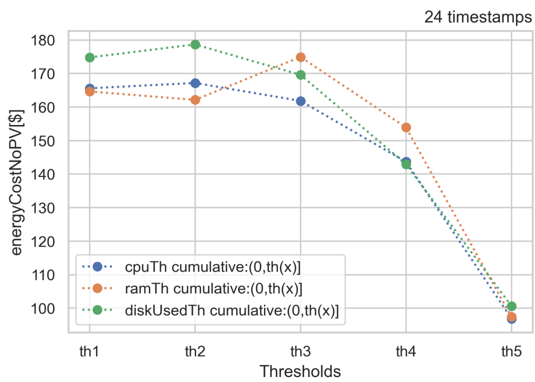

Figure 4.23: Energy costs no PV over cumulative THs: one day

It may be interesting to notice that in this case the transient effect noticed in figure 4.13 is not present anymore. Indeed, in this case the "energyCostsNoPV" variable has a decreasing trend and this may suggest that increasing the number of VM migrated implies a reduction of costs without renewables as well.

From the aggregated plots considered do not come out any relevant relation between features and cost reduction but it seems to be present a relation between the ratio just mentioned and the reduction of costs.

Moreover, in the cumulative case, was evident that more VM migrated implies a reduction of costs and this result may hide an eventual relation with VM features. Indeed as seen in chapter 3 costs are computed considering different contribution and there are many variable factors which influence the simulator. (e.g. like the energy cost which fluctuate during the day and the higher variability of λ and μ and the depletion of energy from renewables and batteries).

All these factors may override the eventual contributions given from a VM characterization which ends to be neglected in comparison to the migration itself and to understand better this situation may be useful a more punctual analysis that will be argued in the next subsection.

The figures representing the energy cost without PV over 168 timestamps and the final energy costs over 24 timestamps are present in the appendix.

4.4.2 Laod migration, PRE/POST analysis

In order to investigate better VM characteristic influences it was adopted a more punctual analysis and the following argumentation regards the output obtained by the application of the PRE/POST approach explained in the methodology.

This case is still different from the aggregated results obtained after 1 timestamp because neglect also the total cost of the DCs and focus only on analyse the perturbation brought by the migration. In this case are neglected many of the external factors of the simulator environment because the migration performance are evaluated for each timestamp and for each thresholds.

Aggregated results already give an idea of the advantages of migration and when they may be higher but they partially hide any particular relation with features that me be hidden from the sum of all the contributions brought by the system.

Now, it is important to remember that the algorithm in this case delete all external contribution and aim at distinguish only the punctual variation at each timestamp. So, this implies results of very small magnitude even in comparison with the one displayed in figure 4.12. Indeed the final aim of this analysis is to focus less on migration performance and more on the influence of VM features considering also that in this case the amount of VM migrated is considered separately for each timestamp

In the steady state condition seen in figure 4.23, the energy costs without renewables seems to decrease accordingly with the number of VM migrated. Now we will focus this punctual analysis only in this worst case condition without PV.

The outputs are obtained from a difference between the energy costs of the system after migration minus the energy costs before it. So, in the following plots a negatives values implies a money saving due to migration. Anyway in these punctual analyses it may be more useful to focus on the dependencies from THs and variation magnitude than the money saving itself which was already argued in the previous subsection.

Figures 4.24 and 4.25 represents the more complete output considering all the THs with cumulative and interval case over a simulation of 24 timestamp.

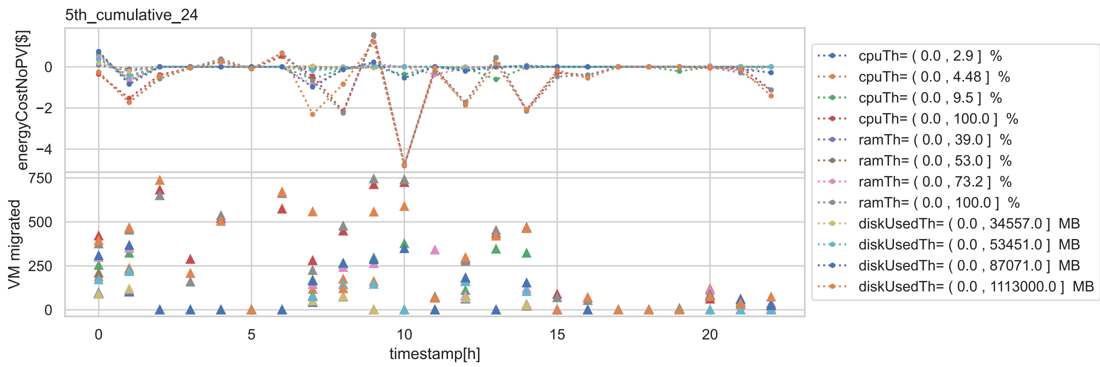

Figure 4.24: "Energy costs no PV in cumulative case with all THs"

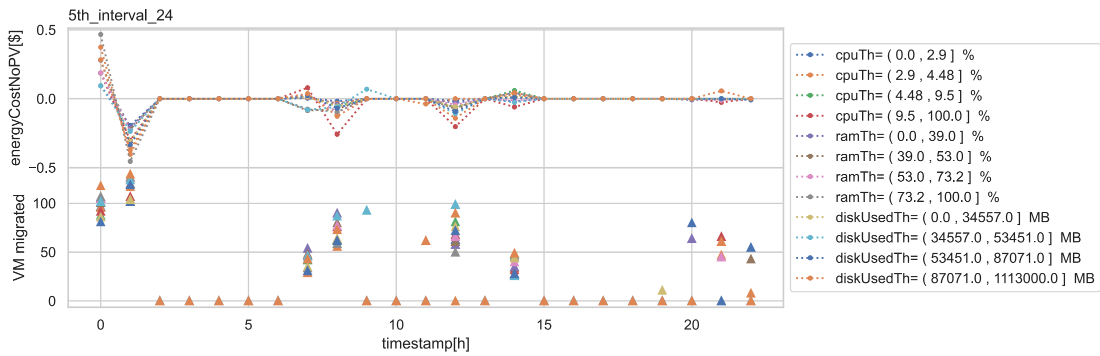

Figure 4.25: "Energy costs no PV in interval case with all THs"

These two plots are just an introduction to grasp better the general idea and make some considerations:

- The number of VM migrated is still greater in the cumulative case but there is not any strong exponential behaviour and the virtual machine migrated increase according to what expected by median implementation.
- As expected there are fluctuation around zero but the majority of peaks are negative and this confirms the migration has still an overall positive effect on energy costs.

- From a comparison between these two graphs, it can be observed again that the higher values of VM migrated increase the migration performances and consequentially reduce the costs.

- The first timestamps shows a positive peak while in the second ones is negative. This is an important confirmation of the consideration done regarding figure 4.12. Indeed, also in the aggregated results the energy costs without PV shows off an increasing trend in the first timestamp as consequence of a transient condition after the initial assignment. Moreover the only exceptions to this consideration regard the two negative points obtained in the first timestamp in correspondence of the two cases with the higher number of VM migrated and this behaviour recalls the one observed for TH4 in the figure just mentioned.

- In both plots there is a negative peak at second timestamp as rebound of the first one.

- In the interval THs approach (4.25)is especially clear that all the peaks are red, corresponding to the higher CPU interval threshold, despite the number of VM migrated does not change in respect to the other metrics. These observations may imply the presence of an expected relation between the migration of VM with higher CPU usage and the energy costs.

- Whenever there is a situation in which there is not present any migration, it has been observed and analysed a proper output file[4] meant to properly debug the simulation and it came out that the reason behind these events are completely different in the two cases:

  – in the cumulative case clearly is not possible to be feature THs policy fault. Indeed the cause is a non-feasible migration for that timestamp. Without constraints on bandwidth the main causes to be accounted for this type of fault are host saturation and delta. In both cases the migration is aborted for all virtual machines.

    In this case anyway it was checked to be delta's fault. Delta was already mentioned and it represents one of the variables checked to see if migration is feasible. It implies a case where the score computed among DCs are too similar. Indeed, this events occur at timestamp 17, 18 and 19 after a series of timestamps in which many virtual machines are migrated.

[4]In order to monitor migration, a series of warn were utilized and it was created also an output file where are reported for each timestamps and for each VM, a series of information about the virtual machine meant to be migrated and the outcome of this attempt.

– in the interval case on the contrary the reason is the features THs approach. In particular from the output file it was observed a tendency of the simulator environment to take together (in the related timestamps) the virtual machines more similar to each other.

  Moreover, there is a propensity to get first the virtual machine with features of higher magnitude this behaviour does not match well with the THs interval approach where both boundaries values increase at each timestamp. In fact, it is evident that these timestamps follows a series of one and more timestamps where the migration actually happened.

  These considerations may be important to outline also the simulator environment behaviour and find hidden pattern or eventual issues.

These are only general considerations on these outputs with many variables, now it may be interesting focusing more on singles features under different THs. In particular CPU is maybe the variables which may results more relevant due to the dependency in the cost function. Plots in Fig.4.26 and Fig.4.27 show migration performances with a simulation over 24 timestamp (one day) with only the highest thresholds considered.

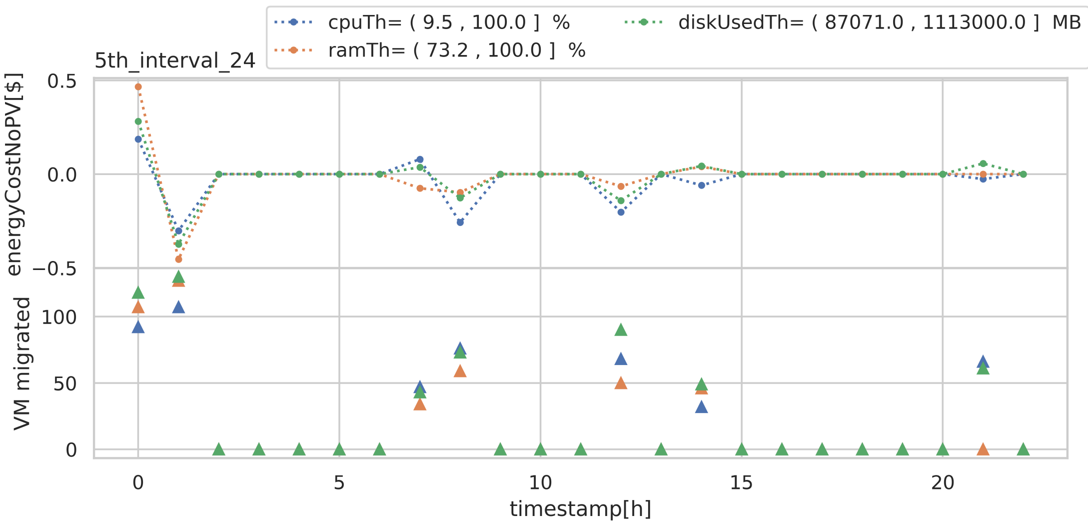

Figure 4.26: "Energy costs with highest THs: interval"

Figure 4.26 represents the interval TH case and indeed the number of VM migrated is lower. In this case it is highlighted the CPU influence that corresponds with the highest peaks despite the related number of VM migrated is similar or even lower than the others.

Figure 4.27 does not show any considerably relation between features and cost reduction and the behaviour are rather similar for all characteristics. Also in the

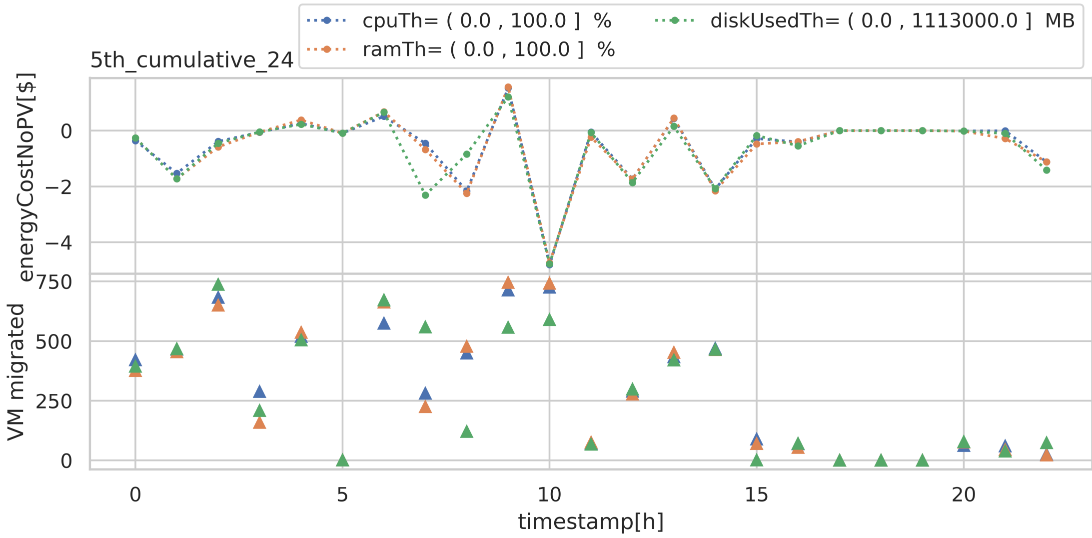

Figure 4.27: "Energy costs with highest THs: cumulative"

corresponding aggregated case after a certain amount of VM migrated the energy costs dropped despite the type of TH considered (4.12).

This confirms the hypothesis made right after the aggregated consideration. In this case as well the relation between features and costs seems to be hidden by the performance obtained by migrating a larger amount of VMs.

There seems to be a relation between CPU and the reduction of costs but there are only a few occurrences. So to have further proofs of these hypothesis, it could be worth to consider a whole week simulation with a TH interval approach where it is considered only the main median as interval threshold.

Figures 4.28, 4.29 and 4.30 show this kind of output where the THs force a binary choice over the VM considered for each feature as consequence of the utilization of only one threshold equal to the main median. Moreover by considering this analyses with multiple timestamp, the number of interesting occurrences increase as well.

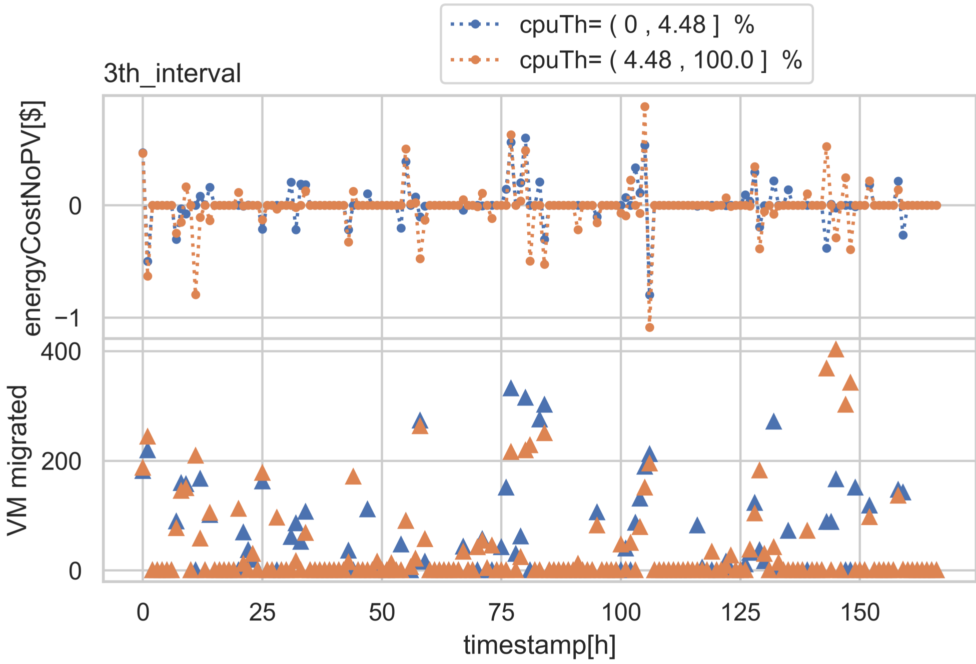

Figure 4.28: Energy costs, higher and lower half:CPU

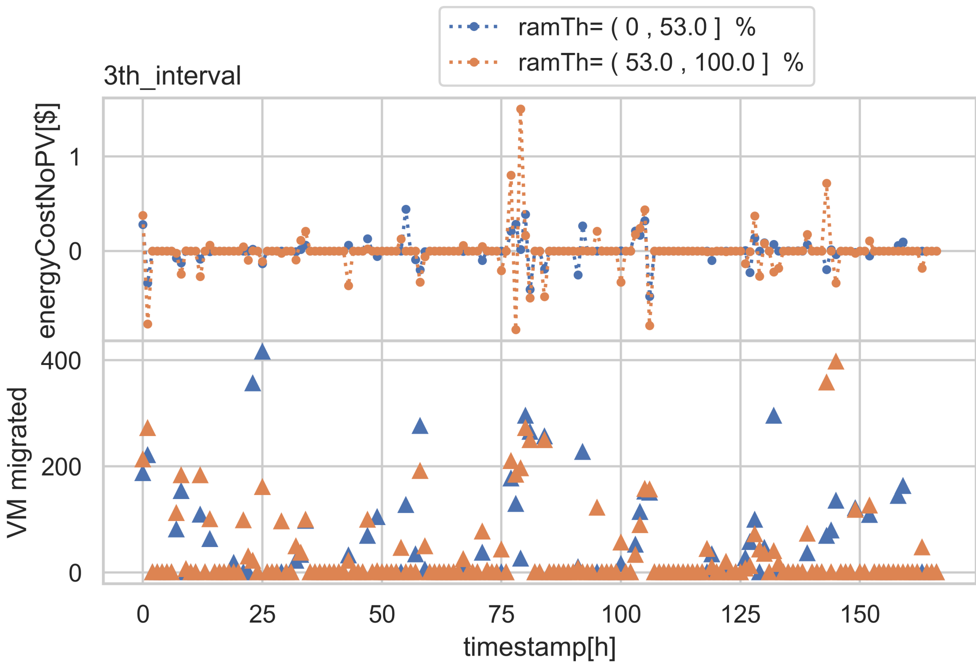

Figure 4.29: Energy costs, higher and lower half:RAM

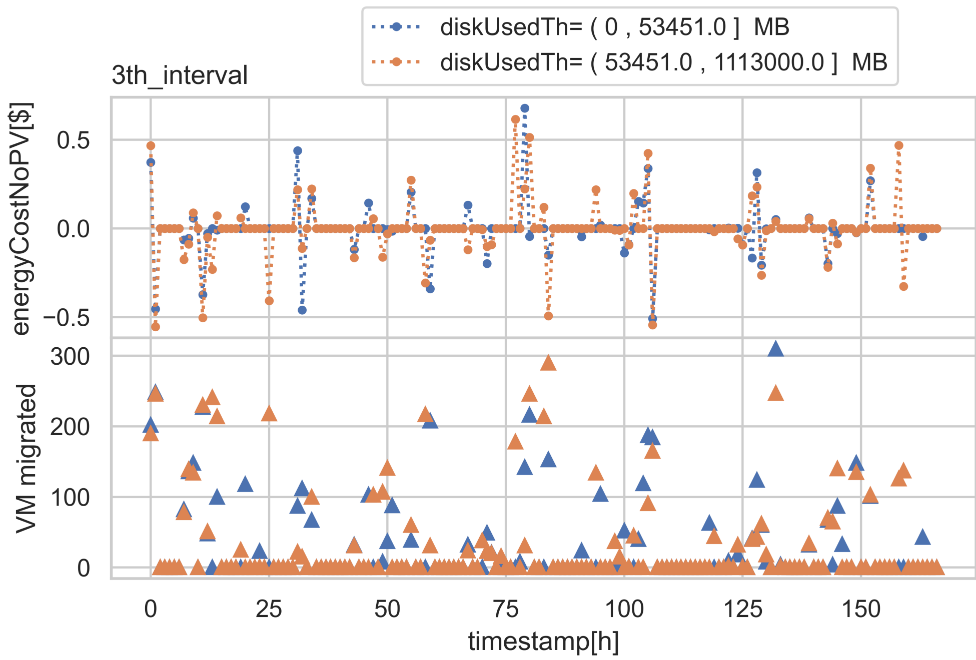

Figure 4.30: Energy costs, higher and lower half:disk

Regardless the number of VM migrated and the low magnitude of the variation of costs it is evident that almost all the peaks in the CPU plot are relative to virtual machines with higher CPU values and this behaviour seems to be present also in RAM charts but with a lower magnitude while in the plot regarding disk usage this trend seems to be almost completely off.

The correlation between the various features was already argued in the previous sections and the plots just considered seem to behave accordingly. Indeed, the CPU confirms to have the major impact due to influence in costs while RAM that is partially correlated to it shows a similar behaviour but with lower magnitude. The last and more uncorrelated feature is disk usage whose corresponding graph shows a different behaviour that can be only slightly associated to the other two.

The last analysis worth to be done regards the relation between costs, CPU and VM migrated and the plot in Fig. 4.31 shows together the number of virtual machine migrated with no constraints (all VM are allowed to migrates) over a iteration of 168 timestamps. As expected, the highest peaks (both negative and positive ones) are in correspondence of the higher bars which represents an higher number of migrated element.

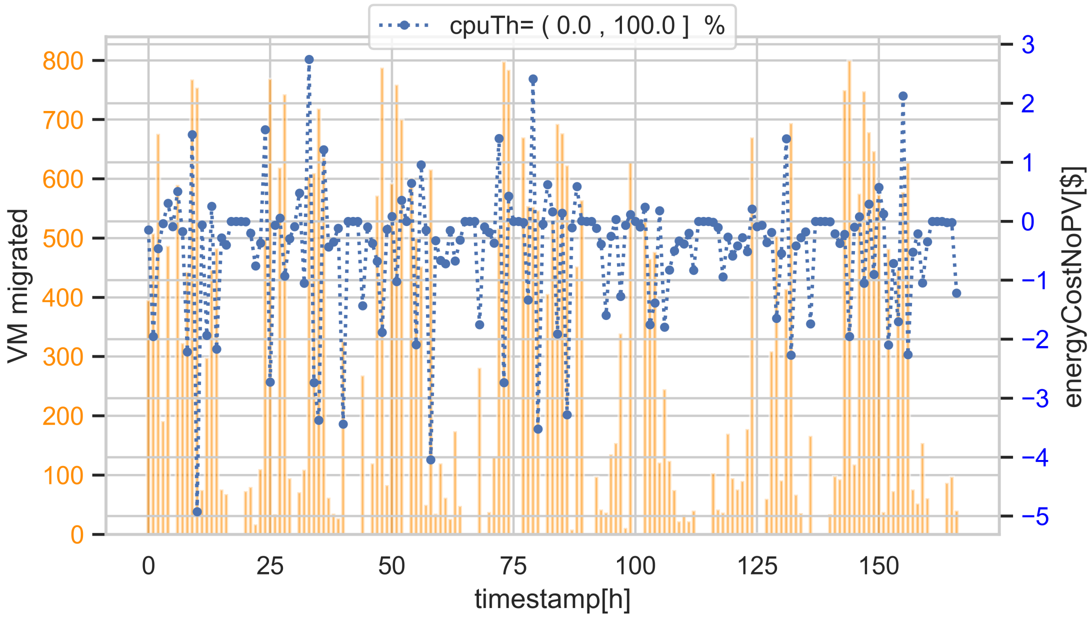

Figure 4.31: Energy costs and VM migrated over 168 timestamps

Chapter 5

Conclusions

Migration demonstrates to be potentially a very useful approach to increase energy efficiency.

The amount of virtual machines migrated shows to be particularly incisive over the migration performances. In particular the aggregated results outlined a relation between the ratio of VMs migrated over the VMs in the system and the migration performances.

To understand better the influence of a proper VM characterization has been analysed three main features which are CPU, RAM and disk usage. After studying the covariance among these three features, has been performed a punctual analysis on the migration performances. The results show in particular an expected relation with CPU which represents one of the main contributions in the energy costs function.

For all the simulations made and the results exposed the cost of migration has been neglected. Moreover the bandwidth considered was incredibly high (200Gbps) in order to neglect this constraint as well. These two parameters create an ideal scenario where all the types of virtual machine are equally convenient and equally feasible to migration. It can be useful to consider what a more realistic conditions may imply.

Figures 5.1 and 5.2 shows the amount of megabytes migrated over one week considering an interval thresholds approach. In this case are considered only three THs: $th1$ =0; $th2$ =the main median(4.48% for CPU, 53% for RAM and 53451 MB for disk usage) $th3$ = max values (100% for RAM and CPU,1113000 MB for disk usage)

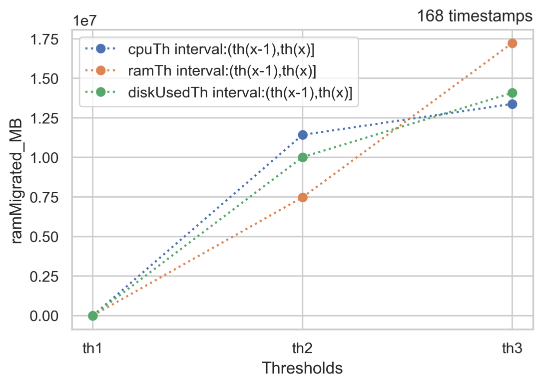

Figure 5.1: RAM migrated over interval THs: one week

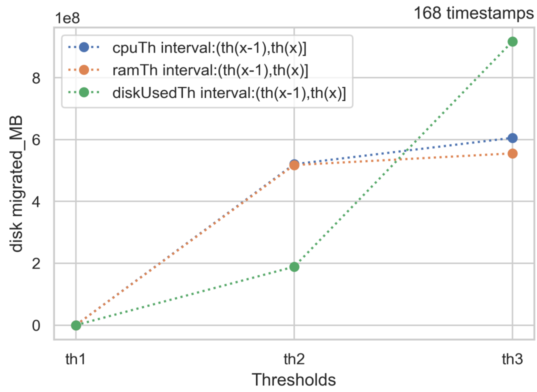

Figure 5.2: disk migrated over interval THs: one week

It can be observed in these two plots that, as expected, the amount of MB migrated increases considerably with the highest THs. Bandwidth and cost migration are particularly affected by VMs with higher values of RAM and disk usage because they are the ones which represents the higher amount of MB to be migrated. The data provided by Telecom about VM do not consider an effective migration cost and the simulator environment neglect it as well. Moreover the migration is performed considering an ideal execution migration time equal to zero and this means that latency is not considered as well.

These considerations are very important to evaluate migration performances in a more real scenario. Indeed in this case the advantages of a proper virtual machines characterization with a relative migration policy increase consequentially. Perform migration, over virtual machines with higher disk usage and RAM utilization, implies higher bandwidth utilization and larger amount of MB to be migrated. These conclusions highlights even more the advantages of migrating virtual machines with proper considerations about theirs features. In particular an ideal approach could be to promote migration of virtual machines with higher CPU utilization and lower disk usage avoiding as much as possible to move large quantity of data.

An interesting research by David Costenaro and Anthony Duer try to study and compute an energetic cost related to internet traffic [40]. This study estimates that total internet energy usage for transmission and communication represents at least 9.5% of the total consumption.

In particular Costenaro and Duer computes an average transmission cost for a 5MB streaming (a YouTube song) of \$0.0004 which means a cost of about $0.00008[\frac{\$}{MB}]$ for transmission purposes only. The plot in Fig. 5.2 has the highest peak over 800 TB and multiplying this value for the energy expenditure just mentioned the final outcome is a transmission cost of 64 thousand dollars.

This is a rough estimation, one further analysis on this thesis topic could be to attain more specific data about migration energy and cost expenditure and include them in the simulator. With this kind of data could be analysed an eventual trade-off between the benefit achieved from migration and its costs.

One more improvement and future work could be to evaluate a new "score" function which may consider some other variables like the migration energy costs and the effective load (not the estimated one) for each data center. This approach could be implemented together with a more punctual policy which evaluates this score after a certain amount of load migrated. Although, this is not an easy task because also compute the score too often can represent an energy expenditure and a computational cost as well, especially considering the magnitude of the problem.

There are a dramatic number of variables parameters and scenarios worth to be considered and analysed regarding data center, migration and virtualization performances worth to be analysed and studied. This thesis continuous a previous work by focusing on some of them but there are certainly a lot more work to do to increase energy efficiency and reduce consumption.

Appendix A

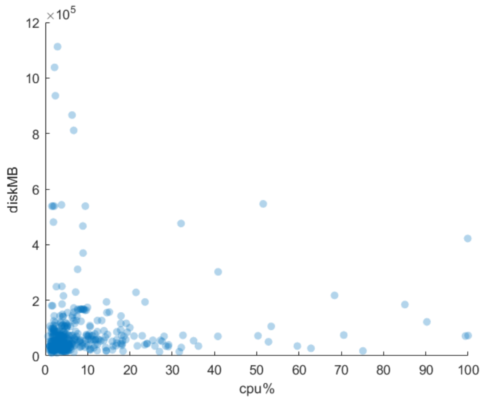

Figure A.1: scatter plot disk over cpu

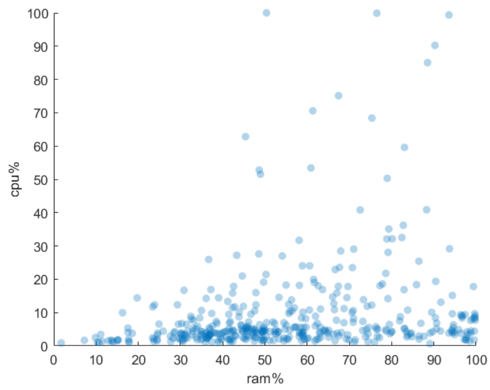

Figure A.2: scatter plot cpu over ram

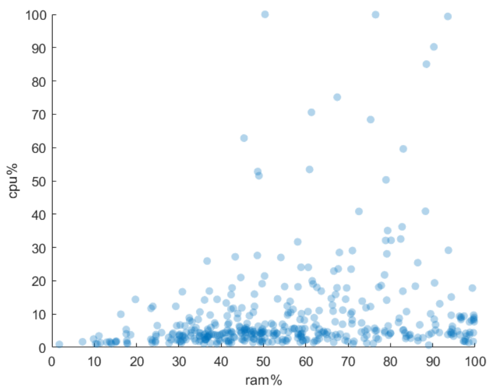

Figure A.3: scatter plot disk over ram

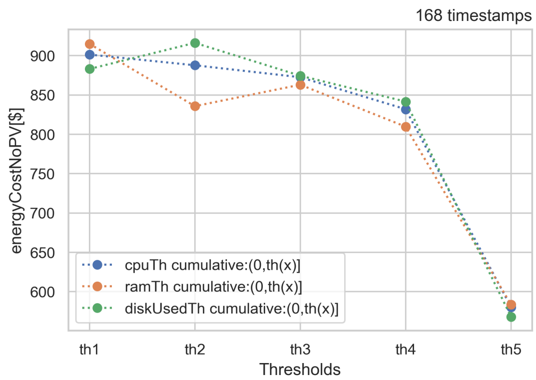

Figure A.4: Energy costs no PV over cumulative thresholds: one week

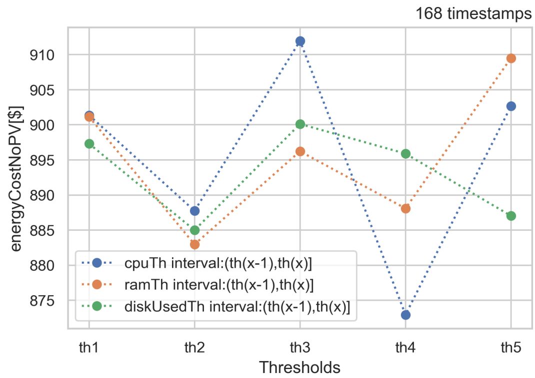

Figure A.5: Energy costs no PV over interval THs: one week

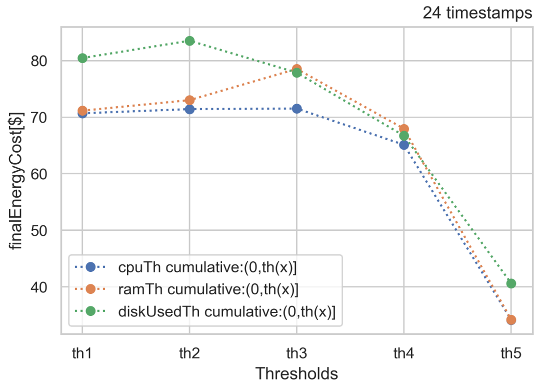

Figure A.6: Final energy costs over cumulative THs: one day

Bibliography

[1] Cisco and affiliates. *VNI Complete Forecast Highlights*. URL: https://www.cisco.com/c/dam/m/en_us/solutions/service-provider/vni-forecast-highlights/pdf/Global_2022_Forecast_Highlights.pdf (cit. on p. 1).

[2] Luca Castellazzi Paolo Bertoldi Maria Avgerinou. *Technical report by the Joint Research Centre (JRC)*. Tech. rep. 2017, pp. 13–15. URL: https://publications.jrc.ec.europa.eu/repository/handle/JRC108354 (cit. on p. 2).

[3] Synergy research group. *The Decade's Megatrends in Numbers*. Tech. rep. Jan. 2020. URL: https://www.srgresearch.com/articles/the-decades-megatrends-in-numbers-part-1 (cit. on pp. 3, 4).

[4] canalys research group. *Global cloud infrastructure market Q4 2019 and full year 2019*. 2020. URL: https://www.canalys.com/newsroom/canalys-worldwide-cloud-infrastructure-Q4-2019-and-full-year-2019 (cit. on p. 3).

[5] Yevgeniy Sverdlik. «Coronavirus Is Bad and Good News for the Data Center Switch Market». In: *Data Center Knowledge* (Apr. 2020). URL: https://www.datacenterknowledge.com/business/coronavirus-bad-and-good-news-data-center-switch-market (cit. on p. 3).

[6] Daphne Leprince-Ringuet. «'You can't just stop.' How data centers are dealing with the coronavirus crisis». In: *zdnet* (Apr. 2020). URL: https://www.zdnet.com/article/you-cant-just-stop-how-data-centers-are-dealing-with-the-coronavirus-crisis/ (cit. on p. 3).

[7] Simon Harris. «The future of data centers in a post Covid-19 world». In: *DCD* (June 2020). URL: https://www.datacenterdynamics.com/en/opinions/future-data-centers-post-covid-19-world/ (cit. on p. 3).

[8] Prof. Sally Eaves. *Le sfide per i Data Center al tempo del Coronavirus*. URL: https://www.kingston.com/italy/it/solutions/servers-data-centers/coronavirus-demands-on-data-centers (cit. on p. 3).

[9] World Health Organization. *Protecting Health from Climate Change, Annex1: Reducing your carbon footprint can be good for your health.* 2008. URL: https://www.who.int/globalchange/publications/factsheets/Kit2008_annex1_2.pdf (cit. on p. 4).

[10] I.E.A. *Data Centres and Data Transmission Networks.* https://www.iea.org/reports/world-energy-investment-2019/power-sector. IEA, Paris, 2020. URL: https://www.iea.org/reports/data-centres-and-data-transmission-networks (cit. on pp. 4–6).

[11] I.E.A. Masanet and E. «Recalibrating global data center energy-use estimates». en. In: *Science* 367.6481 (2020), pp. 984–986. DOI: 10.1126/science.aba3758. (cit. on p. 4).

[12] *Frequently Asked Questions (FAQs) - U.S. Energy Information Administration (EIA).* 2020. URL: https://www.eia.gov/tools/faqs/faq.php?id=74&t=11 (cit. on p. 4).

[13] *Italy CO2 Emissions - Worldometer.* URL: https://www.worldometers.info/co2-emissions/italy-co2-emissions/ (cit. on p. 4).

[14] Hannah Ritchie. «Energy». In: *Our World in Data* (2014). URL: https://ourworldindata.org/energy (cit. on p. 5).

[15] *Google efficiency report.* URL: https://www.google.com/intl/en/about/datacenters/efficiency/ (cit. on p. 6).

[16] I.E.A. *Global data centre energy demand by data centre type, 2010-2022.* IEA. URL: https://www.iea.org/data-and-statistics/charts/global-data-centre-energy-demand-by-data-centre-type-2010-2022 (cit. on p. 6).

[17] I.E.A. *Global data centre energy demand by region, 2010-2022.* IEA. URL: https://www.iea.org/data-and-statistics/charts/global-data-centre-energy-demand-by-region-2010-2022 (cit. on p. 6).

[18] Eric Masanet, Nuoa Lei of the McCormick School of Engineering, and Applied. «How Much Energy Do Data Centers Really Use?» In: *energyinnovation.org* (Mar. 2020). URL: https://energyinnovation.org/2020/03/17/how-much-energy-do-data-centers-really-use/ (cit. on p. 7).

[19] Jodie Van Horn Gary Cook. «How dirty is your data?» In: *Greenpeace International* (Apr. 2011), pp. 01–07. URL: https://storage.googleapis.com/planet4-international-stateless/2011/04/4cceba18-dirty-data-report-greenpeace.pdf (cit. on p. 7).

[20] Gary Cook. *How Clean is You Cloud?* Tech. rep. Apr. 2012, p. 07. URL: https://storage.googleapis.com/planet4-international-stateless/2012/04/e7c8ff21-howcleanisyourcloud.pdf (cit. on p. 7).

[21] Gary Cook and Yifei Wang Tom Dowdal David Pomerantz. *Clicking Clean:How Companies are Creating the Green Internet*. Tech. rep. Apr. 2014, p. 07. URL: https://www.greenpeace.org/usa/wp-content/uploads/legacy/Global/usa/planet3/PDFs/clickingclean.pdf (cit. on p. 7).

[22] Gary Cook and David Pomerantz. *Clicking Clean:A Guide to Building the Green Internet*. Tech. rep. Apr. 2015, p. 08. URL: https://www.greenpeace.org/usa/wp-content/uploads/legacy/Global/usa/planet3/PDFs/2015ClickingClean.pdf (cit. on p. 7).

[23] Gary Cook. *Clicking Clean:who is winning the race to build a green internet?* Tech. rep. Jan. 2017, p. 08. URL: https://storage.googleapis.com/planet4-international-stateless/2017/01/35f0ac1a-clickclean2016-hires.pdf (cit. on p. 7).

[24] P. Torcellini, N. Long, and R. Judkoff. *Consumptive Water Use for U.S Power Production*. en. Tech. rep. 2003. URL: http://www.nrel.gov/docs/fy04osti/33905.pdf (cit. on p. 8).

[25] Willie D. Jones. «How Much Water Does It Take to Make Electricity?» In: *IEEE Spectrum Magazine* (Sept. 2012). URL: https://spectrum.ieee.org/energy/environment/how-much-water-does-it-take-to-make-electricity (cit. on p. 9).

[26] URL: https://engineering.fb.com/data-center-engineering/open-sourcing-pue-wue-dashboards/ (cit. on p. 9).

[27] Drew FitzGerald. «Data Centers and Hidden Water Use». In: *Wall Street Journal* (June 2015). URL: https://www.wsj.com/articles/data-centers-1435168386 (cit. on p. 10).

[28] *The 2020 State of Virtualization Technology*. URL: https://www.spiceworks.com/marketing/reports/state-of-virtualization/ (cit. on pp. 10, 11).

[29] Ashish Dandekar. *8 power usage effectiveness (PUE) best practices for your data center*. Sept. 2011. URL: https://www.computerweekly.com/tip/8-power-usage-effectiveness-PUE-best-practices-for-your-data-center (cit. on p. 14).

[30] Matteo Mezzanotte. *What is Immersion Cooling*. Mar. 2020. URL: https://submer.com/blog/what-is-immersion-cooling/ (cit. on p. 14).

[31] URL: https://natick.research.microsoft.com/ (cit. on p. 15).

[32] L. A. Barroso and U. Hölzle. «The Case for Energy-Proportional Computing». In: *Computer* 40.12 (2007) (cit. on pp. 16, 17).

[33] R. Sen and D. A. Wood. «Energy-Proportional Computing: A New Definition». In: *Computer* 50.8 (2017) (cit. on pp. 17–19).

[34] I.E.A. «Demand Response». In: (2020). URL: https://www.iea.org/reports/demand-response (cit. on p. 19).

[35] M. H. Albadi and E. F. El-Saadany. «Demand Response in Electricity Markets: An Overview». In: *2007 IEEE Power Engineering Society General Meeting*. 2007 (cit. on pp. 19–21).

[36] W. McFadden, Anita Nikolich, Ray Parpart, and Birali Runesha. «Saving on Data Center Energy Bills with {EDEALS}: Electricity Demand-response Easy Adjusted Load Shifting». In: 2016 (cit. on p. 21).

[37] Ana Radovanovic. *Our data centers now work harder when the sun shines and wind blows*. 2020. URL: https://blog.google/inside-google/infrastructure/data-centers-work-harder-sun-shines-wind-blows/ (cit. on p. 22).

[38] Google. *Moving toward 24x7 Carbon-Free Energy at Google Data Centers: Progress and Insights*. https://sustainability.google/progress/projects/24x7/. URL: https://storage.googleapis.com/gweb-sustainability.appspot.com/pdf/24x7%20-carbon-free-energy-data-centers.pdf (cit. on p. 23).

[39] D. Laganà, C. Mastroianni, M. Meo, and D. Renga. «Reducing the Operational Cost of Cloud Data Centers through Renewable Energy». In: *Algorithms* 11 (2018), p. 145 (cit. on pp. 25, 26).

[40] David Costenaro and Anthony Duer. «The Megawatts behind Your Megabytes: Going from Data-Center to Desktop». In: 2012 (cit. on p. 64).