EETCO: a tool to Estimate and Explore the implications of datacenter design choices on the TCO and the environmental impact

Damien Hardy*, Isidoros Sideris*, Ali Saidi†, Yiannakis Sazeides*
*University of Cyprus †ARM

***Abstract*—In this paper, we present EETCO an estimation and exploration tool that provides qualitative trends of data center design decisions on Total-Cost-of-Ownership (TCO) and environmental impact. EETCO has the ability to capture the implications of many parameters including server performance, power, cost, and mean-time-to-failure. The tool includes a model for spare estimation needed due to server failures and performance variability. The paper describes the tool model and its implementation, and presents experiments that explore tradeoffs offered by different server configurations, performance variability and ambient temperature. Some observations from these experiments are: servers with different computing performance and power consumption merit exploration to minimize TCO and the environmental impact; performance variability is desirable if it comes with a drastic cost reduction; and increasing by few degrees the ambient datacenter temperature reduces the environmental impact with a minor increase in the TCO.**

I. Introduction

During the last few years, datacenters have increased in numbers, size and uses. In the meantime, different IT infrastructure configurations are proposed in the market, like blade servers and low-power based servers, requiring designers of datacenters to take decisions with diverse cost implications. Consequently, to deliver a cost-efficient datacenter, designers should be aware of how different decisions affect the Total-Cost-of-Ownership (TCO) of a datacenter. Several cost models have been proposed for guiding datacenters design [1], [2], [3], [4], [5]. The following four main factors determine the TCO:

- *Datacenter Acquisition Cost*: the cost of acquisition of the datacenter building (real estate/development of building) and the power and cooling equipment acquisition cost. It is a cost that depreciates within 10-20 years.
- *Datacenter Operating Expenses (OPEX)*: the cost of electricity for servers and cooling.
- *Server Cost Expenses*: the cost of acquiring the servers, which depreciates within 3-4 years.
- *Maintenance & Staff Expenses*: the cost for repairs and the salaries of the personnel.

While the goal of datacenter designers is to minimize the TCO, another major concern is the energy consumption and the resulting environmental impact of such IT infrastructures. The CO_2 footprint is directly linked to the energy consumption, which represents a significant part of the TCO.

Research and commercial efforts are underway to reduce the energy consumption by choosing low-power based servers [6], [7], by reducing the server idle consumption [8] or by reducing the cooling power, which represents a significant part of the Power Usage Effectiveness (PUE).

These trends render essential tools to assess the benefits and drawbacks of datacenter design choices on the TCO and the environmental impact. To the best of our knowledge, only few public tools are available. The APC [9] provides an online estimator tool while [4] provides a spreadsheet to estimate the TCO. Both tools are not defined to allow easy user exploration of fine grain design choices and examine their implications on the TCO and environmental impact.

In this paper, we present EETCO[1] an estimation and exploration tool to provide qualitative trends of datacenter design decisions on TCO and environmental impact. EETCO enables the exploration of the implications of several data center parameters including server performance, power, cost and mean-time-to-failure (MTTF). The tool includes a model that estimates the cold spares needed due to server failures and hot spares due to servers performance variability.

The tool takes as inputs coarse and fine grain data center design choices like PUE, racks organization, components cost, power consumption and MTTF, and produces outputs related to the operation and organization of a datacenter. The tool contains a kernel estimation component that is used by wrappers to explore design decision tradeoffs on TCO and environmental impact.

In the experimental section of the paper, wrappers are defined to explore traditional vs. low-power based servers as well as the implications of performance variability and changing ambient temperature. These experiments reveal the conditions under which servers with different computing performance, power, cost provide opportunity to reduce either or both the TCO and the CO_2 footprint.

The remainder of the paper is organized as follows. Section II presents the framework overview and computation details are given in Section III. Experimental results are given in Section IV. Finally, Section V concludes the paper.

II. Framework overview

EETCO is divided in two parts as illustrated in Figure 1. The first one is the kernel of the tool, which takes as

[1] EETCO is publicly available: http://www2.cs.ucy.ac.cy/carch/xi/eetco.php

inputs a datacenter configuration (land/building acquisition cost, cooling equipment cost per Watt...) and configurations for different types of server modules (processor and other components cost/power/MTTF, rack configurations...). The kernel produces the TCO and environmental impact estimation and other outputs related with the operation and organization of a datacenter. The second part, illustrated by the wrapper exploration, corresponds to a specific wrapper, which generates datacenter and server modules configurations, keep the kernel's results for each configuration and returns the exploration results for which the wrapper is defined. Different wrappers can be defined according to tradeoffs the user wants to explore. For instance, in the experimental results section, wrappers are defined to explore trends of traditional vs. low-power based servers under different performance ratios and to investigate the effects of changing ambient temperature.

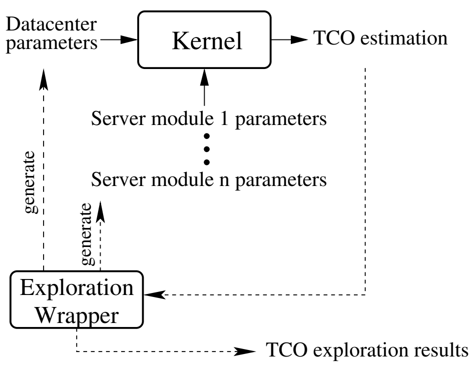

Figure 1. Framework overview

An overview of the kernel framework is shown in Figure 2 and detailed in the next section. For each different server configuration (compute nodes, database nodes, storage nodes...), the estimation starts with a spares estimation for determining the number of hot spares required to reach the number of server modules required for the peak workload and the number of cold spares needed due to server failures. The cold spares are then used to determine the maintenance cost while the hot spares are summed with the required servers to determine the server acquisition cost, the power cost and the datacenter acquisition cost. These costs are then summed to produce the contribution to the TCO of the server configuration and the global TCO is the sum of all server types contribution.

III. TCO ESTIMATION

As shown in the previous section, the TCO estimation is the sum of the datacenter acquisition cost ($C_{acquisition}$), the server acquisition cost (C_{server}), the power cost (C_{power}) and the maintenance cost ($C_{maintenance}$).

$$TCO = C_{acquisition} + C_{server} + C_{power} + C_{maintenance}$$

In this section, the computation model of these different factors is detailed. The following notations are used:

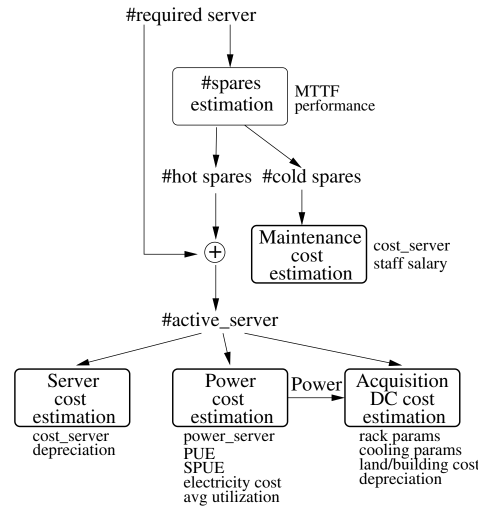

Figure 2. Kernel overview

- N denotes number (number of required server modules, number of spares etc)
- C denotes cost (server module cost, electricity cost etc)
- A denotes area (datacenter area, cooling equipment area, etc)
- K denotes a ratio (server modules per rack etc)
- P denotes power (total server power etc)
- D denotes depreciation (server, data center)

Without loss of generality, a single server configuration is assumed in the following formulas. The resulting TCO of multiple server configurations can be easily determined under the assumption that, the contribution of each server configuration is additive.

In the next subsections, the different computation steps of the estimation are detailed by following the flow described in Figure 2, starting with the spares estimations and followed by the different cost estimations.

A. Hot and cold spares estimation

The distinction between hot and cold spares nodes is necessary, since the hot spares have to be accounted in the power consumption and the cooling requirements whereas the cold spares have to be accounted only in the maintenance cost.

1) Hot spares estimation: In current technology process variations lead to processor performance variations. Thus, in a population of processors, some of them are expected to be affected by a medium/high performance degradation while others will not be affected at all. This performance variation determines the need for hot spares to compensate the performance degradation. For instance, if all cores perform at 90% of their maximum performance and the workload requirements are 10000x throughput (or 10000 cores running separate cloud threads), then we will need (10000/0.9 - 10000) 1111 extra cores to meet our requirements, which translates to extra server costs for acquisition, maintenance and power consumption.

To take into account the performance variation, a *performance variability factor* (PVF) is introduced. PVF takes values from 0 to 1, with 0 meaning no degradation at all and 1 no operation. The performance is thus given by $1 - PVF$. With this factor, the number of hot spares is determined as follows:

$$N_{hotspares} = \frac{N_{srvmodulesreq}}{1 - PVF} - N_{srvmodulesreq}$$

where $N_{srvmodulesreq}$ is the number of server modules required for the peak workload.

In the following, the notation $N_{srvmodules}$ represents the number of active servers. $N_{srvmodules}$ is the sum of $N_{hotspares}$ and $N_{srvmodulesreq}$ modules.

2) Cold spares estimation: Cold spares are server modules needed for replacement when active servers failed. The fault rate of a server can be determined by the MTTF of its components. By assuming a constant fault rate, an exponential distribution can be used to determine the number of cold spares required at a given time t as follows:

$$N_{coldspares}(t) = \frac{N_{srvmodules}}{D_t} - N_{srvmodules}$$

where D_t is defined as follows:

$$D_t = e^{\frac{-t}{MTTF_{allunits}}}$$

According to the exponential distribution, the total MTTF of a server module is obtained using the MTTFs of the components:

$$MTTF_{allunits} = \frac{1}{\sum_i \frac{1}{MTTF_{component_i}}}$$

In the above equations we assume that the server modules with a failure in any component are replaced.

Formula $N_{coldspares}(t)$, assumes that all the replacements are performed at the end of the interval t, which is not accurate. To be more accurate, we can estimate the required cold spare modules in the interval [0,t] by partitioning it in k adequately short time intervals of τ duration ($t = k\tau$), and account for different amount of aging for the newly replaced at the end of each such interval.

By definition, the number of server modules needed for replacement at time 0 is equal to 0 ($CS_0 = 0$) and after the short time interval τ this number is given by:

$$CS_\tau = N_{srvmodules} * (\frac{1}{D_\tau} - 1)$$

Normally, the number of spares for the time interval $[0, 2\tau]$ would be obtained as for the $[0, \tau]$, but since the modules have different age: the CS_τ modules, replaced at time τ, will have age τ and the rest ($N_{srvmodules} - CS_\tau$) will have age 2τ, we have the following formula:

$$CS_{2\tau} = \frac{CS_\tau - CS_0}{D_\tau} + \frac{N_{srvmodules} - CS_\tau}{D_{2\tau}} - N_{srvmodules}$$

And after the elapse of the interval $[0, k\tau]$ the required number of cold spares is given by:

$$CS_{k\tau} = \Big[\sum_{i=1}^{k-1} \frac{CS_{(k-i)\tau} - CS_{(k-i-1)\tau}}{D_{i\tau}}\Big] + \frac{N_{srvmodules} - CS_{(k-1)\tau}}{D_{k\tau}} - N_{srvmodules}$$

By considering $k\tau$ equal to the server depreciation, we obtain the number of cold spares, noted hereafter $N_{coldspares}$, that are considered in the maintenance estimation cost.

B. Cost and environmental impact estimation

The different costs are simply derived from the number of server modules and number of cold spares as explained next.

1) The maintenance cost per month is determined as follows:

$$C_{maintenance} = \frac{N_{coldspares} * C_{srvmodule}}{D_{srv} * 12} + N_{racks} * C_{salaryperrackpermonth}$$

where $C_{srvmodule}$ is the cost of one server module, D_{srv} is the server depreciation in years, $C_{salaryperrackpermonth}$ is the salary cost of datacenter staff per rack per month and N_{racks} is the number of racks determined as follows:

$$N_{racks} = \lceil \frac{N_{srvmodules}}{K_{modulesperrack}} \rceil$$

where $K_{modulesperrack}$ is the number of server modules per rack.

2) The server acquisition cost per month is determined as follows:

$$C_{server} = \frac{N_{srvmodules} * C_{srvmodule}}{D_{srv} * 12}$$

3) The power cost per month is determined as follows:

$$C_{power} = PUE * SPUE * P_{total} * \frac{C_{elecperKWh} * 30 * 24}{1000}$$

where PUE is the power usage effectiveness of the datacenter (the ratio of total power of the datacenter to the IT power), SPUE [10] is the server power usage effectiveness (The ratio of total power of a server to the power of pure electronic components) and $C_{elecperKWh}$ is the electricity cost per KWh.

Finally P_{total} is the total power consumption of all the server modules to consider for the power cost estimation. Depending on how the service provider is charged for the energy they consumed [11]: the peak power consumption or the actual consumption, the peak power (P_{total_peak}) or the average power (P_{total_avg}) have to be used.

$$P_{total\_peak} = N_{srvmodules} * P_{srv\_peak}$$

$$P_{total\_avg} = N_{srvmodules} * (u * P_{srv\_peak} + (1-u) * P_{srv\_idle})$$

where P_{srv_peak} is the peak power consumed by a server, P_{srv_idle} is the power idle consumption of a server and u is the average utilization.

4) The datacenter acquisition cost per month is determined as follows:

$$C_{acquisition} = \frac{C\_building + C_{cooling\_equipment}}{D_{dc} * 12}$$

where $C\_building$ is the land/building acquisition cost, $C_{cooling\_equipment}$ is the cooling equipment cost and D_{dc} is the datacenter depreciation in years.

$$C\_building = A_{perrack} * N_{racks} * K_{coolingarea} * C_{buildingpersqm}$$

where $A_{perrack}$ is the area of one rack, $K_{coolingarea}$ is a factor accounting for more space for the cooling equipment and $C_{buildingpersqm}$ is the cost of land acquisition/building deployment per square meter.

$$C_{cooling\_equipment} = N_{srvmodules} * P_{srv\_peak} * C_{cooling\_eqperW}$$

where $C_{cooling\_eqperW}$ is the cost of cooling infrastructure per Watt.

5) Environmental impact estimation: Conversion factor [12] can be used to translate the KWh consumption into the emission of CO_2 in kg. Thus, the environmental impact per year can be estimate as follows:

$$\frac{P_{total\_avg} * PUE * SPUE * 24 * 365}{1000} * 0.54522$$

IV. Case studies

In this section, some case studies using the EETCO tool are presented. We first describe the experimental assumptions (IV-A) and then we present and analyze results of our experiments (IV-B).

A. Experimental setup

The experiments were conducted using two different server configurations s1 and s2. s1 represents a low-power based server configuration while, s2 represents a traditional server configuration. Table I summaries the datacenter configuration, common server characteristics and the server configurations. The parameters are fixed to representative values from published papers and public industrial data: [10], [4], [2], [1] for the datacenter configuration, [13], [10] for the common server configuration and [14], [15], [16], [8], [17] for the server configurations. The $MTTF_{allunits}$ is computed assuming 1 disk (respectively 2 disks) with 100 years MTTF [17], 4GB DRAM (respectively 8GB DRAM) with 200 years MTTF [16] and one processor with 100 years MTTF [17] for s1 (respectively s2).

parameter	value
$C_{buildingpersqm}$	3000\$/m^2
$C_{cooling\_eqperW}$	15\$/W
$C_{elecperKWh}$	0.07\$
$K_{coolingarea}$	1.2
$C_{salaryperrackpermonth}$	200\$
D_{dc}	15 years
PUE	1.3

1) Data center configuration

parameter	value
Rack	42U
$A_{perrack}$	1.44m^2 (with: 0.6m ; depth 1.2m ; used distance 1.2m)
$K_{modulesperrack}$	84 (6 blades per rack ; 14 modules per blade)
SPUE	1.3
u	0.2
D_{srv}	3 years
τ	1 day
PVF	0

2) Common server configuration

	cost	power	power idle	$MTTF_{allunits}$
s1	500\$	13W	8W	25 years
s2	2500\$	125W	45W	14.29 years

3) Server configurations

Table I
Data Center Common Configuration Parameters

Different cases are studied, in the following to show: the breakdown of the two server configurations, the benefits of our cold spares estimation model, the impact of performance, power and cost between servers, the effect of performance variation and the implications of ambient temperature on the TCO and the environment. For each experiment, 50000 servers are assumed and the peak power consumption (noted peak) and the actual power consumption (noted average) is used to compute the power cost, when it makes a difference.

B. Experimental results

Breakdown of the two server configurations: The normalized breakdown of the TCO for both server configurations is shown in Figure 3.

As shown in the Figure, the server cost represent the most important part of the TCO around 63%-68% for each configuration followed by the maintenance cost (19% s1 and

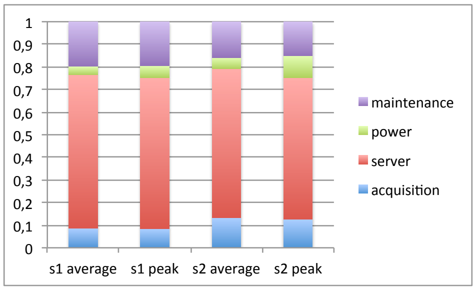

Figure 3. Breakdown of s1 and s2 configurations, average and peak power consumption

15% s2) and the acquisition cost (8% s1 and 13% s2). The power cost differs when the average (3% s1 and 4% s2) and the peak (5% s1 and 9% s2) consumption is assumed. This result is explained by considering the ratio of power consumption and the power consumption at idle time, which is more significant for the s2 configuration.

Benefits of the cold spares estimation model: The benefits of more accurate cold spare estimation is shown in Figure 4 as a function of τ.

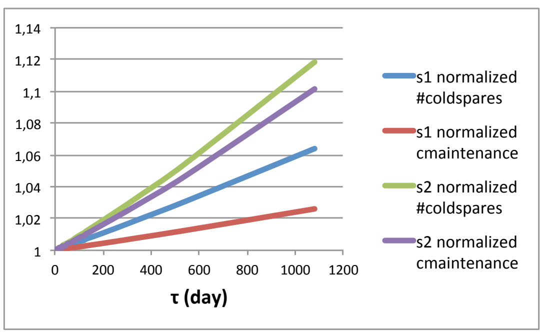

Figure 4. Impact of time step τ on the #coldspares and the maintenance cost for s1 and s2. results normalized with $\tau = 1$ day

As shown in the Figure, taking into account that at any given time not all modules have the same age can reduce significantly the estimated number of cold spares up to 6% for s1 and 12% for s2) and consequently the maintenance cost (up to 2% for s1 and 10% for s2).

Different computing performance between servers: The TCO breakdown is not sufficient to compare the two configurations since they may not have the same computing performance. Let us assume that s1 server configuration will required more units to reach the same computing performance as s2. The equivalent performance coefficient (epc) (defined to be how many s1 servers are required to reach the computing performance of one s2 server) can vary for different configurations, epc is assumed to be from 1 to 8 in this experiment to observe the trends. Results are presented in Figure 5 and the values are normalized with the TCO and the environmental impact obtained with s2.

As shown in the Figure, when epc is relatively small, the TCO obtained with the low-power configuration (s1) is better. At a given point, (epc around 5.5 in our case) the

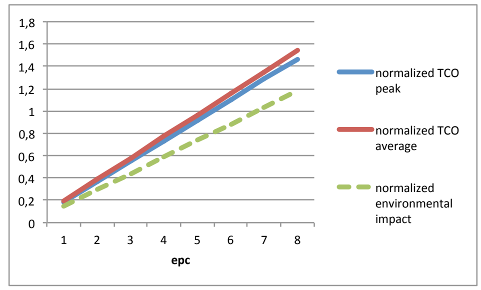

Figure 5. Different computing performance between s1 and s2. s1 results normalized with s2 results. $N_{srvmodulesreq}(s1) = epc * N_{srvmodulesreq}(s2)$

TCO of both configurations are equivalent. Nevertheless, in that case the resulting environmental impact is lower with s1. s1 is thus preferred for the environment for equivalent TCO. After that point, s2 is a better choice for both the TCO and the environment. Awareness of such trend seems can be useful for designing datacenters with reduced TCO and environmental impact.

Impact of performance variation: As discussed in the previous section, process variations lead to processor performance variations. This variation does not only affect the performance but also the processor's cost [18]. In this experiment, a performance variations (PVF) is assumed from 0 to 0.1 and a processor's cost reduction function ($\frac{1}{(1+PVF)^n}$ with n=1, 5 and 10) is applied to the processor's cost with PVF=0, which is assumed to be 1/5 of the server's cost. Results are illustrated in Figure 6 and Figure 7 for s1 and s2 respectively.

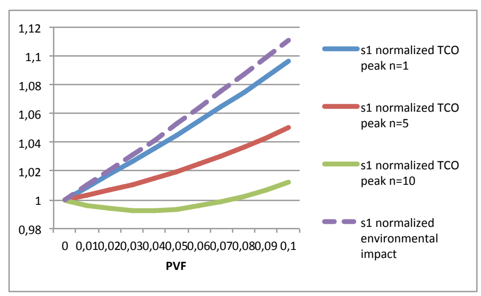

Figure 6. Impact of performance variation, s1 configuration, results normalized with those obtained when PVF=0

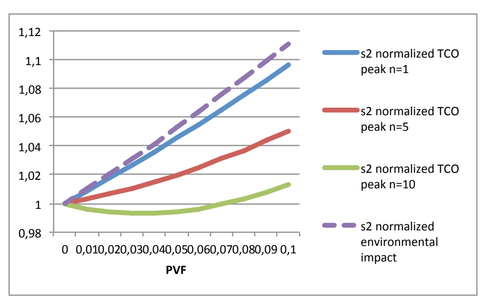

Figure 7. Impact of performance variation, s2 configuration, results normalized with those obtained when PVF=0

As shown in the Figures, trends are similar for s1 and s2. This behavior is explained by the nearly-linear behavior of the TCO according to the number of servers and the same reduction cost ratio. Furthermore, as we can observe, when the processor's reduction is significant (n=10), a TCO reduction for PVF values below 0.07. This positive impact of performance variability comes at the price of an environmental impact increase. In fact, the higher the performance variability, the higher the number of active servers needed, which results inevitably in a higher energy consumption and thus higher CO_2 emissions.

Impact of ambient temperature: The last experiment addressed is the effect of ambient temperature (generally 20°C) on the TCO and the CO_2 emissions. The fact to increase the ambient temperature from 20°C to 30°C has a positive impact on the cooling power consumption (PUE can scale from 2 to 1.65 [19]) while the mean time to failure can be reduced by half [19], [20]. In this experiment, we assess this positive and negative impact by assuming a linear reduction of the PUE and the MTTF per degree and a constant server's power consumption. Finally, two different maintenance server costs are assumed m=1 and m=0.5, representing that the cost of a cold spare is equal to the initial server's cost in the first case and half the price in the second one. The latter is introduced to consider the server's cost reduction over time and the replacement of component instead of changing the server when a failure occurs. Results are illustrated in Figure 8 and Figure 9 for s1 and s2 respectively.

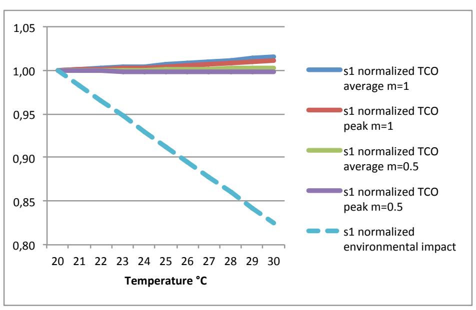

Figure 8. Impact of temperature, s1 configuration, results normalized with those obtained when T=20°C

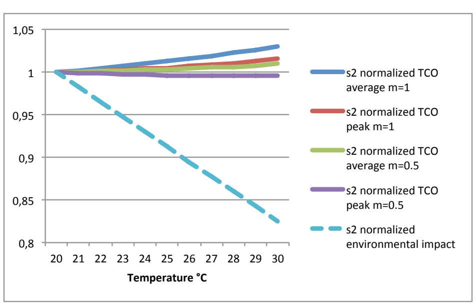

Figure 9. Impact of temperature, s2 configuration, results normalized with those obtained when T=20°C

As shown in the Figures, the CO_2 emissions is significantly reduced for both configurations while we can observe a small TCO overhead when m=1 and a constant TCO or a TCO reduction when m=0.5. In conclusion, increasing by few degrees the ambient datacenter temperature appears to be a good tradeoff to reduce the environmental impact without increasing the TCO.

V. Conclusion

In this paper, we have presented EETCO an estimation and exploration tool to provide qualitative trends of datacenter design decisions on TCO and environmental impact. This tool models the effects of performance variability, estimates cold spares needed due to server failures and captures impact of varying ambient temperature. Different case studies have been performed to assess tradeoffs between server configurations, performance variability and datacenter ambient temperature to better design datacenters with the goal to minimize the TCO and reduce the environmental impact.

Acknowledgment

The research leading to this paper is supported by the European Commission FP7 project "Energy-conscious 3D Server-on-Chip for Green Cloud Services (Project No:247779 "EuroCloud")"

References

[1] C. Patel and A. Shah, "Cost model for planning, development and operation of a data center," HP Laboratories Palo Alto, Tech. Rep., 2005.

[2] J. Karidis, J. E. Moreira, and J. Moreno, "True value: assessing and optimizing the cost of computing at the data center level," in *Proceedings of the 6th ACM conference on Computing frontiers*, ser. CF '09. New York, NY, USA: ACM, 2009, pp. 185–192. [Online]. Available: http://doi.acm.org/10.1145/1531743.1531773

[3] J. Moore, J. Chase, P. Ranganathan, and R. Sharma, "Making scheduling "cool": temperature-aware workload placement in data centers," in *Proceedings of the annual conference on USENIX Annual Technical Conference*, ser. ATEC '05. Berkeley, CA, USA: USENIX Association, 2005, pp. 5–5. [Online]. Available: http://portal.acm.org/citation.cfm?id=1247360.1247365

[4] J. Koomey, K. Brill, P. Turner, J. Stanley, and B. Taylor, "A simple model for determining true total cost of ownership for data centers," 2007.

[5] K. V. Vishwanath, A. Greenberg, and D. A. Reed, "Modular data centers: how to design them?" in *Proceedings of the 1st ACM workshop on Large-Scale system and application performance*, ser. LSAP '09. New York, NY, USA: ACM, 2009, pp. 3–10. [Online]. Available: http://doi.acm.org/10.1145/1552272.1552275

[6] E. Ozer and et al., "Eurocloud: Energy-conscious 3d server-on-chip for green cloud services," in *2nd Workshop on Architectural Concerns in Large Datacenters*, 2010.

[7] Calxeda, "http://www.calxeda.com."

[8] D. Meisner, B. T. Gold, and T. F. Wenisch, “Powernap: eliminating server idle power,” in *Proceeding of the 14th international conference on Architectural support for programming languages and operating systems*, ser. ASPLOS ’09. New York, NY, USA: ACM, 2009, pp. 205–216. [Online]. Available: http://doi.acm.org/10.1145/1508244.1508269

[9] APC, “http://www.apc.com/tools/isx/tco/.”

[10] L. A. Barroso and U. Holzle, “The datacenter as a computer: An introduction to the design of warehouse-scale machines, morgan and claypool publishers, 2009.”

[11] K. Le, J. Zhang, J. Meng, R. Bianchini, T. D. Nguyen, and Y. Jaluria, “Reducing electricity cost through virtual machine placement in high performance computing clouds,” in *Proceedings of Super Computing (SC11), To appear*, 2011.

[12] Defra, “http://archive.defra.gov.uk/environment/business/reporting/pdf/101006-guidelines-ghg-conversion-factors.pdf.”

[13] “Barcelona supercomputing center http://www.bsc.es/plantillaa.php?cat_id=202.”

[14] seamicro, “http://www.seamicro.com/.”

[15] Dell, “http://www.dell.com/.”

[16] Micron, “http://download.micron.com/pdf/technotes/tn0018.pdf.”

[17] K. Bergman, S. Borkar, D. Campbell, W. Carlson, W. Dally, M. Denneau, P. Franzon, W. Harrod, J. Hiller, S. Karp, S. Keckler, D. Klein, R. Lucas, M. Richards, A. Scarpelli, S. Scott, A. Snavely, T. Sterling, R. S. Williams, K. Yelick, K. Bergman, S. Borkar, D. Campbell, W. Carlson, W. Dally, M. Denneau, P. Franzon, W. Harrod, J. Hiller, S. Keckler, D. Klein, P. Kogge, R. S. Williams, and K. Yelick, “Exascale computing study: Technology challenges in achieving exascale systems peter kogge, editor & study lead,” 2008.

[18] A. Das, B. Ozisikyilmaz, S. Ozdemir, G. Memik, J. Zambreno, and A. Choudhary, “Evaluating the effects of cache redundancy on profit,” in *Proceedings of the 41st annual IEEE/ACM International Symposium on Microarchitecture*, ser. MICRO 41. Washington, DC, USA: IEEE Computer Society, 2008, pp. 388–398. [Online]. Available: http://dx.doi.org/10.1109/MICRO.2008.4771807

[19] M. K. Patterson, “The effect of data center temperature on energy efficiency,” in *11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*, 2008, pp. pages 1167–1174.

[20] “Western digital wdc.ph/wdproducts/library/other/2579-001134.pdf.”