Proceedings Paper:
Garraghan, P, Al-Anii, Y orcid.org/0000-0002-4382-1998, Summers, J orcid.org/0000-0001-8266-5038 et al. (3 more authors) (2016) A unified model for holistic power usage in cloud datacenter servers. In: UCC '16: Proceedings of the 9th International Conference on Utility and Cloud Computing. 9th IEEE/ACM International Conference on Utility and Cloud Computing (UCC 2016), 06-09 Dec 2016, Tongji University – Shanghai, China. Association for Computing Machinery , pp. 11-19. ISBN 978-1-4503-4616-0

https://doi.org/10.1145/2996890.2996896

A Unified Model for Holistic Power Usage in Cloud Datacenter Servers

Peter Garraghan[1], Yaser Al-Anii[2], Jon Summers[2], Harvey Thompson[2], Nik Kapur[2], Karim Djemame[3]

School of Computing & Communications[1]
Lancaster University, UK
p.garraghan@lancaster.ac.uk

School of Mechanical Engineering[2]
University of Leeds, UK
{mnytaa, j.l.summers, h.m.thompson, n.kapur} @leeds.ac.uk

School of Computing[3]
University of Leeds, UK
k.djemame @leeds.ac.uk

Abstract — **Cloud datacenters are compute facilities formed by hundreds and thousands of heterogeneous servers requiring significant power requirements to operate effectively. Servers are composed by multiple interacting sub-systems including applications, microelectronic processors, and cooling which reflect their respective power profiles via different parameters. What is presently unknown is how to accurately model the holistic power usage of the entire server when including all these sub-systems together. This becomes increasingly challenging when considering diverse utilization patterns, server hardware characteristics, air and liquid cooling techniques, and importantly quantifying the non-electrical energy cost imposed by cooling operation. Such a challenge arises due to the need for multi-disciplinary expertise required to study server operation holistically. This work provides a unified model for capturing holistic power usage within Cloud datacenter servers. Constructed through controlled laboratory experiments, the model captures the relationship of server power usage between software, hardware, and cooling agnostic of architecture and cooling type (air and liquid). An exciting prospect is the ability to quantify the amount of non-electrical power consumed through cooling, allowing for more realistic and accurate server power profiles. This work represents the first empirically supported analysis and modeling of holistic power usage for Cloud datacenter servers, and bridges a significant gap between computer science and mechanical engineering research. Model validation through experiments demonstrates an average standard error of 3% for server power usage within both air and liquid cooled environments.**

Keywords- Cloud datacenter, holistic energy, server power modeling.

I. INTRODUCTION

Cloud datacenters form the backbone of modern Internet infrastructure globally, and are critical for provisioning digital services to consumers. These systems are driven by diverse user behavior and are composed by numerous interacting physical (servers, cooling), and virtual (applications, resource schedulers) sub-systems. Datacenters require vast amounts of compute and storage power to facilitate Internet-scale workload, and subsequently consume 1.8% electrical energy globally [1].

Datacenters suffer numerous challenges towards achieving high energy-efficiency, stemming from cooling load [2], failures [3], and underutilization as low as 10%-25% [4], with idle datacenters (i.e. servers running with minimal usage) consuming almost half of their peak power [5]. Suppliers to the datacenter industry are starting to address aspects of energy-efficiency through layout for efficient cooling [7], efficient IT components and workload scheduling [8], consolidation, and resource throttling [9]. However, datacenters are mission critical facilities with implicit requirements including reliability, capacity, as well as explicit requirements including availability enforced through Service Level Agreements (SLAs). As a result Cloud datacenter providers are reluctant to deploy energy–efficiency mechanisms without precisely understanding its implication towards operational performance and business objectives.

In order to achieve such an understanding, it is imperative to study holistic power profiles of servers – specifically the co-relation between electrical and thermal energy produced in both hardware and cooling under different operational conditions driven by software utilization. Servers alone consume 29-31% of the total datacenter energy, requiring an additional 34-38% in facility-level cooling for heat removal from the facility [10], [11]. However, ascertaining such knowledge is challenging due to the complexity in identifying and analyzing key parameters within each sub-system, and importantly their interrelation (i.e. an increase in software resource utilization results in higher power consumption, resulting in higher chip temperature, thus requiring more server cooling load).

This problem is inherently multi-disciplinary in nature due to the diversity of solutions which requires in-depth knowledge and understanding of the both virtual (software stacks) as well as physical (IT hardware, cooling) systems. Work towards energy-efficient datacenters primarily focuses on either the perspective of computing (scheduling, workload management, software, etc.) [14-16] or mechanical (cooling systems, UPS, layout, etc.) [2],[16], [18],[19] This has resulted in knowledge gaps as depicted in Fig. 1 towards understanding explicitly the holistic power profiles and sub-system interaction across both the virtual and physical layers within servers. As a result, there is a strong need to conduct experiments in controlled laboratory environments to study the holistic and correlative power profiles of server sub-systems.

Determining this relationship and their respective power profiles would enable the creation of a unified model for holistic

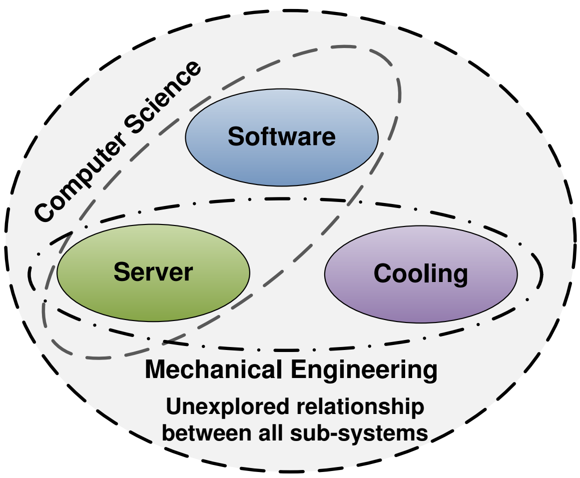

Fig. 1. Knowledge gap in energy-efficient Cloud datacenters across disciplines.

power usage for Cloud datacenter servers. The effectiveness of such a model would require accurately capturing power profiles of all sub-systems agnostic of hardware specification and cooling type, enabling Cloud researchers and engineers to determine more accurate power profiles for computing servers when considering the power draw from its internal cooling capability.

The objective of this work is to analyze and model holistic power usage within Cloud datacenter servers. Specifically, we propose a unified model comprising software utilization, server architecture and cooling type and how each contribute towards the total power usage of the server under different operational conditions. To our knowledge, this is the first endeavor towards empirically studying and modeling holistic Cloud datacenter server power profiles including cooling properties. This work provides a step change in bridging the gap between Computer Science and Mechanical Engineering – research disciplines which have closely aligned interested towards energy-efficient Cloud datacenters. Contributions are summarized as follows:

Analysis of Cloud datacenter server sub-system power profiles. Through numerous experiments conducted we investigate the operational profile of the entire server including software utilization, hardware and cooling by extracting parameters including performance, resource usage, core temperature, and power consumption for hardware and cooling.

Unifying model for Cloud datacenter server power usage. We propose an empirically validated statistical model for capturing sub-system power profiles and their co-relation that operates agnostic of server hardware and cooling type. Importantly, in addition to electrical power of server operation, we are able to capture the thermal energy rejected by cooling fans and/or pumps. Furthermore, we present a number of applications of the analysis findings and proposed model.

Section 2 provides background of this work; Section 3 discusses related work; Section 4 presents the experiment methodology; Section 5 presents the analysis findings; Section 6 constructs the holistic power model; Section 7 presents the model validation; Section 8 discusses model application; Section 9 details conclusions and future work.

II. BACKGROUND

A. The Holistic Energy Chain

Cloud datacenter servers (and by extension the greater facility) are formed by numerous interacting sub-systems. These sub-systems exhibit both implicit and explicit interactions with respect to their operational characteristics. As a result, a logical deduction is that changes to a particular sub-system will impact other sub-systems within the server.

With respect to energy usage, as the user drives the operational characteristics and resource usage of the application software in the datacenter [6], it is intuitive to assume that alterations to application operation will impose a cascading affect throughout the entire system. As shown in Fig. 2, software comprises specific characteristics of energy cost (computation per watt) and performance (MIPS). This results in the generation of heat on the CPU chip that requires cooling for heat rejection from the motherboard.

While it is possible to directly measure each sub-system and produce the sum of parameters, calculating the total power consumption is challenging when a specific sub-system changes.

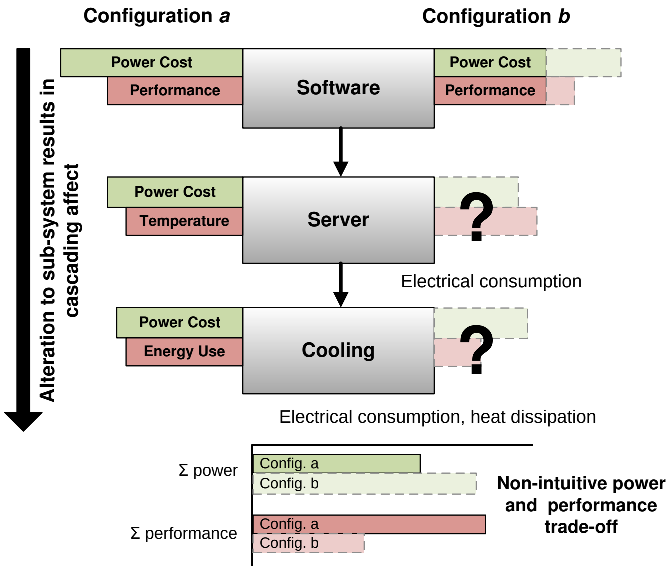

Fig. 2. Depiction of Energy Chain in Cloud Datacenters.

For example, if the software throughput increases, it is not obvious how this effect cascades throughout the server. This is particularly important when considering sub-systems may not necessarily follow a linear relationship for resource utilization and energy use [20], and that servers can be composed of air or liquid cooling reflecting different power profiles. This becomes even more challenging outside the boundary of a single server and within context of the entire facility. As a result, it is critical that not only is the holistic power usage of a datacenter is quantified, but more importantly how alterations within a specific sub-system impact such parameters.

III. RELATED WORK

Power consumption of servers can be measured, estimated, or a combination of both. Measurements include devices that capture the electrical power of a server, while estimation are parameters that are inferred by correlated parameters (i.e. CPU utilization) [14]. There are numerous works that model sever power consumption comprising processors, VMs and servers.

Works such as [21-23] quantify the datacenter power consumption of a server by measuring the voltage drop of resistors inserted on power rails across the server motherboard. As detailed in [14], direct measurement of datacenter servers requires precise insight into the mainboard layout so that supply lines are assigned to the correct subsystem, and relevant to this work are unable to capture subsystem power consumption that cannot be directly measured. Numerous works study and model the power profiles of computing servers.

Heath et al. [24] propose an energy model for heterogeneous server clusters, where power consumption of an individual server is estimated using a linear model constructed from resource utilization parameters. The model is evaluated through simulation from trace requests from 1998 World Cup trace data, demonstrating an average error estimation of 1.3%.

Economou et al. [25] propose a linear model to estimate server power consumption using the input parameters of CPU, memory, disk I/O and network rate. Through the use of custom benchmarks, they stress individual motherboard subsystems and measure their respective power profile. The proposed power models were validated through experimentation, demonstrating model error between 0 – 15% for a blade server.

Fan et al. [26] propose a nonlinear power model using CPU utilization at the input, and include an error correction factor determined based on system characteristics learnt during calibration. Model validation was performed on several hundred servers, reporting an average estimation error less than 1% on average.

Harton et al. [27] propose a real-time power consumption model for a software defined datacenter. Their model requires inputs from utilization and electrical power from the Power Supply Unit. By using a MIMO/MISO model and machine learning regression techniques, they model server power consumption. They demonstrate through experiments that their model is capable of accurately predicting server power utilization within 5% margin of error.

While each of these works provide accurate measurements and model validation of server electrical power consumption and validated models for server power consumption of numerous hardware components, none of them are able to quantify non-electrical power consumed due to cooling requirements of the server itself.

Beitelmal and Fabris [28] presented new IT efficiency metric based on a thermodynamic approach. This approach defined the ratio between IT work, which is the outcome of running the server, and the total power including cooling load. This metric restricted the efficiency of ICT, however indicated the need to potentially redefine the current metric standard for datacenters.

A study of thermal effects of servers was analyzed by Sampath [20] and Pandiyan [33]. The analysis included measuring server component response to different levels of system load. The research focused on ambient temperature effects on power consumption and related cooling techniques. CFD models were used to investigate the feasibility of a power consumption prediction formula predict empirical formula for power consumption. All simulation results were validated with experimental work, and the error percentages were less than 10%. They demonstrate a linear relationship between power consumption and processor utilization within their experiments. This results could be used to extrapolate energy-efficient studies at the datacentre level.

Investigation of chiller-less cooling technology of datacenters has been conducted intensively [29], [32], [34]. A liquid-cooled server with an economizer based system was used to better understand the sectors of power through the datacenter. Results showed an excellent saving of energy of 25% when using energy-centric configuration for cooling. Further to these studies, David et al [30] investigated the operation conditions and scenarios effects on overall power consumption of liquid-cooled rack of servers. Ham et al [31] presented an investigation of a simplified model to model datacenter cooling and energy consumption. The model focused on the effects of interior thermal management techniques for reducing energy bills.

These works present a gap between studying the power utilization of servers and its respective cooling in a unified manner. Towards highlighting this challenges, there has been concentrated efforts towards studying datacenter power holistically.

Shoukourrian, et al. [12] propose an evaluation toolset PowerDAM capable for a unified collection of energy profiles from HPC datacenter sub-systems. The system uses remote scripts to collect input data from system monitoring tools, resource management, and sensor data. They demonstrate that applying their approach within the LRZ HPC system, computation and cooling constitutes 84% and 4% of the total consume energy, respectively per user.

Pelly et al. [13] presents an analytics framework for modeling total datacenter power towards understanding and abstracting total datacenter power. They integrate numerous parametric power models for datacenter component simulation including servers, chiller plants and cooling towers. They provide a case study for hypothetical datacenter and propose a technique for intelligent cooling management.

IV. METHODOLOGY

A. Approach

In order to construct a unifying model for Cloud datacenter server power profiles, it is first necessary to study the operational characteristics of the server holistically. This is performed by capturing parameters across all sub-systems. These parameters are categorized as virtual (i.e. CPU, application throughput), and physical (i.e. server power consumption, hardware temperature) as summarized in Table 1. As these sub-systems focus on a specific subset of system operation, they each entail a bespoke technique for parameter extraction ranging from built-in functions and external monitoring devices, and values derived from empirically validated mathematical models.

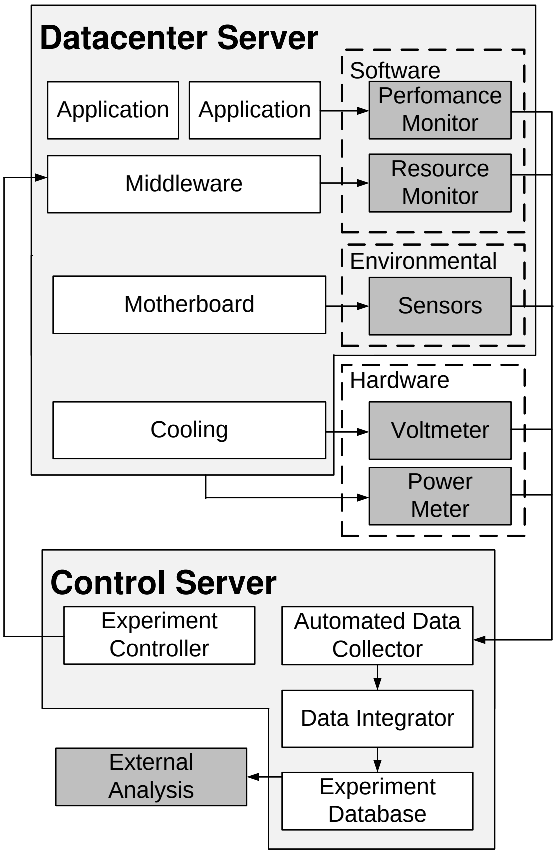

Fig. 3. Architecture for parameter extraction

Table 1. Extracted server sub-system parameters.

Category	Sub-system	Parameter
Virtual	Application	Throughput
	Middleware	CPU utilization
Physical	Server Hardware	Power (W)
	Server Cooling	Voltage (V) & Current (A)
	CPU Core	Temperature (ºC)
	Environment	Ambient Temperature (ºC)

Fig. 3 provides an overview of sub-system parameter extraction and data collection. Parameters from numerous sub-systems are transmitted to the Control Server. The Control Server is responsible for integrating unstructured data files detailing recorded parameters into CSV files with their respective timestamps. These CSV files are then stored in a database which can be queried for conducting analysis.

B. Experiment Setup

We conducted controlled laboratory experiments shown in Fig. 4 to study sub-system operation of a Sun Fire V20z, 2AMD Operton processor, 64-bit x86, 1GBx8, running Debian Ubuntu. We executed a generic three-tier web application to drive utilization. The three-tier web application is composed of a set of VM images, with its load balancer image implemented via HAProxy [36] that distributes load between application servers. These application servers are comprised of a single JBoss web container running within a Java VM and a pre-installed photo album application. The photo album application stores and retrieves data within a single MySQL image. The application enables to control the performance and throughput of the applications, creating additional threads thus increased server CPU utilization.

Metrics were collected using the method illustrated in Fig. 3. Specifically, application performance was collected from the generation of system log files detailing application performance at fixed time intervals. Server resource usage was collected through the top command within Linux. Server temperature was determined by using the PCMI command to ls-sensors, and server electrical power collected using a Vol-tech 9000 meter.

We also collected the power consumption of cooling equipment within the servers. The installed cooling pump is a CoolIT ECO III – 120 DCLC, and its power was measured from its manual specification. On the other hand, while the server has

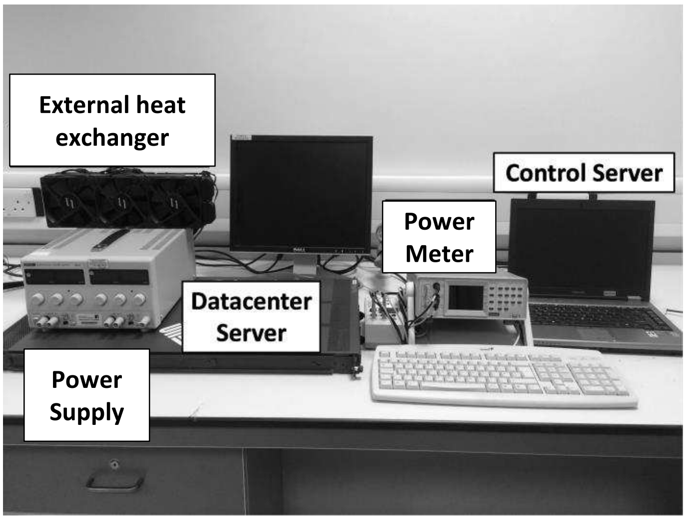

Fig. 4. Experiment layout for datacenter server sub-system parameter extraction.

a built-in tachometer sensor to capture fan speed, the amount of power drawn by these fans is still unknown. To ascertain this, samples of Delta FFB0412SHN were tested in order to correlate speed (rpm) and electrical power (W). We studied the operation of 3 x DELTA FFB 40x40x28mm series fans. The test procedure included powering the fan using a controllable power supply within the ranges of 4.48 – 13.83V and measuring the power consumption at each level. Simultaneously, an iParaAiluRy digital tachometer is used to record the rotational speed of the fan at each level. This digital tachometer sends a laser beam to the blade of the fan and capture the reflection using a silver light-reflective sticker attached to a fan blade. Depending on the number of laser signal detected indicates the revolution per minute, the resolution of the tachometer 1 rpm is capture with an error rate of ±0.05%.

Fig. 5 shows the correlation between fan speed N and power P_{fan} using the three cases, where it is observable that the correlation follows a polynomial function. As a result, we are able to capture this correlation as

$$P_{fan} = f_1 N - f_2 N^2 + f_3 N^3 \quad (1)$$

In order to determine the correlation parameters f_1, f_2 and f_3, the collected data were analyzed using non-linear regression. The estimation is based on Gauss-Newton method. This method starts with an initial approximation of values and then performs linearization around this selected value. This will require minimizer function Γ in order to achieve convergence as shown in Equation 2.

$$\Gamma = \sum_{i=0}^{i_{max}} \left(r_i\left(x^{(k)}\right) + \nabla r_i\left(x^{(k)}\right)^T \left(x - x^{(k)}\right) \right)^2 \quad (2)$$

Where k is the number of current iteration, x is the values of current approximation, and r is the evaluated function in matrix form at the assumed guess values. Once convergence occurs, after a number of iterations, the empirical constraints are ready to compute as shown in Table 2, with a proportion of variance (R = 0.993).

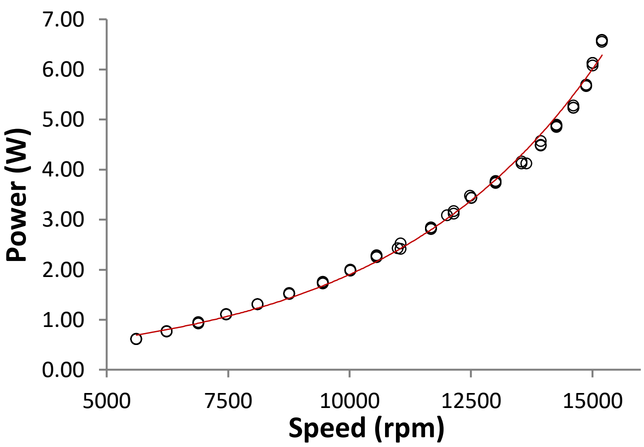

Fig. 5. Fan power measurement and modeling.

Table 2. Empirical constant for the fan formula.

Constant	f_1	f_2	f_3
Value	2.46e-04	– 3.70e-08	3.11e-12

Table 3. Datacenter Sub-system Parameters.

	Virtual		Physical					
Case	**Software perf.**	**CPU**	**Power (W)**	**Core 0 temp. (º C)**	**Core 1 temp. (ºC)**	**Ambient temp. (º C)**	**Fan (W)**	**Pump (W)**
A	0.00	0.00	136.00	37.15	35.63	18.04	1.37	2.88
B	48.42	20.00	152.99	37.26	35.86	18.17	1.55	2.88
C	97.22	41.43	162.97	38.49	36.43	18.13	2.32	2.88
D	183.98	59.15	184.25	41.59	40.09	18.11	2.98	2.88
E	228.78	78.65	206.17	43.49	41.42	18.01	3.78	2.88
F	250.08	89.21	211.42	43.85	41.90	18.07	3.90	2.88
G	265.46	100.00	222.52	45.19	43.54	18.17	4.62	2.88

This equation can be used to directly map the fan speed of the server with its corresponding electrical output.

In order to create different operational conditions for experiments, we used cpu-limit to enable the ability to control the CPU utilization of the server. Each experiment was conducted 10 times for a period of 5 minutes, resulting in a total of 70 experiment cases. All sub-system metrics after experiments were automatically transferred to a Windows machine via bash scripts and scp for conducting data analysis using RStudio.

V. Analysis

Table 3 summarizes collected parameters across all experiment cases. It is observable that there exists a strong correlation between application throughput and all respective parameters indicated by a Pearson correlation value > 0.9. This result is intuitive given that users drive resource usage, and subsequent power and temperature profiles within the server. This is demonstrated by varying levels of CPU and server power consumption for each experiment case as shown in Fig.s 6(a) and 6(b), respectively.

While all parameters with the exception of ambient temperature exhibits strong positive correlation, this relationship is not strictly linear as shown in Fig. 7 depicting change of parameters from idle server operation (i.e. no workload execution). Temperature of both CPU cores increases dramatically from 37.15ºC at 40% utilization to 45.19ºC at 100% utilization for Core 0. Furthermore, we observe that server

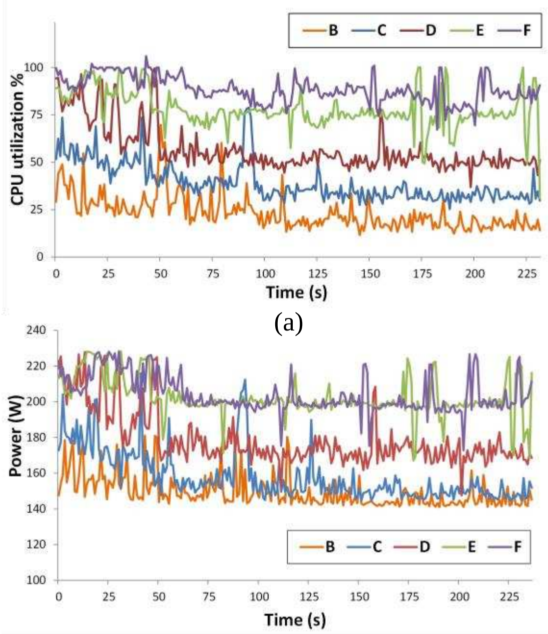

Fig. 6. Datacenter server experiment cases (a) CPU utilization, (b) Power.

power usage and application performance is not perfectly linear – with performance degradation occurring at 80%+ utilization.

Altering the system utilization levels of a server at different intensity results in a non-proportional change in corresponding parameters within the server in terms of performance, temperature and power consumption. This result indicates that alteration to software efficiency (i.e. reduced resource usage or increased throughput) results in different changes with respect to power. This is particularly noticeable when studying the change in utilization at step intervals as illustrated in Fig. 8. While these changes might appear minimal within a single server, such behavior becomes increasingly important within the context of large-scale systems composed of hundreds and thousands of servers over extended periods of time.

The total electrical power of cooling systems consumed comprises cold plate pumps and fans. The cold plate pumps, require 2.88W to operate across all experiment cases. As mentioned previously in Section 4, the fan speed is correlated to

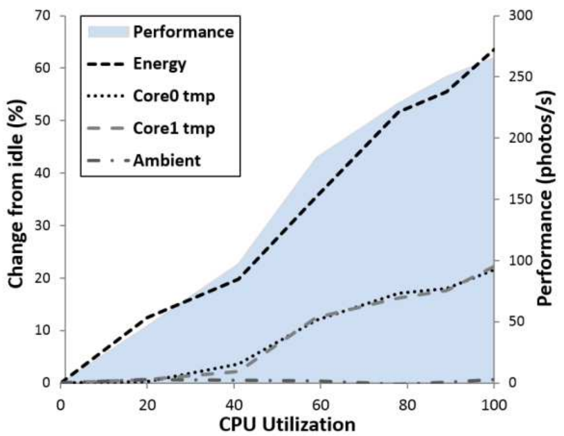

Fig. 7. Change in sub-system parameters from idle.

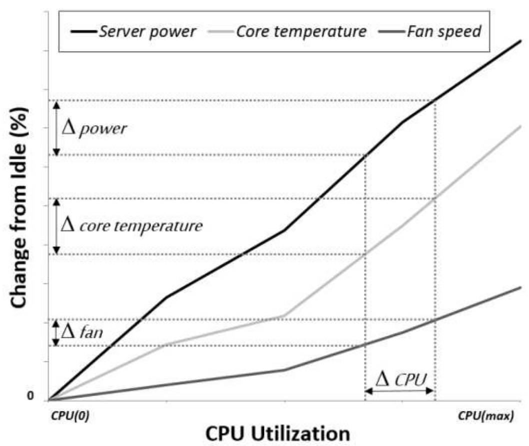

Fig. 8. Illustration of disproportionate parameter change.

the power consumption. Therefore, the instantaneous power consumption of the fans can be recorded based on monitoring of fan speed. The average power consumption by fans is 3.81W (at 100% utilization). On average, the cooling corresponds to 3.07% of the datacenter server electrical power consumption.

This analysis demonstrates three items of interest. First, while there exists a positive correlation between an increase in CPU utilization and sub-system operation, this does not reflect a strictly linear within sub-system operation. Secondly, power consumed by the cooling systems themselves represent a non-negligible amount of power consumption. Thirdly, such analysis and measurement of sub-system operation does not capture the non-electrical power expended for server cooling. Such observations highlight that there is an opportunity to further energy-efficient research in Computer Science by the inclusion of additional power to provide more realistic power profiles. Even if such cooling represents a small amount of additional power, such increments result in large power profiles consumed within the context of tens of thousands of datacenter servers.

VI. Modeling

A. Construction

We propose a unified model for capturing the relationship between resource utilization, microelectronic processor power, and cooling comprising electrical and non-electrical power.

The research focuses on the relationship between thermal characteristics of servers and the corresponding reflect on the cooling system. The first part of model construction entails sectoring microelectronic processor power into three categories: Idle power, dynamic power, and static power. Idle power P_{idle} represents the power consumed by the motherboard and server processor when it is idle (i.e. ~ 0% CPU utilization). Dynamic power $P_{dynamic}$ is the server power consumption driven by an increase in CPU resource utilization u. Static power P_{static} is the power consumed by the CPU driven by its die temperature T (i.e. physical temperature of the core). Each of these are interrelated to one another, and although their respective relationship appears intuitive, it is derived from complex electronics with static leakage power of CPU consumption (solely dependent on temperature and voltage) and a physics phenomenon known as the Poole-Frenkel effect.

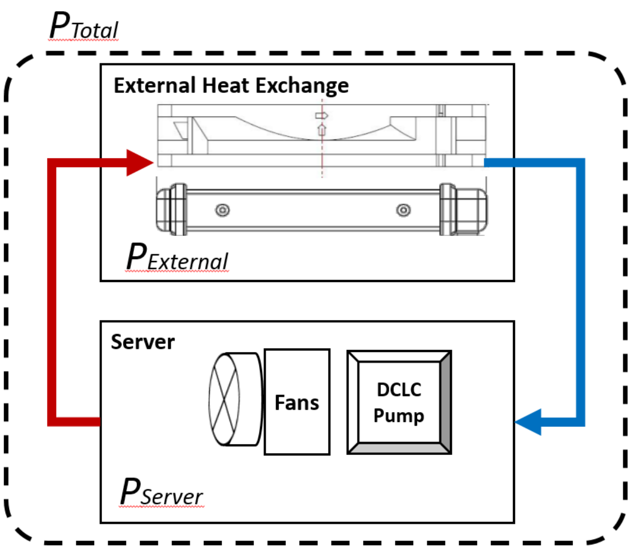

Fig. 9. Illustration of P_{total} breakdown.

Each of these categories together form the total power consumption consumed by IT processes P_{IT}.

$$P_{IT} = P_{idle} + P_{dynamic} + P_{static} \tag{3}$$

$$P_{dynamic} = a_1 u \tag{4}$$

$$P_{static} = a_2 T_{CPU} + a_3 T_{CPU}^2 \tag{5}$$

The constants a_1, a_2, a_3 for dynamic and static power are calculated based on the method detailed Equation (2) in Section 4 producing parameters in Table 4 with an R value of 0.983 indicating high accuracy.

Table 4. Empirical constants for the power formula

Constant	a_1	a_2	a_3
Value	0.7605	0.05	0.005

Datacenter server power consumption additionally comprises the operation of pumps P_{pump} and/or fans P_{fan} for cooling. This results in the total power consumption of the entire server P_{Server} as

$$P_{Server} = P_{IT} + \sum_{1}^{m} m\, P_{pump} + \sum_{1}^{n} n\, P_{fan} \tag{6}$$

where m and n represent the total number of pumps and fan within the server, respectively. This summation of cooling components allows for the model to capture the power profiles of server cooling comprising liquid, air or a combination of both.

The total power drawn by a server P_{Total} includes both power consumed by the microelectronics components, P_{Server}, and power losses in the power supply units PSUs, P_{PSU}. It is possible for PSUs to exhibit varied operational efficiencies (i.e. the ratio between input and output of electricity).

$$\sum_{1}^{l} l\,(1-\eta) P_{PSU} \tag{7}$$

where η represents the power supply unit efficiency percentage and l represents the number of PSU used to power the server.

While P_{Total} represents the consumed power by server microelectronic components, such components require cooling power internal and external of the server to function. Internally, there exist two types of cooling: (1) pumps which circulate water inside a cold-plate to collect heat from CPUs, and (2) fans which move colder air to pass over microelectronic components. The hot working fluid requires mechanical work to reject heat and back to the supply condition. This is performed by the heat exchanger and fans in conjunction, and represents the external power $P_{external}$. As shown in Fig. 9, $P_{external}$ represents the amount of power expended to achieve temperature homeostasis between the input and output temperature of the server (heat transfer). As a result, the summation of internal and external power consumption of cooling $P_{cooling}$ forms

$$P_{cooling} = \sum_{1}^{m} m\, P_{pump} + \sum_{1}^{n} n\, P_{fan} + P_{external} \tag{8}$$

Thus, the total power consumption of a datacenter when considering all sub-systems holistically is represented as:

$$P_{Total} = P_{IT} + \sum_{1}^{l} l\,(1-\eta)P_{PSU} + P_{cooling} \quad (9)$$

With this equation and its respective parts, it is possible to calculate the partial power utilization effectiveness (pPUE) of the server. The pPUE is defined as the ratio of power consumed within the IT and cooling, and aims to determine the effectiveness of server power usage. We have presented two methods to measure pPUE: exclusion and inclusion of internal cooling as shown in Equation 10 and 11, respectively.

$$pPUE_I = \frac{P_{IT} + P_{cooling}}{P_{IT}} \quad (10)$$

$$pPUE_{II} = \frac{P_{Server} + P_{external}}{P_{Server}} \quad (11)$$

The reason for providing two measurements for pPUE is due to limitations of knowledge pertaining to server operation. Feasibly, if the specification for fan and pump characteristics are unavailable or immeasurable it is not possible to fully construct a unified model. As a result, $pPUE_I$ assumes full knowledge of sub-system operation and provides a higher degree of accuracy. $pPUE_{II}$ instead measures the temperature difference of the server ($P_{external}$), treating the internal cooling architecture as a black box.

B. Validation

In order to validate our proposed unifying model for server power consumption, we conducted numerous experiments to study model accuracy.

The experimental setup used for model validation is a Sun Fire V20z, 2AMD Operton processor server using (1) air-cooling only, and (2) liquid-cooling only. This is important as it drives different cooling behavior within the system (for example, the operation temperature directly effects the power consumption of the CPU cores). The temperature of supply and return water from the external heat exchanger were recorded at each level of the experiment. The recorded water temperature assists in determining the precise calculation of predicted power. In addition, it is required to measure the flow rate of the pumps, which is captured by timing the collection period of 300ml of water and dividing the volume by time.

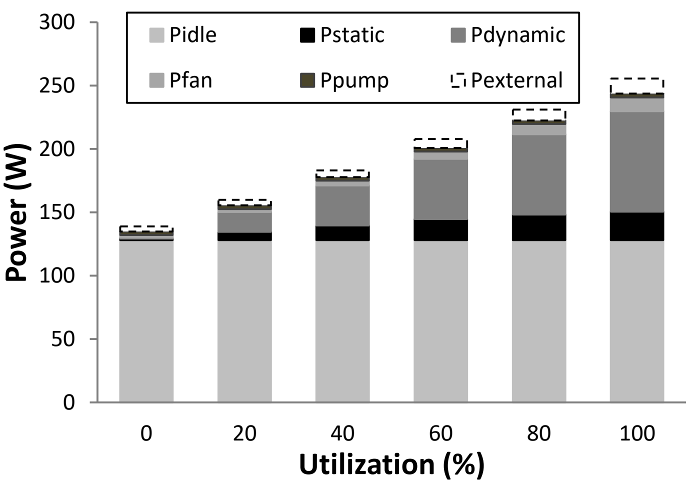

Fig. 12. Breakdown of server sub-system power usage for liquid cooling.

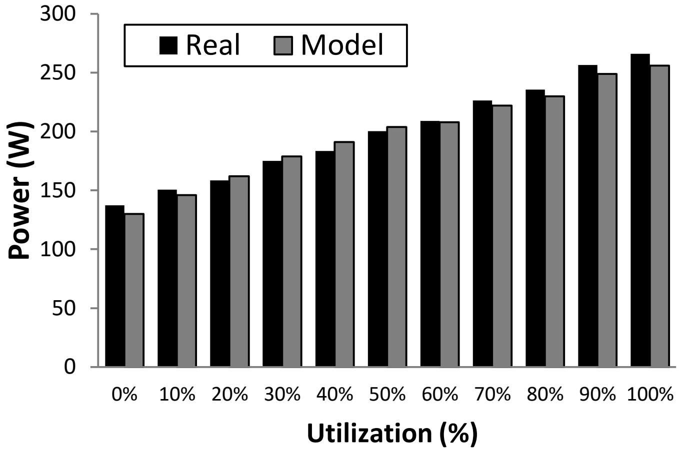

Fig. 10. Model accuracy vs. experiments with air cooling.

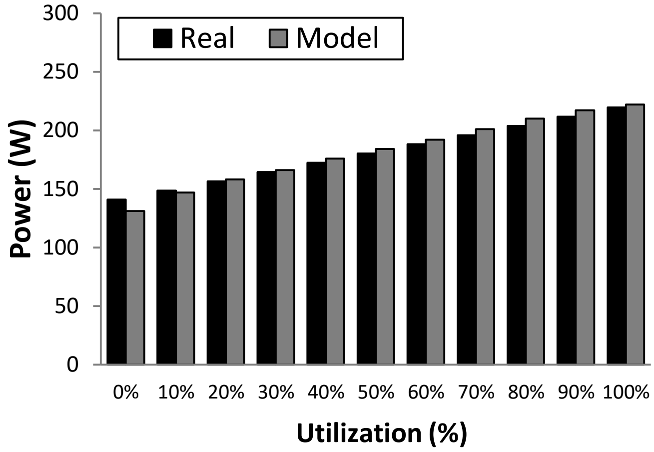

Fig. 11. Model accuracy vs. experiments with liquid cooling.

We used the SPECPower benchmark [35] in order to impose workload to control the server utilization between 0 – 100% at 10% increments. At each level of utilization level the server is monitored and both the CPU die temperature and fan speed are recorded. The recorded data is analyzed, averaged and fitted into the proposed model to evaluate server power consumption.

Fig. 10 and 11 contrasts the modelled server power profile against measured experiment outputs for air and liquid cooling, respectively. The model is able to successfully capture server power consumption with an average relative error rate of 0.98% and 1.62%, with values ranging between -5.60 – 3.92% (air) and -0.97 – 2.94% (liquid). The weakness of the model arises for idle server power usage, where the error rate deviates up to 7% for liquid cooling. The reason for this deviation is due to parameter selection for Equation 4 to represent $P_{Dynamic}$ (i.e. a zero value for utilization with the applied constant results in a large discrepancy). While the model accurately model power consumption at various utilization level, there is future room for improvement for idle utilization through error correction.

Fig. 12 demonstrates the breakdown of power usage within each server sub-system in liquid cooling. It is observable that increasing the resource utilization results in power usage within each sub-system increasing at different rates. $P_{Dynamic}$ experiences the largest growth from 1% to 33%, driven by increased resource utilization of the server. Furthermore, it indicates that the total cooling constitutes an additional 5.9 – 10% power for the datacenter server to operate, increasing with

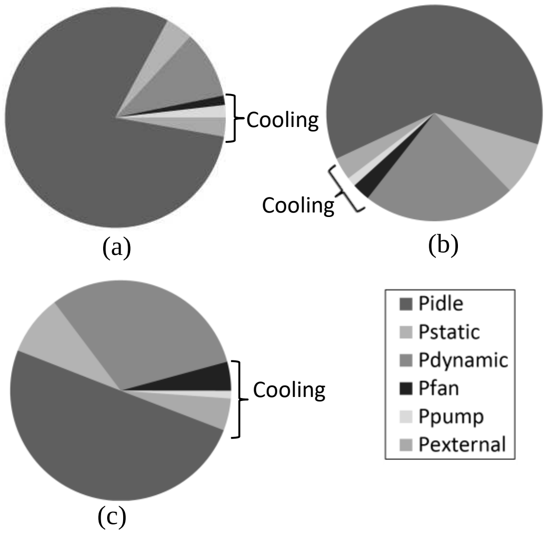

Fig. 13. Power usage of liquid cooled server sub-systems at utilization (a) 20%, (b) 60%, (c) 100%.

higher levels of resource utilization as shown in Fig. 13. In contrast, air cooling results in 3.44% – 13.44% additional power consumption, with similar breakdown of sub-system power usage observed within liquid cooling.

Air cooling consumes on average 7% additional power in comparison to liquid, with the largest deviation of 13% at 100% utilization. When applying $pPUE_{II}$ to servers with air and liquid cooling, we observe a PUE on average of 1.16 and 1.09, respectively. This is due to the capability of liquid to remove heat is greater than air and sustain a lower temperature due to its thermal capacity [19].

VII. APPLICATION OF MODEL

The findings ascertained within this work combined with the proposed model indicates that it is possible to capture the total power consumption of Cloud datacenter servers comprising both hardware and cooling. We describe three potential applications of these findings within the Cloud computing and datacenter research community:

A. Power profiling

The proposed model is capable of producing detailed power profiles of all server sub-systems, and more importantly quantify the cooling power expended under different utilization levels and cooling techniques. By demonstrating that up to 10% and 13% of the total server power usage is driven by cooling for air and liquid (non-detectable by server power meters), we envision the enhancement of numerous energy-efficient scheduling algorithms capable of capturing this new power profile.

B. Air and liquid cooling evaluation

The unified power model allows for studying and experimenting with different server architectures using both liquid and air cooling. We envision that the proposed model can be exploited by the community in order to rapidly evaluate various datacenter server deployments under different cooling configurations. As demonstrated from our analysis and validation, each cooling type will result in different temperature and power profiles, therefore providing more accurate trade-off analysis for application performance and total facility cooling.

C. Enhanced pPUE accuracy

The pPUE metrics detailed in Equations (10) and (11) applied to the analysis results in Section 4 results in values of 1.062 and 1.096 for $pPUE_I$ and $pPUE_{II}$, respectively. While this difference is minimal, it is worth highlighting that $pPUE_{II}$ provides increased accuracy for measuring energy-efficiency, and will be magnified within the context of thousands of servers. Moreover, the external heat rejection is very close to the IT equipment, whereas in larger scale systems the liquid/air must travel significant distance to reject heat. Furthermore, in many scenarios it is not practical or economically feasible to collect data from every sub-system for energy-aware decision making. Therefore the metrics combined withim our model allow for providers to determine and select which metric is most suitable for their requirements.

VIII. CONCLUSIONS

In this paper we have presented a unified model of power usage within a Cloud datacenter server comprising software, hardware and cooling holistically. Through controlled experiments within a laboratory environment, we analyzed the alteration in operational characteristics for multiple sub-systems and propose a model capable of capturing sub-system power usage cross-cutting the entire architecture at various utilization levels. We validate our model in air and liquid cooling experiments, demonstrating high model accuracy under numerous system conditions. We foresee that this model can be rapidly integrated into existing and future server power models for enhanced accuracy. Our contributions are summarized as follows:

Different power profiles within heterogeneous utilization, architectures and cooling systems. Through our analysis we demonstrate that while there exists an intuitive relationship between sub-systems parameters from application performance, microprocessor temperature, fan speed and server power, this relationship is not strictly linear and is dependent on utilization levels and cooling type. This requires rethinking the cascading effects of improvements in application efficiency and its impact onto other sub-systems (and vice versa).

Cooling represents a non-negligible amount of server power usage. Our findings show that the actual power consumed and rejected by cooling systems within the server represents up to 10% and 13% within the total server, varying dependent on cooling type and utilization levels. Such an assumption is frequently omitted in energy-aware scheduling in computer science due to solely measuring server electrical power. Our model allows for researchers to easily integrate this characteristic into their assumptions for server power usage.

Future work includes applying our method and model to a greater number of server architectures and cooling techniques to further validation. Furthermore, we believe that there is potential to apply error correction factoring to further enhance model accuracy for idle server usage to be in line with other utilization levels. Finally, we intend to exploit this model to create new energy-aware scheduling algorithms as well study its integration with facility level power models.

ACKNOWLEDGMENTS

We would like to thanks Django Armstrong for supporting the experiment setup. This work is partly supported by the European Commission under H2020-ICT-20152 contract

687584 - Transparent heterogeneous hardware Architecture deployment for eNergy Gain in Operation (TANGO) project, the University of Anbar, and HCED Iraq. Yaser Al-Anii is the corresponding author.

REFERENCES

[1] J. K, Koomey, Growth in Data Center Electricity Use 2005 to 2010, 2011, Analytics Press. Internet: http://www.analyticspress.com/datacenters.html

[2] Y. Joshi, P. Kumar, Energy Efficient Thermal Management of Data Centers, 2012, Springer.

[3] P. Garraghan, I. S. Moreno, P. Townend, J. Xu, An Analysis of Failure-Related Energy Waste in a Large-Scale Cloud Environment, IEEE Transactions on Emerging Topics in Computing, 2014, vol.2, pp.166-180.

[4] A. Greenberg et al, he Cost of a Cloud: Research Problems in Data Center Networks, Computer Communication Review, 2009, pp.68-73.

[5] C. Gough et al., Energy Efficient Servers: Blueprints for Data Center Optimization, 978-1-4302-6638-9, 2015

[6] I.S.Moreno, P.Garraghan, P.Townend, J. Xu, Analysis, Modeling and Simulation of Workload Patterns in a Large-Scale Utility Cloud, IEEE Transactions on Cloud Computing, 2014, vol. 2, pp. 208-221.

[7] M. Martin, M. Khattar, High-density heat containment. ASHRAE Journal, 2007, pp. 38-43.

[8] S. Albers, Energy-Efficient Algorithms, Communications of the ACM, 53(5), 2010, pp.86-96.

[9] M. Stillwell D. Schanzenbach, F. Vivien, H. Casanova, Resource Allocation Using Virtual Clusters, CCGrid, 2009, pp.260–267.

[10] L.A. Barroso, J. Clidaras, U. Hölzle, The Datacenter as a Computer: An Introduction to the Design of Warehouse-scale Machines, Synthesis Lectures on Computer Architecture, 2013.

[11] Emerson Network, Power Energy Efficient Cooling Solutions for Data Centers, White Paper, 2007.

[12] H. Shoukourian, et al. Monitoring Power Data: A First Step Towards a Unified Energy Efficiency Evaluation Toolset for HPC Data Centers, Environmental Modelling and Software, 2014, vol. 56, pp. 13-26.

[13] S. Pelley, et al. Understanding and Abstracting Total Data Center Power, Proc. of the 2009 Workshop on Energy Efficient Design, WEED, 2009.

[14] C. Möbius, W. Dargie, A. Schill, Power Consumption Estimation Models for Processors, Virtual Machines, and Servers, IEEE Transactions on Parallel and Distributed Systems, 2014, vol. 25, no. 6, pp. 1600-1614.

[15] X, Li, Z. Qian, S. Lu, J. Wu, Energy Efficient Virtual Machine Placement Algorithm with Balanced and Improved Resource Utilization in a Data Center, Mathematical and Computer Modelling, 2013, 58(5), pp. 1222-1235.

[16] Y. Fulpagare, A. Bhargav, Advances in Data Center Thermal Management, Renewable and Sustainable Energy Reviews 43 , 2015, pp. 981-996.

[17] J. Huusho, et al. Energy Efficient Data Centers. First International Workshop, E2DC. 2012.

[18] C.D. Patel, A Vision of Energy Aware Computing from Chips to Data Centers, International Symposium on Micro-Mechanical Engineering. 2003.

[19] J. Dai, et al., Optimum Cooling of Data Centers: Application of Risk Assessment and Mitigation Techniques. 2014, Springer.

[20] S. Sampath, Thermal Analysis of High End Servers Based on Development of Detail Model and Experiments, Master of Science in Mechanical Engineering, The University of Texas at Alrington, 2012.

[21] P.Bohrer, E.N.Elnozahy, T.Keller, M.Kistler, C.Lefurgy, C. McDowell, R. Rajamony, Power Aware Computing , R. Graybill and R. Melhem, Eds. 2002, pp. 261-289.

[22] W. Dargie, A. Schill, Analysis of the Power and Hardware Resource Consumption of Servers Under Different Load Balancing Policies, IEEE International Confencence on Cloud Computing (CLOUD), 2012, pp.772-778.

[23] W.L. Bircher, L.K. John, Measurement Based Power Phase Analysis of a Commercial Workload, in Unique Chips and Systems, E. John and J. Rubio, Eds. CRC Press, 2007, pp. 217-237.

[24] T. Heath, B. Diniz, E.V. Carrera, W.M., Jr., R. Bianchini, Energy Conservation in Heterogeneous Server Clusters, ACM SIGPLAN Symposium on Principles and Practise on Parallel Programming (PPoPP), 2005, pp. 186-195.

[25] D. Economou, S. Rivoire, C. Kozyrakis, P. Ranganathan, Full-System Power Analysis and Modeling for Server Environments, in Proceedings of 2nd Workshop MoBS, 2006, pp. 70-77.

[26] X. Fan, W. D. Weber, L.A. Barroso. Power Provisioning for a Warehouse-Sized Computer. ACM SIGARCH Computer Architecture News [Online], 2007, 35(2), pp. 13-23.

[27] T. W. Harton, C. Walker, M. O' Sullivan, Towards Power Consumption Modeling for Servers at Scale, IEEE/ACM International Conference on Utility and Cloud Computing (UCC), 2015, pp. 315 – 321.

[28] A. H. Beitelmal, D. Fabris, Server and Data Centers Energy Performance Metrics. Energy and Building, Energy and Building, 2014, pp. 562-569.

[29] D. P. Mines, M. Iyengar, P. Parida, R. Simons, M. Schultz, M. Gaynes, R. Schmidt, T. Chainer, Experimental Characterization of an Energy Efficient Chiller-less Data Center Test Facility with Warm Water Cooled Servers, IEEE Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2012, pp. 232-237.

[30] D. P. Mines, M. Iyengar, P. Parida, R. Simons, M. Schultz, M. Gaynes, R. Schmidt, T. Chainer, Impact of Operating Conditions on a chiller-less Data Center Test Facility with Liquid Cooled Servers, Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2012, pp. 562-573.

[31] S. Ham, M. Kim, B. Coi, J. Jeong, Simplified Server Model to Simulate Data Center Cooling Energy Consumption, Energy and Building, 2015, vol 86, pp. 328-339.

[32] M. Iyengar, et al. Server Liquid Cooling with Chiller-less Data Center Design to Enable Significant Energy Savings, Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2012, pp. 212-223.

[33] V. Pandiyan, Development of Detailed Computational Flow Model of High End Server and Validation Using Experimental Methods, Master of Science in Mechanical Engineering, The University of Texas at Arlington, 2012.

[34] P. Parida, et al., Experimental Investigation of Water Cooled Server Microprocessors and Memory Devices in an Energy Efficient Chiller-less Data Center, Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2012, pp. 224 - 231.

[35] Standard Performance Evaluation Corporation, (2012), SPECpower_ssj2008 Results [Online] Available: http://www.spec.org/power_ssj2008/results/

[36] HAProxy - A Reliable, High Performance TCP/HTTP Load Balancer, 2015, Internet: http://www.haproxy.org/