Titel der Arbeit

Art der Arbeit

Arbeitsbereich Wissenschaftliches Rechnen
Fachbereich Informatik
Fakultät für Mathematik, Informatik und Naturwissenschaften
Universität Hamburg

Vorgelegt von:	Alexander Leib
E-Mail-Adresse:	2leib@informatik.uni-hamburg.de
Matrikelnummer:	6418345
Studiengang:	Wirtschaftsinformatik
Betreuer:	Reza Heidari

Hamburg, den 31.03.2016

Abstract

While performance is a major and name giving element of HPC development, the high costs associated with operating such facilities lead to cost efficiency as another important aspect driving research in the HPC domain. The enormous electricity requirements of supercomputing centers suggest a high correlation of cost and energy efficiency and introduce the electricity service providers as a party of interest. This paper represents an examination of different means to increase energy efficiency with the focus being on adjusting job-scheduling strategies. The high impact of HPC facilities on the electrical grid can provide an opportunity to engage in mutually beneficial partnerships with the providers. The consideration of particular types of agreements can lead to cost reducing solutions of which some can seem counterintuitive at first.

Contents

1 Introduction **4**

2 Problem: Costs **5**
2.1 Acquisition costs 5
2.2 Operational costs 5

3 Energy Efficiency **7**
3.1 Energy 7
3.2 Energy efficiency in HPC environments 7

4 Cost Efficiency **10**
4.1 Cost reduction potential 10
4.2 Electricity Service Providers (ESP) 10

5 Conclusion **13**

6 Literature **14**

Appendices **15**

List of Figures **16**

1 Introduction

The term “efficiency” is commonly confused with “effectivity” and it is advisable to briefly make the distinction. Effectivity is focused mostly on the result. A method of achieving a goal faster, if more effective than a slower one. Efficiency can be described as a ratio of input and output. Contributing factors to higher efficiency are using less resources, producing less waste and (as with effectivity) reaching a better result. In the context of HPC the term effectivity can be used to describe performance when comparing different HPC centers. The unit of measurement used is floating point operations per second (**FLOPS**). Efficiency can be used to describe different characteristics and should be used with an appropriate reference. Energy efficiency compares performance to power usage and is commonly measured in FLOPS per watt (**FLOPS/W**).

Developments and research concerning HPC centers have traditionally been almost exclusively focused on pushing the boundaries of computing performance capabilities. Since 1993 the famous TOP500 list ranks HPC systems by their performance [6] and it is frequently referred to in the HPC domain.

But in light of climate change and the increasing importance of green IT, energy efficiency has become more important. This trend can be seen by rising number of publications on energy efficiency in HPC. There is even a Green500 list that emerged in 2005 [7] and ranks supercomputers by their energy efficiency. One of the purposes of the Green500 list is to raise awareness for energy efficiency as a desirable quality and “to encourage supercomputing stakeholders to ensure that supercomputers are only simulating climate change and not creating climate change.” [7]

Furthermore, energy efficiency is not only good for the planet, but it also is a contributing factor to cost efficiency. Cost efficiency is not only the focus of this paper, but it also is an important consideration for businesses setting up HPC facilities. Of course cost efficiency is no less important to research facilities (especially, if they are chronically underfunded). In the following chapters, we will first take a look at the problem, which is the cost of HPC centers in Chapter 2 . We will then address the solution by examining energy efficiency in Chapter 3 and how that relates to cost efficiency in Chapter 4.

2 Problem: Costs

Many different entities of science, business, government, etc. may at some point consider setting up and operating large scale high performance computing facilities. The reasons to do so can be numerous, while the counter-arguments tend to be focused on the high costs. Seeing the gigantic amounts of money that are associated with HPC systems can be deterring at first. However, these are not just millions down the drain. HPC environments are investments that can have unexpectedly high returns on investment. That is probably one of the reasons they exist outside the realm of scientific research at all. Nevertheless, reducing the costs of such assets, while minimizing the impact on their performance is worth looking into.

The expenses for HPC facilities can be broken down into two major components: **Acquisition Costs** and **Operational Costs**.

2.1 Acquisition costs

All amounts of money that are spent before a supercomputer is turned on for the first time, can be summarized under as acquisition costs. These expenses constitute a onetime initial investment. Appropriate facilities to house the actual HPC components either have to be acquired or may already be available to use. Furthermore, suitable power supply and data transmission infrastructure has to be provided. This can require expensive new wiring or modifications on existing connections. Moreover, the components for the HPC system itself need to be purchased and the working hours for the assembly have to be paid. Some of these examples may or may not have to be factored into acquisition costs, depending on already existing assets.

A lot of other acquisition costs can accumulate during the setup of an HPC center and the point of this section is not to provide a comprehensive breakdown of all possible expenses. It is rather meant to convey the key characteristics that differentiate acquisition from operational costs.

2.2 Operational costs

The operational costs of an HPC center can be characterized as continuous expenses for operation and maintenance. In contrast to acquisition costs, operational costs are of a recurring or even constant nature. They can include but are not limited to spending for spare components, cooling water, personnel and possibly rent. One of the highest contributions to the sum of operational costs comes from power consumption.

Finding statistically significant numbers on the size of operational costs or their ratio to acquisition costs proves to be quite hard. However, the United States' National Science Foundation offers funding awards for organizations planning to set up HPC facilities. One of the conditions is, that their annual operational costs do not exceed 20% of their acquisition costs [5]. That means that after about 5 years the accumulated operational costs equal the initial investment. It also implies that numerous HPC centers reach that point in a shorter period of time (by exceeding said 20% limit).

3 Energy Efficiency

3.1 Energy

In the previous chapter we addressed the costs involved in HPC centers and the costs for electrical energy were identified as one of the major contributors. Assuming a constant price per kilowatt-hour, reducing energy consumption will lead to lower operational costs. Simply lowering consumption does however not always lead to higher efficiency and energy prices can vary. It seems clear that energy efficiency and costs efficiency do correlate and our next step will be to take a closer look at energy. We will attend to varying prices and cost efficiency in general in the chapter after that.

The terms “power” and “energy” (often in combination with their “consumption”) are frequently used as synonyms. However, they do not mean exactly the same and a distinction between the two has to be made. The exact definitions of the two are less important to us, than their connection to each other.

Power is measured in watts (W) is a ratio of energy per time. Energy is measured in joules and can be calculated by multiplying power with time. The unit kilowatt hours (kWh), which is used by electricity service providers to calculate usage prices is therefore a unit of energy. This means that two electronic devices with different power consumptions can consume the same amount of energy in different time spans. In the Introduction, the unit FLOPS/W was mentioned as a measure of energy efficiency. The unit watt in there can convey the impression of referring to power. But the acronym spelled out means floating point operations per second per watt or simplyfied floating point operations per watt-second. One watt-second equals one joule, making FLOPS/W a unit of energy efficiency.

3.2 Energy efficiency in HPC environments

Having briefly dealt with energy and power, we can now move on to discussing ways of reducing their consumption in the HPC context. There are different approaches to achieving higher energy efficiency ratings and we will now look at a few of them.

3.2.1 Energy efficient hardware

One basic and almost trivial way to increase efficiency is by using hardware that in itself is more energy efficient compared to alternatives with equal performance. However, this method is not applicable to HPC systems that are already in operation. Furthermore, we can assume more energy efficient hardware to be more expensive. This tradeoff needs to be mitigated over the operational life span.

3.2.2 Dynamic Voltage Frequency Scaling (DVFS)

DVFS is based on the idea of slowing down CPUs to adjust to lower workloads. This is accomplished by reducing the clock frequencies through supply voltage, resulting in reduced energy consumption. DVFS can be used to adjust energy consumption proportionally to the system's workload.

3.2.3 Dynamic Power Management (DPM)

DPM is a very effective but aggressive means of power saving ([3] P.5). When the workload is low and only a portion of the available resources is needed to handle it, individual components are deactivated. As soon as the resources are needed again (i.e. the workload is higher) they are subsequently reactivated. The high effectiveness of this dynamic approach is due to reports that "an idle machine consumes about 2/3 of the peak load" ([3], P.5). DPM may not have the desired results, if the alterations between high and low workloads are very frequent. Turning the components on again takes time in which they are not actually doing any work. This causes delays and power wasted on reactivation.

3.2.4 Job Scheduler

Using different algorithms to schedule jobs in an HPC environment can reduce its energy consumption. This was successfully shown by Mämmelä et al. in [2] and we will now take a closer look at the setup and results. Following algorithms were used in the experiment:

First in, first out (FIFO): This simplest of scheduling algorithms uses a basic queue to line up the jobs in the order of their arrival. If there are not enough resources available to handle the first job in the queue, all jobs must wait.

Backfilling first fit (BFF): This algorithm basically works like FIFO, but handles insufficient resources for the first job differently. Instead of having the entire queue wait for resources, the resource requirements of the following jobs are checked. The first job that fits the available resources is then executed.

Backfilling best fit (BBF): BBF can be described as a BFF algorithm that is extended by additional criteria. Instead of scheduling the first possible job, all jobs that can be executed with the given resources are checked. The potential backfill jobs are then matched against said criteria (e.g. shortest or longest job).

By adding a routine to turn off idle nodes, the authors created energy aware versions of these algorithms (**E-FIFO**, **E-BFF** and **E-BBF**). The test results of these 6 job scheduling strategies were then compared and had following results:

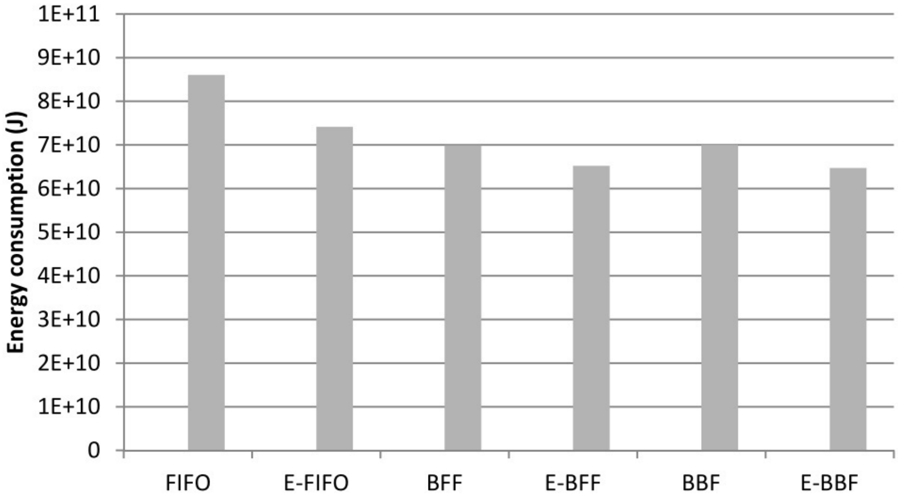

Figure 3.1: Energy Consumption

As we can see in figure 3.1. the energy aware versions did indeed lower energy consumption. It was also demonstrated that FIFO took considerably longer to finish the simulation. “On average, backfilling can decrease the duration by 23% [. . .] compared to FIFO” ([2], P. 8). The findings on the difference in execution speed can be explained by improved resource utilization of the backfilling algorithms. As we can see, the job scheduler offers a good point to dynamically adjust power consumption. By changing its scheduling behaviour, power consumption can be controlled to some extent.

For our purposes the following should be noted: Even though FIFO consumed more **energy** in total, the average **power** consumption was lower.

4 Cost Efficiency

4.1 Cost reduction potential

While potentials to reduce the expenses can be found at of the most factors mentioned under Chapter 2, the real challenge is to save money while simultaneously preventing the loss of performance. Substituting expensive hardware with lower priced alternatives can be expected to involve a tradeoff in quality and consequently reduce efficiency. The same goes for most acquisition costs: Unsuitable locations and facilities can come at lower prices, but the tradeoff rarely outweighs the potential benefits. There may be substantial saving opportunities yet to be identified and future works on this topic might prove successful. However, the focus in this particular paper is on operational costs.

Therefore, a rough breakdown of these expenses should help to identify possible reduction potentials. The costs for spare components, cooling water and personnel offer rather slim opportunities to increase overall efficiency. After all, replacing defective hardware with low cost alternatives is unadvisable and cooling water is simply mandatory. Even worse, cutting or underpaying the operating staff easily leads to losses in efficiency and is also a bad policy from a moral standpoint. As for acquisition costs there might be saving potential here. It has to be evaluated if cost efficiency can be increased by implementing appropriate changes, but discussing those is not the intention here.

We already have addressed energy consumption and possible ways to reduce it. Now we take a closer look at the actual energy costs. They are typically calculated by using a factor that consists of monetary units per kilowatt hour. Simply put, multiplying that factor with the actual consumption (during a specific time span) results in the amount to be paid.

This offers a second approach to reducing energy costs. Instead of consuming less energy, we could also try to use cheaper energy. The prices per kilowatt hour are set by the electricity service providers and therefore cannot directly be influenced. But working together with the providers can result in mutually beneficial agreements and partnerships. In [1] the relationships between HPC centers and electricity providers have been examined and evaluated.

4.2 Electricity Service Providers (ESP)

For most private persons and households, ESPs are not more than organizations that provide electricity and take money according to the usage. Comparing different providers may result in lower prices or clean energy, depending on personal preferences. Switching providers is relatively easy, can be done online and contact to these companies is very limited. The same can be true for big customers from the business or industry sector, that consume high amounts of electrical energy. However, ESPs are not shapeless unapproachable entities. They have their own goals and knowing these can prove very valuable.

/bigskip First of all, ESPs aim “supply efficient and reliable generation, transmission, and distribution of electricity” ([1], P. 2). This includes keeping the electrical grid stable and provide a steady supply of energy. Considering that power consumption is a very dynamic and unpredictable variable, this task seems very hard to achieve. Sudden rises or falls of power demand in the grid can lead to fluctuations that have to be mitigated somehow. The demand for clean, renewable sources of energy can also lead to problems with integrating them into the net. Wind or solar energy depend on external factors and are not always available.

The providers react to these and other problems with different methods like promoting energy efficiency. Customers can be offered financial incentives to switch to more energy efficient electrical devices, for example. Dynamic pricing is another method of regulating overall power usage, where energy prices vary during the time of the day. This approach can be used to shift consumption away from regular times of high demand.

According to [1] “10–50% of electricity costs could potentially be saved“, if HPC centers adapted ESP incentives. But there can be a lot more potential for reducing electricity costs in approaching ESPs and working out individual deals and agreements. By complying to requests from the providers, energy prices could be reduced drastically.

Constant Power Consumption(CPC)

One possibility for an agreement with the ESP is to keep power consumption constant. That means adjusting the power consumption of the facility in such a way, that power fluctuations are kept at a minimum. The result can be a permanent usage of for example roughly 2 MW all day and every day. The advantage for the provider with this method is a constant and predictable consumption and no more fluctuations from that specific HPC center. The ESP can in turn compensate this by lower prices.

Dynamically Adjusted Consumption(DAC)

Another approach for an agreement can come from close cooperation with the provider. By implementing methods of adjusting power consumption dynamically, requests to increase or decrease it during specific times can be fulfilled. Different time frames for this kind of agreement are possible and can reach from weeks to hours in advance. Very quick responses of the HPC centers are thinkable, if necessary capabilities to adjust the consumption are present. Shorter reaction times lead to higher usefulness in reacting to short term and unforeseen fluctuations in the grid as a whole. This usefulness for ESPs in turn can lead to lower energy prices.

Adjusting Consumption As we discussed in Chapter 3 different approaches to lower energy consumption are possible. In order to implement the CPC or DAC approaches it would be possible to utilize more or less energy efficiency methods at a given time. For example, DVFS, DPM or different job scheduling algorithms can be used to adjust power consumption. This means that it can be more cost efficient to turn some of these off at certain times. Other methods like shifting workload to different times can also be used for consumption adjustments.

Current levels of integration

The above considerations require the stakeholders of HPC centers to not just implement means of adjusting consumption. They also have to be willing to sacrifice performance at specific times, to fulfil ESP requests. This is not always the case.

In [1] Bates et al. provided some of the highest ranked supercomputing centers in the USA with a questionnaire to examine this issue further. The questionnaire was designed to find out about total power usage, peak demand and fluctuations as well as current relationships with ESPs.

They found that in the U.S. the interaction between HPC centers and ESPs is very low. About half of the respondents did not have any interaction at all. Those who did, had mostly talked about common incentives offered by ESPs (e.g. dynamic pricing). It also seems that these incentives are not being thought of as interesting by many. In [1] it has been argued that this may be organizational. For example, financial issues are often decided outside the reach of HPC operators.

It was also revealed that complying to ESP requests is not deemed feasible in a financial sense. One respondent wrote that "if a site spent $ 100M for a computer that will remain in production for 60 months, then the apparent benefit [. . .] could easily be outweighed by lost productivity of the consumable resource."

However, there were also responses suggesting that some HPC centers and even whole universities are currently working successfully together with providers.

5 Conclusion

Even though [1] suggests that partnerships between ESPs and HPC sites are rare, advanced integration is possible. Partnerships of this kind combined with energy efficiency methods can have very positive impact on cost efficiency as a whole. The possibilities for integrating HPC sites into the electrical grid are far from being completely explored and a lot of future work on this topic can be successful. It may well be possible to have HPC centers monitor and adjust their power usage on the job scheduler level. This would open the doors wide for new automated integration methods. Developing sophisticated controlling software of this kind further could lead to the point, where the facilities can react to power fluctuations in real-time. Given an appropriate connection to ESPs, the HPC facilities could then be used as resources to stabilize the electrical grid. Turning off energy efficiency mechanisms at times may seem like a bad idea for the environment. However, the fluctuations that come from renewable energy sources could be offset in this case.

The currently low level of cooperation seems unnecessary and the arguments presented against engaging in even just discussions with ESPs seem weak. A lot of potential is wasted by not using the chance to negotiate some form of agreement with the providers. Considering the level of dedication to more energy efficient means of operating HPC environments, the electricity service providers should be included in the advancement of green IT. Hopefully awareness for better communication between these parties will rise in the future.

6 Literature

1 N. Bates et al.: Electrical Grid and Supercomputing Centers: An Investigative Analysis of Emerging Opportunities and Challenges

2 O. Mämmelä et al.: Energy-aware job scheduler for high-performance computing

3 A. Chandio et al.: A comparative study on resource allocation and energy efficient job scheduling strategies in large-scale parallel computing systems

4 D. Yang et al.: Parallel-machine scheduling with controllable processing times and rate-modifying activities to minimise total cost involving total completion time and job compressions

5 http://www.nsf.gov/pubs/2014/nsf14536/nsf14536.pdf

6 http://www.top500.org/project/

7 http://www.green500.org/about

Appendices

List of Figures

3.1 Energy Consumption 9

Eidesstattliche Versicherung

Hiermit versichere ich an Eides statt, dass ich die vorliegende Arbeit im Studiengang XXX selbstständig verfasst und keine anderen als die angegebenen Hilfsmittel – insbesondere keine im Quellenverzeichnis nicht benannten Internet-Quellen – benutzt habe. Alle Stellen, die wörtlich oder sinngemäß aus Veröffentlichungen entnommen wurden, sind als solche kenntlich gemacht. Ich versichere weiterhin, dass ich die Arbeit vorher nicht in einem anderen Prüfungsverfahren eingereicht habe und die eingereichte schriftliche Fassung der auf dem elektronischen Speichermedium entspricht.

Ich bin damit einverstanden, dass meine Abschlussarbeit in den Bestand der Fachbereichsbibliothek eingestellt wird.

Ort, Datum	Unterschrift