IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received May 30, 2017, accepted June 15, 2017, date of publication July 11, 2017, date of current version August 22, 2017.

Digital Object Identifier 10.1109/ACCESS.2017.2725898

Analysis, Modelling and Characterisation of Zombie Servers in Large-Scale Cloud Datacentres

JOHN PANNEERSELVAM, LU LIU, (Member, IEEE), JAMES HARDY, AND NICK ANTONOPOULOS
Department of Engineering and Technology, University of Derby, Derby DE22 1GB, U.K.

Corresponding author: John Panneerselvam (j.panneerselvam@derby.ac.uk)

This work was supported in part by the National Natural Science Foundation of China under Grant 61502209 and Grant 61502207, and in part by the U.K.–China Knowledge Economy Education Partnership.

ABSTRACT Cloud datacentres are acknowledged as being massive energy consumers, which may have significant environment impacts. Service providers have an ethical responsibility to reduce the environmental impact of server resources and a simultaneous and complementary commercial desire to reduce energy costs. Zombie servers in the datacenters are one of the primary sources of undesirable energy expenditures by incurring idle resources during task execution. This paper investigates the cause, impact and energy-related implications of zombie servers. Important outcomes of this paper are the characterization of the diversity among the workload behaviors in resource consumption and the quantification of the presence of idle CPU and memory resources during task execution causing server zombieness. The undesirable power consumption of zombie servers is determined based on the profiles of currently available servers and their corresponding environmental implications are illustrated in this paper. Empirical analysis shows that cloud workloads are highly heterogeneous in resource consumption pattern and CPU resources may display 75.6% of idleness relative to their allocated level, while memory is 25.5% idle. The report concludes that significant reductions in power consumption and CO2 emission can be achieved by provisioning a realistic level of resources to servers, which are scaled to suit the anticipated workloads.

INDEX TERMS Energy-aware systems, data exploration, power management, ubiquitous computing.

I. INTRODUCTION

The emergence of Cloud Computing over the recent years has achieved tremendous exposure in both academia and industry. Cloud datacentre resources are witnessed to consume tremendous amounts of energy and are generating large amounts of carbon footprints. Cloud providers are contractually committed to provide services without violating the terms of the initially negotiated SLA (Service Level Agreement) [1]. The SLA is paramount in quantifying the attributes used to define, measure and maintain the Quality of Service (QoS) at a desired level and for setting and managing the Quality of Expectations (QoE) of the end users. In the Cloud Computing service model, client demand requirements are satisfied by provisioning resources, this is typically at a level which exceeds [2] the actual amounts of resource necessary to process their prospective job. User requests arriving at the datacentres in the form of job submissions are usually scheduled and allocated onto Virtual Machines (VMs) deployed on physical servers. Providers allocate the resource levels for task execution based on the intensity of the user demands. Since the allocated resource levels within VMs are not often completely utilised, this directly causes the physical servers to operate below their actual capacities. This results in an increased proportion of server resources being idle and as a result of which, approximately 46% of machines [3] are operating below maximum capability for a significant amount of time with their resources being reserved for the over-estimated resource requirements. These resources remain idle and unutilised, consuming unnecessary energy and without contributing to the execution of the workloads. Servers with increased proportions of idle resources are termed as ‘comatose’ or ‘zombie’ [4] servers, which are usually powered on and consuming electricity but delivering no useful information services. Unfortunately zombie servers can remain unnoticed in datacentres since they do not have service affecting failures, their utilisation metrics do not generate exceedance alarms and they are therefore unlikely to appear in any ''top 10'' reporting lists. It has been shown that an idle server may consume approximately two-thirds of the energy used by the same server operating at full load [5]–[7]. It is worthy of note that the power overhead incurred by resource idleness vary significantly for different hypervisors on alternative physical servers.

While it is possible to simultaneously reduce energy impacts by reducing the active resources at the datacentres, however this may result in lower levels of resource availability which could in turn causes temporal loss of service, jeopardising the agreed SLA, and damaging the commercial credibility of the provider. Service unavailability is an unacceptable event for the Cloud users as they have purchased a contract based on the structure and service appeal of Cloud Computing, namely an illusion of infinite resource availability. Exposing the cause and impacts of zombie resources in Cloud datacentres is an integral requirement of the Cloud service model not only to achieve energy efficiency but also to maintain optimised performance quality.

The imprecise knowledge and understanding of the characteristics of both the datacentres and the Cloud workloads [8] are primary causes for the less than perfect efficiency of existing methodologies of energy efficiency [5], [9]–[12]. Improvements can be made if the workload and datacentre behavioural characteristics and corresponding energy implication is quantified for the entire system. To this end, extensive analysis has been conducted in the recent past with the motivation of exploring the undesirable energy consumption within the Cloud server resources. Previous analysis [13]–[19] has emphasised phenomenon such as long tails, task failures, rollback and performance interference as well as latency sensitiveness of the workloads, different run-time tasks etc., as being primary causes for undesirable energy consumption. The main focus of the majority of the cited works is on the energy implications resulting from failure related events in the Cloud datacentres. Despite the existing works, an in-depth analytics of the workload characteristics is important for two reasons. Firstly, analysing the inner distribution of resource consumption trend and task length helps exposing the job behaviour heterogeneity in resource consumption pattern. Secondly, studying the requested-to-utilised resource ratio helps to quantify the proportions of idle resource times for identifying the zombie server resources.

The major contributions of this paper include:

1. An empirical analysis of the characteristics of the Cloud jobs in consuming the provisioned server resources. This analysis uncovers the diversity among the submitted jobs in consuming the previsioned level of resources and encompassed task duration and exposes the cause for energy wastages among the job execution.
2. The first comprehensive analysis of the presence of proportional idle resource time for CPU and memory during every single task execution session in a large-scale datacentre operation to spot server zombieness. This analysis can be applied when allocating optimum level of resources for task executions to reduce over-allocation of resources. Also, this analysis can find applications in the prediction models aiming to predict the resource utilisation levels of the user requests. Minimising the allocated-to-utilised resource ratio will reduce the resource idleness among the server resources, thereby reducing undesirable energy wastage.
3. Mensuration of the power consumption and carbon footprints incurred by the proportional idle resource times during tasks execution, with our comprehensive analysis conducted based on in-trend server power profiles.

The rest of the paper is organised as follows: Section II presents the related works and Section III gives an overview of the Cloud workloads. Section IV presents the methodology of our analysis and Section V is covered with our analysis on resource consumption trend of the jobs. Section VI includes the analysis of the presence of the idle resource time in terms of CPU and memory resources during the task execution sessions. The energy wastage incurred by the power consumption of the zombie resources are illustrated in Section VII and their corresponding environmental implications are detailed in Section VIII. Section IX presents the application of our work and Section X concludes this paper along with defining the future direction of our research work.

II. RELATED WORKS

Usually, undesirable energy consumption is observed at various levels within the processes and components in a large-scale datacentre environment. The impact of failure related energy consumption are investigated in the works of [14] and [19] where they quantified the various termination events and consequent energy implications based on the Google trace logs. Furthermore, they investigated the task terminations in relation to the priority levels of the tasks. This analysis insists that large numbers of task kill and evict events are prevalent in Cloud datacentres, thereby causing undesirable energy expenditures. The cause for such terminations has also been analysed by the same work and the failure and repair times incurred energy expenditures have been exhibited. Apart from this work, the undesirable energy consumption incurred by task failures and terminations have been extensively studied [16], [20]–[24] from various perspectives in similar environments such as grids, large-scale MapReduce applications in Cloud environments etc.

A latency-aware analysis has been conducted in our earlier work [17], where the various task terminations events are quantified in accordance with the latency sensitivity levels of the submitted tasks. This work leads us to infer that tasks characterised by higher level of latency sensitivity suffer increased terminations and thus account for excess energy consumption. Apart from the network and dispatch latency [25], the in-house computing latency of the tasks has significant impact on the overall energy consumption. Furthermore, the impacts of latency on the Cloud environments are presented in the works of [26]–[28] revealing that

TABLE 1. Summary of related works.

Authors	Cause	Power Model	Trend
Garraghan [14][19]	Task terminations	Yes	2012-13
Panneerselvam [17]	Latency sensitivity	No	NA
Garraghan [13]	Long tails	No	NA
Yang [15], Moreno [29]	Performance interference	Yes	iVIC (2013)
Chen [18]	Diverse run-time tasks	Yes	2012
Baliga [30]	Transmission and switching	Yes	2010
Lin [31]	Communication components	Yes	2015
Shaoling [32]	Equipment dormancy	Yes	2015
Schroeder [20], Nguyen [21], Quiane-Ruiz [22], Liang [23], Li [24]	Execution failures	No	NA
Jin [6]	Server virtualisation	Yes	2012
Our analysis	**Zombie resources**	**Yes**	**2016**

latency sensitivity levels of the tasks have significant impact on QoS, energy consumption, termination events and end user tolerances.

The phenomenon of longs tail have been investigated in the works of [13] and it has been demonstrated that the presence of smaller proportions of long tails can significantly impact the completion times of the tasks. Long tails are the stragglers those significantly delay the completion of the entire tasks, thereby incurring excess energy consumption. The performance interference effects on energy consumption have been studied [15], [29] and the impact of such non-negligible performance interference overheads was negotiated with a workload placement model for reducing energy wastage. In general, aggressively consolidating workloads having similar resource intensiveness (CPU or memory) will lead to such performance interference effects. Energy consumption of different run-time tasks has been investigated by exploring the correlation of energy consumption with computational tasks [18]. The energy-aware analysis of the computational correlations among the tasks helps to achieve energy-efficient task placement and optimal resource management. Energy consumption of the sending and receiving network switches [30], and communication components [31] in Cloud datacentres has been studied. These communication components can consume considerable amounts of energy whilst connecting the users to the Clouds. The energy consumption of dormant server equipment [32] in Cloud datacentres has been analysed to develop a model for optimising energy consumption of the network components. This work states that individual server resources can achieve the lowest possible energy consumption state without affecting other components working under normal conditions. The effect of virtualisation [6] on overall datacentre energy consumption has been investigated by analysing server energy usage under various hypervisor configurations. All such works demonstrated that there is an increased scope for reducing excess energy consumption at the Cloud datacentres.

From state-of-the-art research works on energy analysis, it is clear that there is no special emphasis given to energy consumption incurred by the zombie resources resulting from unutilised idle resources in the Cloud datacentres. Most studies are completely focused on identifying the cause, effect and implications of task failures on the overall energy consumption. Given that server resources are being under-utilised in a Cloud datacentre, idle resource times contribute a significant proportion of the overall energy consumption. The overall datacentre energy consumption is an accumulation of the energy consumed by various components and their corresponding events. Ignoring the energy consumption of zombie servers leaves a large proportion of energy waste unnoticed, and so a complete datacentre energy consumption profile cannot be obtained.

This necessitates the need for an extensive analysis of the presence, cause, and implications of the zombie servers in large-scale datacentres. With this in mind, this paper dwells into the presence, cause and implications of the idle CPU and memory resource times whilst executing tasks in large-scale datacentres for exposing the energy related implications of zombie servers. The distinctive contribution of this paper from the related works is illustrated in Table 1.

III. BACKGROUND

A. CLOUD WORKLOADS

A typical Cloud workload [17], [33], [34] arrives at the Cloud datacentre in the form of jobs submitted by the users. Every job includes certain self-defining attributes such as the submission time, user identity and resource requirements in terms of CPU, memory and disk space. A single job may contain one or more tasks, which are scheduled for processing at the Cloud servers. A single task may have one or more process requirements. Tasks may have varied service requirements and characteristics such as throughput, latency, jitter, etc., even though they belong to the same job. The provider generally records the resource utilisation levels of every scheduled task and maintains the user profiles.

Cloud workloads behave distinctively with different server architectures and this behaviour of workloads in the Cloud processing environment is more strongly correlated with the

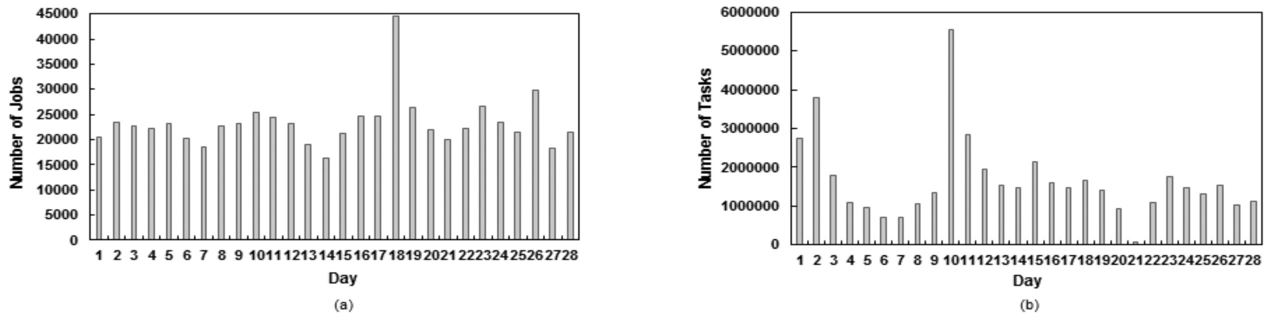

FIGURE 1. **User request submission (a) Jobs. (b) Tasks.**

CPU cores than with RAM capacity of the machines at the server level. While the capacity levels of CPU and memory in a physical server usually remain static, the utilisation levels are more dynamic and vary abruptly with different workloads under different server architectures. The dynamic parameters of the server architectures are usually calculated as the measure of the number of cycles per instruction for CPU, and memory access per instruction for memory utilisations respectively. Thus the task resource utilisation profile (r_t) can be expressed as a multi-dimensional composite as shown in equation 1.

$$r_t = \{f(c), f(m), f(d), f(t)\} \tag{1}$$

where c is the CPU usage in core counts, m is the memory usage in bytes and d is the disk space utilisation and t is the task execution duration.

TABLE 2. **Trace log statistics.**

Observation	Value
Number of Days	28
Total Number of Job Submissions	650892
Total Number of Task Submissions	46093201
Number of Operating Servers	12500
Average Number of Users per Day	190

B. DATA SAMPLE

This research work explores the Cloud trace logs [35] released by Google, featuring more than 650000 job submissions across 46093201 tasks over 28 days of datacentre execution. The Google trace logs are investigated in order to explore and observe patterns of workload heterogeneity, and resource provision and resource consumption profiles of tasks at the Cloud datacentre. The analysis is intended to extract information pertinent to energy efficient operation of the datacentres. The event analysis of the trace logs based on our previous research [17] has been presented in Table 2 and Fig. 1. In order to include the impact of unsuccessful events on the overall energy consumption, the submission events presented in this paper include resubmissions resulting from the workload terminations.

IV. METHODOLOGY

This section outlines the methodology of our analysis in quantifying the user requests, their resource utilisation levels and the power consumption incurred by the proportional idleness of the allocated resources during task execution sessions causing server zombieness. A power model is described to measure the proportional idle resource time during task execution and their corresponding energy expenditures.

A. WORKLOAD SAMPLING

The trace log data has been sampled on a per day basis with a single day spanning across 24 hours starting from 12.00 am for a given day in order to accurately model the time-of-the-day behaviour of the workloads.

The Cloud trace logs provide the explicit information [36] of the task events including start and end time, job ID, user, task priority levels, and resource requests in terms of CPU, Memory and Disk space, and the corresponding amounts of resources consumed during every execution session. The task usage information of the trace logs required for our analysis includes mean CPU usage rate, maximum CPU usage, assigned memory and maximum memory usage for every single task execution session. With the assigned and used amounts of resources being explicitly provided, the unutilised resources is computed using the maximum to mean utilisation ratio for every task execution session.

The elapsed time period of a task execution is calculated from the start and end time of an execution session. The elapsed time period are not necessarily the completion time of the tasks, since a task execution session may be successful or might have faced a termination resulting in resubmission of the terminated task. Thus the elapsed time calculated is the time period when the status of the corresponding tasks is updated with either completion or termination. However, providers allocate resources for all the task execution sessions, and resubmissions of the tasks will cause the providers to allocate resources again for the terminated tasks.

B. HETEROGENEITY ANALYSIS

Both the Cloud workloads and the datacentre resources [37] exhibit extreme dynamic characteristics. The workloads being submitted show heterogeneities and dynamism in their actual resource consumption pattern at the datacentres. For instance, workloads with similar resource requests may not be similar in their actual resource utilisation patterns at the back end Cloud servers. In addition to quantifying the resource provision and consumption trend for the entire days, analysing the characteristics of individual jobs and their impacts on the overall resource consumption trend is important to understand the job heterogeneity within a single datacentre. In spite of exposing the job heterogeneity, jobs encompassing different number of tasks (50, 100, 200, 500 and 1050, named Job 0 through to Job 5) have been analysed as representatives of different job behaviours and encompassing number of tasks. The distribution of the data trend of the provisioned and utilised resources have been evaluated against notable theoretical distribution such as Normal, Lognormal, Exponential, Gumbel, Gamma, Weibull etc., and further the best fit distribution is presented for the actual data trend in terms of their CDFs within every studied jobs based on Anderson-Darling test. Firstly, the amounts of CPU and memory resources provisioned has been evaluated against the trend of actual resource consumption for every individual tasks within their respective jobs. Secondly, the task duration within a single job has been analysed to observe the task heterogeneity within jobs.

In spite of the extensive analysis conducted on server resource heterogeneity [14], [38] in the previous works, quantifying the datacentre heterogeneity to any further degree is out of scope of this paper.

C. PROFILING RESOURCE UTILISATION

Resource utilisation levels has two measurement viewpoints. First is the job level resource utilisation which is usually the measure of CPU, memory resources consumed by the workloads whilst executed in the VMs. Secondly, machine level utilisation is the measure of ratio of the actual usage level of a machine to its maximum usage capacity. An accumulation of the job level utilisation of all the VMs directly reflects the machine level utilisation of the corresponding physical server.

The resource utilisation for task execution sessions are characterised in terms of the CPU and memory resource consumption at the task level. With the duration of the task execution session being calculated for every task, the resource utilisation profiles of the tasks are enumerated to exhibit the proportional idle resource times over the actual level of resource allocation for every task execution session. Before quantifying the resource idleness, it is necessary to compute the total amount of resource time actually allocated for a task execution session, from which the idle resource times can be measured. The unutilised memory resources in a given task execution session is a direct measurement by subtracting the maximum memory usage from the assigned memory. Due to the data ambiguity of the CPU resource utilisation levels provided in the trace logs, it is necessary to make an assumption whilst profiling the CPU resource utilisations. We assume that the maximum CPU usage level in an execution session as the maximum allocated CPU resources for that task execution session. This maximum resource usage level is the highest peak of resource usage level in a given session and the mean resource usage is the mean value of the remaining usage level.

As the maximum to mean resource usage ratio $r_{i(ratio)}$ increases, the presence of idle resource time in an execution session shows a corresponding increase. Conversely, lower values of $r_{i(ratio)}$ correspond to higher values of Power Usage Effectiveness (PUE) of the servers. PUE [39], [40] defines how effectively a server is using its electricity, which is always desirable to be at an optimum level to achieve energy-efficient computing. A PUE closer to 1.0 is very effective, implying that almost all the energy is transformed into computing power. A PUE of 2.0 means that every computationally useful watt of input power will require an additional watt for cooling, lighting, power distribution, etc. Now for profiling the CPU resource utilisation levels, the total amount of resources time consumed and the presence of idle resource time in a task execution session is calculated using equations 2 to 4.

$$t_r = \int r_a \cdot t \tag{2}$$

$$r_i = r_a - \int r_m \tag{3}$$

$$r_{i(ratio)} = \frac{(r_a - r_m)}{r_a} * 100 \tag{4}$$

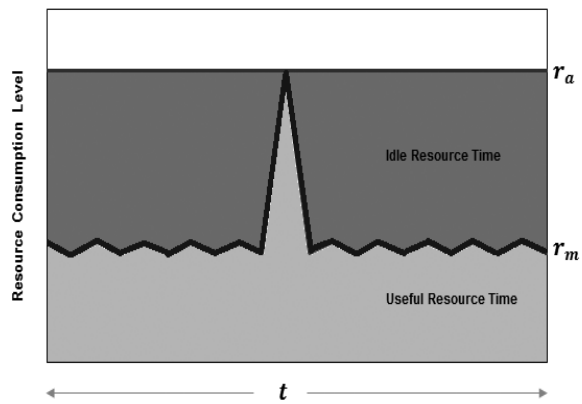

FIGURE 2. Idle resource time proportion.

where, t_r is the total amount of resource time allocated for a task execution session, r_i is the amounts of unutilised idle resources in an execution session whilst executing tasks, r_a is the maximum level of allocated resources for an execution session, r_m is the mean resource usage rate, $r_{i(ratio)}$ is the proportional idle resource time and t is the duration of the task execution session. The proportional presence of the idle resource time in a task execution session is illustrated in Fig. 2.

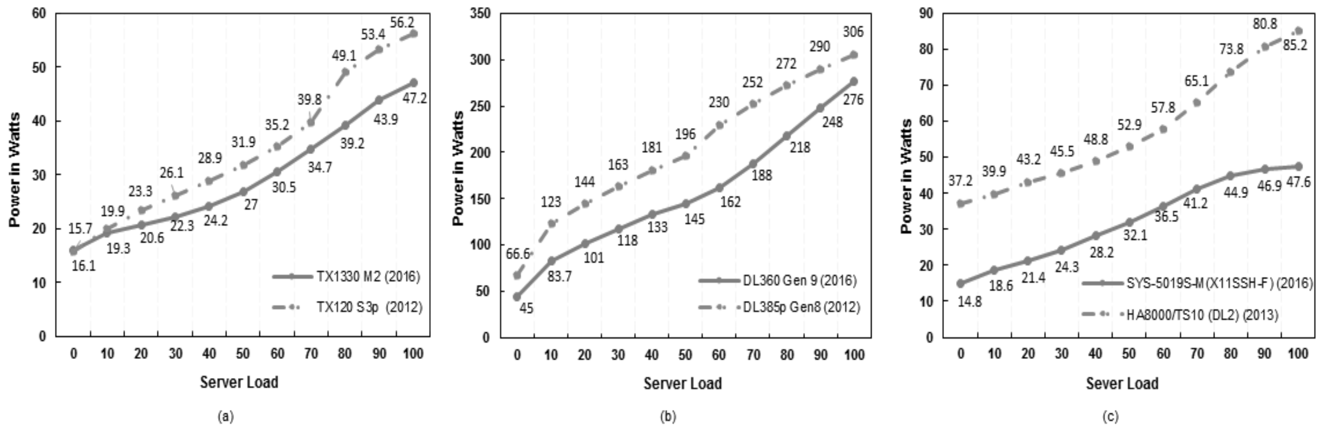

FIGURE 3. Power profile comparison of servers: (a) server A, (b) server B, and (c) server C.

TABLE 3. Server capacity comparison.

Plat-form Type	2016					2012-2013				
	Server	CPU	CPU Capacity (ssjops)	Total Memory (GB)	Max Temp (100% load)	Server	CPU	CPU Capacity (ssjops)	Total Memory (GB)	Max Temp (100% load)
A	TX1330 M2	Intel Xeon E3-1240L v5	484,122	16	22.6	TX120 S3p	AMD Opteron 6380	1,660,274	8	20.6
B	DL360 Gen 9	Intel Xeon E3-1265V2	420,255	64	21.4	DL385p Gen8	Intel E3-1260L v5	528,348	64	20.4
C	SYS-5019S-M(X11SSH-F)	Intel Xeon E5-2699 v3	3,159,419	16	21.2	HA8000/TS10 (DL2)	Intel Xeon E3-1280 v2	495,083	8	23.1

D. POWER MODEL

The energy consumption of the zombie servers incurred by the idle resource times is computed based on the unutilised amounts of resources in every task execution session. Such unutilised resource capacities might result from the early completion of the tasks, in which case, we refer the unutilised resources as over-allocated resources. Another factor causing the end of the execution session is the termination events, where all the allocated resources are referred to as undesirable resource wastage. All the unutilised resources consume energy without contributing towards the actual task execution and are the primary source of zombie servers.

Potentially unrelated to the demand for energy efficient computing, Cloud service providers are also in the process of enhancing the energy profiles of their operating servers. The energy profiles of the servers developed by Oracle Corporation and IBM are used in our analysis for the purpose of measuring the energy consumption of the zombie resources in order to match the profiles of datacentres in practice today. It is notable that the energy consumption of the selected servers has reduced since 2012 despite their increasing capacities of CPU and memory resources. The improving energy efficiency of the server power profiles over the recent years is depicted in Fig. 3, comparing the energy profiles of the current server profiles (2016) with those used in 2012. The server profiles for this comparative analysis have been selected based on a close match with their CPU type and capacity in order to compare their power consumption trend, as shown in Table 3. The utilisation levels of the CPU is measured in terms of Server Side Java operations (ssjops) and the power consumption is depicted in Watts. Though the trace logs are obtained from Google servers, we perform our analysis based on three different in-practice server profiles to match the current energy consumption trend of the server resources according to the SpecPower 2008 benchmark [41]. The active idle power consumption of the servers is required to maintain the server turned on and the power consumption of the server increases with increasing percentage of the server load. A task execution will consume energy equivalent to the product of the respective proportional power of the current server load and the duration of the task execution session. When there is an increase in the server load, the

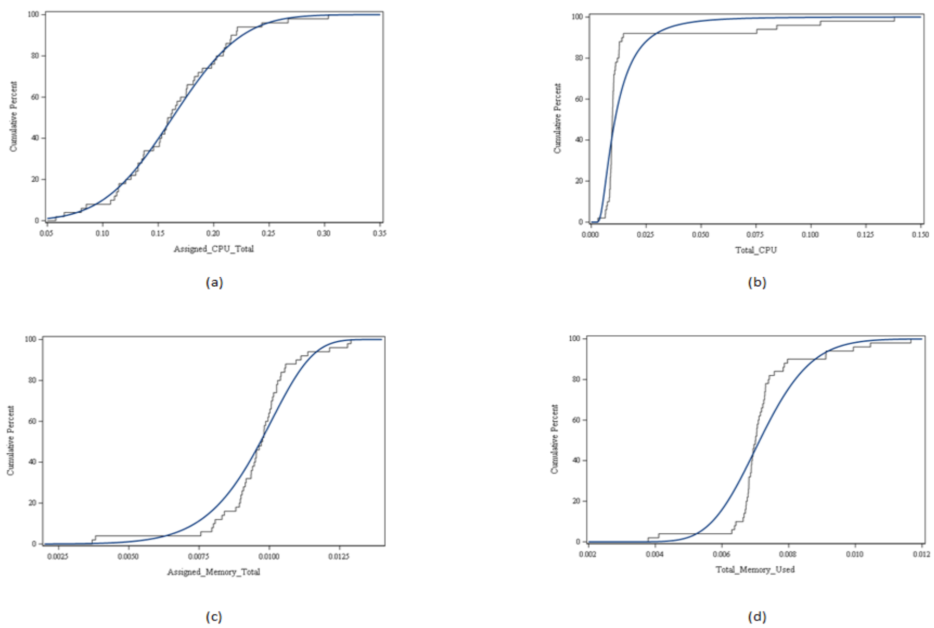

FIGURE 4. CDF of Job 0: (a) Assigned CPU, (b) CPU consumed, (c) Assigned memory, and (d) Memory consumed.

TABLE 4. Resource consumption heterogeneity.

Job Name	CPU					Memory				
	Assigned		Consumed		Proportion Wasted	Assigned		Consumed		Proportion Wasted
	Distribution	Parameters	Distribution	Parameters		Distribution	Parameters	Distribution	Parameters	
Job 0	Normal	μ=0.16276 σ=0.04914	3P Lognormal	ϴ= 0.0025 ζ=-4.788 σ= 0.8408	90.5%	3P Weibull	ϴ= 0.006 ζ=0.0166 c= 12.02	3P Lognormal	ϴ= -0.003 ζ=-4.608 σ= 0.1204	23.5%
Job 1	Normal	μ =0.0822 σ = 0.0439	3P Lognormal	ϴ= 0.0005 ζ=-5.56 σ=1.7625	87.8%	Gumbel	μ=0.004 σ=0.001	3P Lognormal	ϴ=0.0002 ζ=-5.791 σ= 0.525	24.5%
Job 2	3P Lognormal	ϴ=-0.033 ζ=-2.869 σ= 0.177	Normal	μ =0.0026 σ = 0.0015	88.8%	3P Lognormal	ϴ= -0.018 ζ=-3.893 σ=0.0028	Lognormal	ζ=-5.791 σ= 0.525	42.5%
Job 3	Normal	μ = 0.0459 σ = 0.0266	3P Lognormal	ϴ=376E-7 ζ=-6.471 σ=1.1861	91.6%	N/A	N/A	3P Weibull	ϴ=-0.014 ζ=0.0618 c=31.124	31.2%
Job 4	Normal	μ = 0.0588 σ = 0.0261	Normal	μ = 0.0132 σ = 0.0102	80.8%	3P Lognormal	ϴ= 0.0016 ζ=-8.162 σ=0.1728	Weibull	ζ=0.0016 c=5.3478	21.8%
Job 5	Lognormal	ζ=-0.208 σ=0.2412	Lognormal	ζ=-0.411 σ= 0.2398	16.15%	Normal	μ = 0.078 σ = 0.0134	Normal	μ = 0.0858 σ = 0.016	9.75%

power consumption of the servers will increase with a decrease in the amounts of idle resources. Thus the total power consumed in a task execution session is computed by equation 5.

$$P(r) = P(r_s) * t \tag{5}$$

where, $P(r)$ is the total power consumption during the task execution session and $P(r_s)$ is the server power proportional to current load and t is the task execution duration. Now, the presence of the proportional idle resource time (zombieness) and its corresponding power consumption $P(r_i)$, can be computed using equation 6 and 7.

$$t_{ri} = t * r_{i(ratio)} \tag{6}$$

$$P(r_i) = P(r_{ai}) * t_{ri} \tag{7}$$

where, t_{ri} is the proportional idle resource time resulting from the unutilised server capacities and $P(r_{ai})$ is the server power on active idle. The power consumption for every task execution session is computed and we obtain the total amount of idle resource proportion and the power consumed by the presence of idle resource times on a per day basis over the observed period of 28 days. Since the tasks are processed in virtualised servers, the power consumption is computed based on the allocated resources in the virtual cores running across the physical server resources.

V. HETEROGENEITY ANALYSIS

This section presents the analysis of the diversity among the user submitted requests across the sampled dataset in terms of the provisioned to consumed amount of resources for both CPU and memory resources. Table 4 presents the

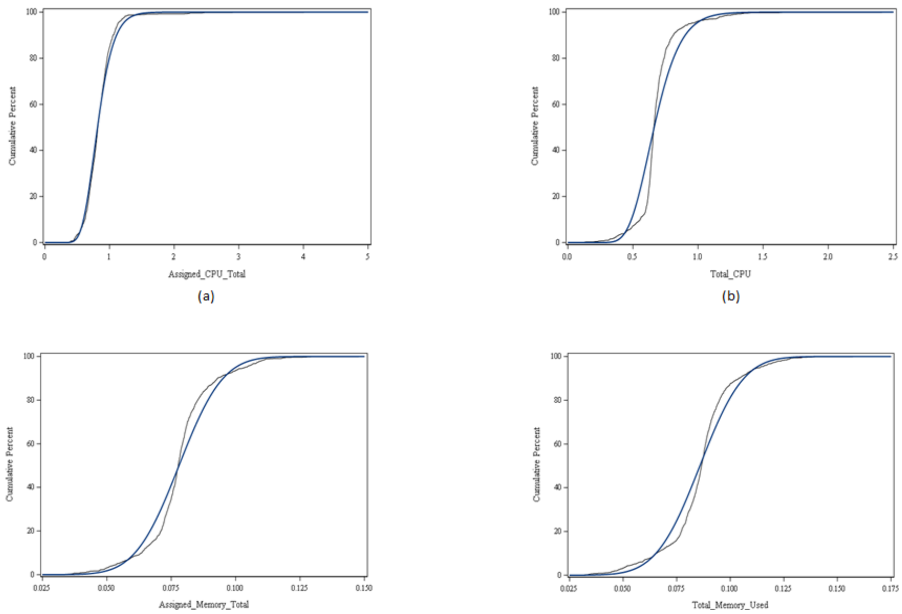

FIGURE 5. CDF of Job 5: (a) Assigned CPU, (b) CPU consumed, (c) Assigned memory, and (d) Memory consumed.

distribution statistics for all the studied jobs in terms of the proportions of resources provisioned and consumed within every studied job execution. It can be postulated that the distribution trend of the assigned and utilised resources can determine the resources wasted accordingly.

From Table 4, it can be observed that the studied jobs are heterogeneous in effectively utilising the assigned resources without wastage of resources, predominantly following different distributions. For space constraints, only Job 0 and Job 5 have been chosen to display their inner distribution in Fig. 4 and Fig. 5 respectively, as they exhibit different extremism of resource consumption trend. From Fig. 4 and Table 4, the assigned CPU within Job 0 predominantly follows normal distribution and the consumed CPU predominantly follows 3P lognormal distribution and the curve is right skewed. Whilst the CPU cores have been provisioned and distributed equivalently across the encompassed tasks, a majority (around 90%) of the encompassed tasks consumed only a marginal proportion of the assigned resources and only a minority (around 10%) of the tasks has consumed a reasonable margin of the provisioned resources. An immediate implication is that the former 90% tasks are vulnerable to leave most of the provisioned resources utilised, causing 90.5% of CPU idleness since the provisioned and consumed curves are extremely heterogeneous. In addition, the memory assigned to the tasks follows 3P Weibull and memory consumed follows 3P Lognormal distribution respectively, with both the curves are slightly left skewed. Idleness in the memory resources are witnessed at just around 23.5% since both the curves follow a similar distribution trend.

From Fig. 5, both the CPU assigned and consumed for Job 5 curves predominantly follow Lognormal distributions and are left-skewed. Whilst more than 90% of the encompassed tasks within Job 5 are provisioned with less than around 1.3 core counts, 90% of the tasks have consumed just less than around 1 core counts each. Thus the CPU assigned and consumed trend are nearly homogenous and the curves share a close enough distribution, whereby reducing the CPU idleness to 16.15%. Furthermore, a similar behavioural trend is evident in the trend of memory resources, since both the memory assigned and consumed curves follow normal distribution. It can be arguably climbed that both the curves of memory trend are nearly identical, which has reflected in a significant reduction in the amounts of resources wasted accounting only at 9.75% of the provisioned resources. Thus it can be postulated that jobs with heterogeneous distributions between the resource provision and consumption trend are vulnerable to leave most of the provisioned resources unutilised, causing a significant proportions of resource wastages. For achieving an energy efficient job execution, a close-enough distribution should be achieved between the trend of resources provisioned and the resources consumed. Though, the distribution trend between the provisioned and consumed trend of jobs can only be achieved post the job execution.

Table 5 presents the observations of the task length for tasks encompassed within the studied jobs and Fig. 6 displays the CDF with the best fit distributions for tasks within Job 0, Job 5 and Job 3. Both Job 0 and Job 5 predominantly follows a 3P Weibull distribution and the curves are significantly left-skewed, insisting the fact that both the jobs encompass tasks with shorter, medium and long running tasks causing an increased heterogeneity among the task length within a single job. Within Job 0, 10% of the tasks characterise a task

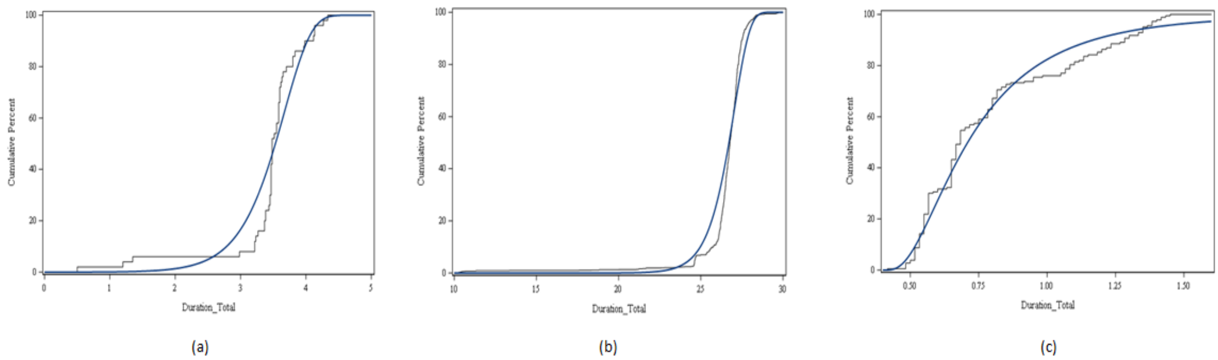

FIGURE 6. CDF of task duration: (a) Job 0, (b) Job 5, and (c) Job 3.

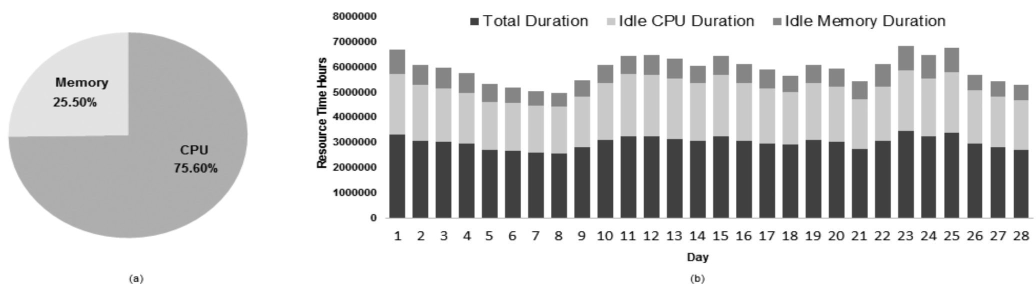

FIGURE 7. Idle resource time analysis (a) idle resource proportion and (b) day-wise resource time in hours.

TABLE 5. Task length heterogeneity.

Job Name	Duration	
	Distribution	**Parameters**
Job 0	3P Weibull	ϴ=-21.25 ζ=24.92 c= 61.54
Job 1	Normal	μ = 2.5038 σ = 0.401
Job 2	3P Weibull	ϴ= -0.594 ζ=1.644 c=8.4604
Job 3	3P Lognormal	ϴ= 0.3881 ζ=-1.133 σ= 0.6901
Job 4	Weibull	ζ=3.0382 c= 2.2869
Job 5	3P Weibull	ϴ=-57.39 ζ=84.448 c=89.825

length of around 3 minutes and another group of 10% of tasks characterise a task length of more than 4 minutes, and the remaining 80% of the tasks runs between 3 and 4 minutes respectively. Within Job 5, 10% of the tasks characterise a task length of around 26 minutes and another group of 10% of tasks characterise a task length of more than 28 minutes, and the remaining 80% of the tasks runs between 26 and 28 minutes respectively. Here the task length is fairly homogeneous which reduces the presence of an increased proportions of long tails which would delay the job completion time. Conversely Job 3 predominantly follows a 3P Lognormal distribution and the curve is right-skewed. Here the task length is heterogeneous across the encompassed tasks within Job 3, with around 70% of tasks runs for less than a minute and the remaining 30% runs for more than a minute up to a maximum of 1.7 minutes. From these observations, it is clear that the jobs are increasingly heterogeneous further the tasks encompassed within a single job may exhibit an increased diversity in terms of their resource consumption. Task may or may not exhibit homogeneity in terms of their running task length within a single job. Both the jobs and every tasks within jobs should be uniquely treated whilst attempting to optimise their resource usage profiles for achieving energy efficiency. CPU resource provision are increasingly vulnerable to leave most of the provisioned resources unutilised and the provisioned memory resources are fairly utilised.

VI. IDLE RESOURCE TIME ANALYSIS

This section analyses the presence of idle resource time among the allocated CPU and memory resources during the task execution causing server zombieness, which are in essence, can accommodate more workloads. The idle CPU and memory resource times compared with the actual allocated resource time is illustrated in Fig. 7. Fig. 8 depicts the presence of idle CPU and Memory resource percentages over the observed period of 28 days.

The usage measurement sessions in the trace logs include the usage measurements for periods when no process belonging to the task execution was running in the task's container. For this reason, we normalise the values presented in Fig. 8, by obtaining a mean value of the proportional resource time for every measurement period and normalise it against the overall duration in a given day in order to enhance the

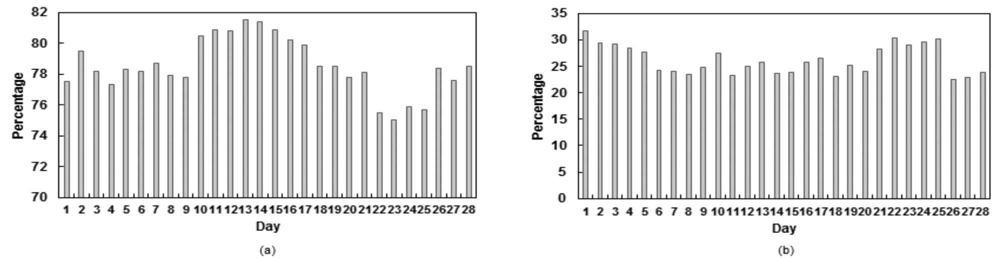

FIGURE 8. Day wise idle resource time percentage (a) CPU and (b) memory.

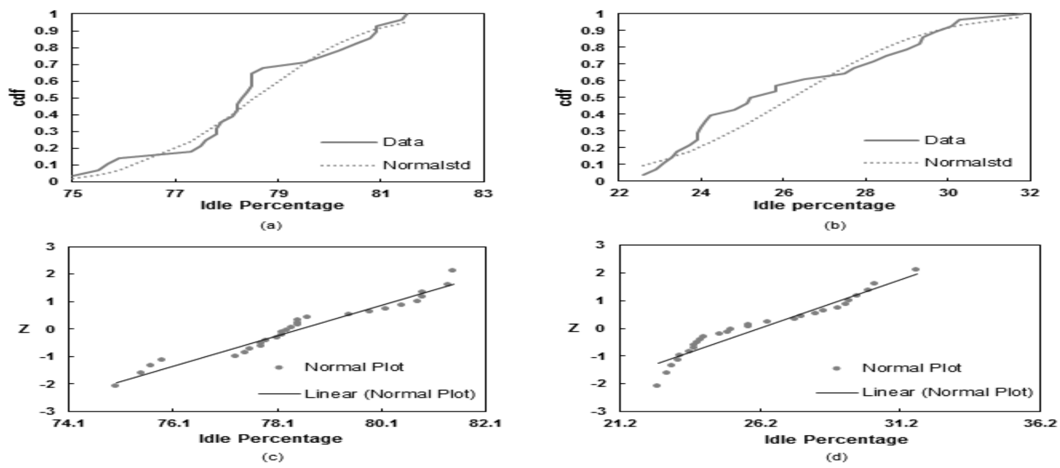

FIGURE 9. Idle resource time distribution analysis: (a) idle CPU percent, (b) idle memory percent, (c) idle CPU AD test, and (d) idle memory AD test.

computation accuracy. It can be observed that an average of 75.6% of the resource execution time being wasted caused by the idle CPU cores. This suggests that the CPU resources are commonly being over-allocated, leaving many of the CPU cores active without any actual contribution towards task execution. The immediate implication of the presence of excessive idle CPU time is that the server capability is always under-utilised because tasks with a specific processor demand are being allocated to servers with capabilities that far exceed the minimum requirement. For maximum greening in datacentres, server CPU resources should be allocated with a task load corresponding to their peak processing capability. This strategy will reduce the server zombiness and help the providers to achieve a PUE factor closer to 1.

We can observe an average of 25.5% of idle memory resource time amongst the allocated memory resources. Though the wasted memory resources are considerably less than those of CPU, memory resources are still over-assigned. The actual CDF fitted with a normal standard distribution of idle resource time percentage for both CPU and memory resource is depicted in Fig. 9, along with the probability plot of the Anderson Darling Goodness of Fit test. It is evident that the CDF distribution of both the CPU and the memory idle resource times are showing measurable fluctuation from that of the normal standard distribution. This supports our observations of day-wise loose correlation and increased fluctuations among both the CPU and memory idle resource times. We repeated the Anderson Darling test for data normality on both the CPU and Memory idle resource time distribution on a daily basis. The test shows that both CPU and Memory idle presence follows a near to normal distribution, with the CPU distribution negatively skewed at -0.108 and the memory distribution positively skewed at 0.434.

A. RESOURCE TIME FLUCTUATION ANALYSIS

From Fig. 9, there is no correlation evident between the idle memory and CPU resource times, suggesting that the CPU and memory consumption levels of the tasks are independent to each other, though this may be dependent on the nature of the tasks. Since the CPU is often the largest power consumer in the purely digital system, the increased presence of idle CPU resources will incur a more significant energy consumption. Further, the allocation and consumption of CPU cores exhibits increased fluctuations. This increased fluctuation in the usage of CPU cores over the memory resources is illustrated using the standard deviation function of the idle resource time, as shown in Fig. 10. This fluctuation is measured as the function of deviation evident among the presence of idle resource time among the co-located execution session within every day.

From Fig. 10, the average standard deviation of the idle CPU and memory resource time are 22.9 and 35.2,

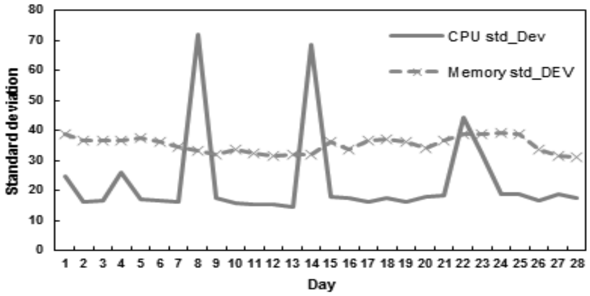

FIGURE 10. Idle resource time fluctuation.

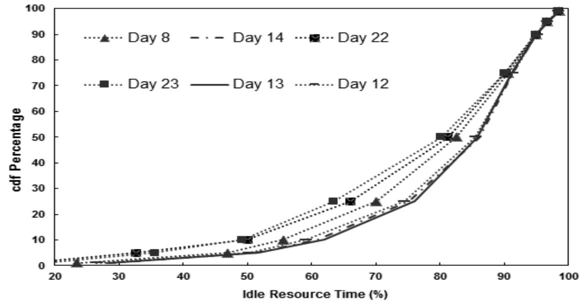

FIGURE 11. CPU idle time cdf distribution.

TABLE 6. CPU idleness statistics.

Day	Skewness	Median	Std.dev
Day 13	-2.1519	85.87	14.5
Day 12	-4.4658	85.57	15.18
Day 23	-3416.4	75.03	32.28
Day 22	-492.85	75.48	44.12
Day 14	-5730.2	81.35	68.83
Day 8	-680.39	82.49	72.12

respectively, supporting the observation that CPU resource time has a closer correlation than memory resource time among the execution sessions within a single day. This would infer that the idle CPU resource time has a better trend of predictability than the memory counterpart. However, within the idle resource time behaviours of the entire month, there are abrupt spikes in the idle CPU resource time at uneven intervals, while the idle memory resource time has an almost saturated curve. Day 8, 14, 22 and 23 are examples of abrupt spikes in idle CPU resource times, their memory counterparts do not show any notable fluctuations. There is no temporal correlation evident between the idle resource time duration and task submissions. Furthermore, on a day-wise observation, the behaviour of CPU cores show unpredictable trend of resource consumption among the co-located execution sessions which is in contrast to the trend of the memory resources. For this reason, we perform further investigation into days showing abnormal fluctuations of idle CPU resources among the co-located task execution session within the same day. We chose the days showing high resource time fluctuation (Day 8, 14, 22, 23) along with days characterised by minimum fluctuations (Day 13, 12) to observe the behavioural trend of CPU idleness. Fig. 11 illustrates the CDF of the idle CPU resource percentage and Table 6 presents the observed statistics for the selected days. It is evident that the CDF curves of all the observed days are negatively skewed, which is significant for the days showing increased fluctuations among the co-located sessions of the same day. The curve skewness is insignificant for the Day 12 and 13, insisting that the distribution curve is deviating away from normal with increasing fluctuation in the idle resource time among the co-located execution session. The median is observed to be between 75 to 85 percent for all of the days of interest, insisting the fact that at least half of the allocated CPU resources suffer 75% of minimal idleness. Since task submissions are loosely correlated with resource idleness, user requests are not observed to be affecting the presence of idle resources to a considerable margin.

VII. POWER CONSUMPTION ANALYSIS

This section presents our analysis of the power consumption incurred by the presence of zombie server resources. Such undesirable power consumption will increase the non-computing or overhead energy, which is an undeniable waste of input energy. The power consumed by idle CPU and memory resources has been computed based on their respective proportions of idleness during task execution based on the server capacities as shown in Table 3.

Based on the compute capacities of the servers, we classify the server platform A and C as mid-range servers and platform B as a high-spec server. From the trace log analysis, the idle resource times presented in section VI are measured across a total number of 12,500 active servers. We therefore consider 12,500 active servers across the datacentre, for profiling the power consumption of the zombie resources in an entire datacentre. The measured power consumptions of the idle CPU and memory resources are presented for the three server profiles in Fig. 12 accordingly. The average power consumption incurred by the idle CPU resource time across the observed period are 1.461, 4084 and 1343 (presented in kW-hours) respectively for platform A, B and C. It is notable that the power consumption levels of the servers are correspondingly higher with increasing process capacity. An interesting observation is that the CPU compute capacity of server platform B is around 6 times higher than the other two servers, but the idle CPU power consumption of server profile B is measured at only around 3 times more than profile A and C. This suggests that the power consumption levels of the active CPU resources are better optimised with increasing server capability, thereby proportionally reducing the power consumed. The idle resource time is the time during which the resource is either not being realised or is under-utilised, though the resource is still active and immediately available to use (not switched to any lower power "sleep" state). Such resource time could be effectively utilised by either increasing the intensity of task computation or

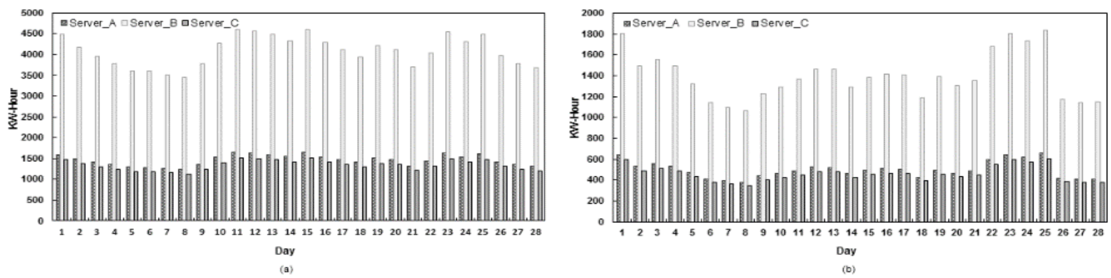

FIGURE 12. Power consumption of zombie resources (a) idle CPU and (b) idle memory.

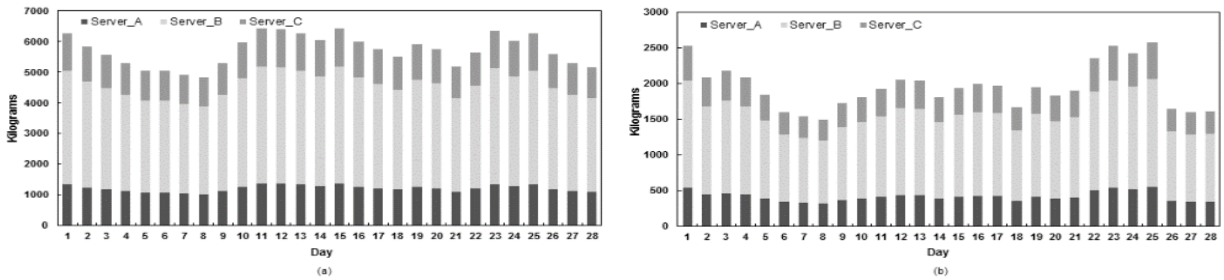

FIGURE 13. CO2 emissions of zombie resources (a) idle CPU and (b) idle memory.

increasing the server task load. Though both actions would increase the power consumption of the servers, their computation capacities can be effectively realised and an optimum utilisation of the available resources can be achieved.

The characteristics of idle memory resource time are viewed from a different perspectives to that of CPU, since the memory usage would most often exhibit less minimal fluctuation within a single execution session than those of CPU usage. The power consumption levels of the idle memory resource time are measured at 499, 1394 and 458 for server platform A, B and C respectively. It can be observed that the power consumption trend of memory resource is similar to that of CPU, with both exhibiting the trend of proportional power reduction of the active resources with increasing server capabilities. While the memory capacity of server platform B is 4 times larger than A and C, its idle memory resources are incurring a power consumption which is only 3 times higher. Further, the power consumption levels of the idle memory resources are observed to be much lower than that of idle CPU due to the fact that the idle memory resource times are observed to be just around one third of the idle CPU resource times. The input power to the servers also incurs power supply losses from AC/DC and DC/DC conversion, with AC/DC losses being much higher [42] than DC/DC losses. Besides the PSU, CPU and memory resources, the internal power consumption of the servers also includes power consumed by fans, drives, PCI cards, chip set etc. Though their power consumption levels are noticeable, we only acknowledge their presence but consider them to be outside of the scope of this paper.

VIII. ENVIRONMENTAL IMPACTS

The power consumption of zombie resources has a direct environmental impact through the means of their CO2 emissions. While the CO2 emissions created along the path of the power consumption cannot be completely avoided, there is seemingly always scope for reducing the level of emission. Some statistics insist that datacentres [43] can reduce CO2 emissions by an incredible 88% or more. This section presents the environmental implications of the zombie servers for investigating the potential for reducing the intensities of the datacentre CO2 emissions by the way of effectively reducing the presence of idle resources.

Fig. 13 presents the amounts of CO2 emissions incurred by the power consumed (kW-hours) by the idle CPU and memory resource time for the three server profiles under consideration. The CO2 emissions are measured according to the statistics [44] of the U.S. Energy Information Administration. The average CO2 emissions incurred by the power consumed during the idle CPU resource time are 1213, 3390 and 1115 (presented in kg) for server platforms A, B and C respectively. The average CO2 emissions of the memory counterpart are 414, 1157 and 381 respectively for the three servers.

Findings presented earlier in this paper mean that it is no surprise that idle CPU resource time has a significantly higher level of environmental implication than idle memory resource time. The manufacturer data also confirms that high-spec servers consume more power than mid-range servers, thereby resulting in increased amounts of C02 emissions. The statistics presented in Fig. 13 confirm that datacentres are one of the contributors toward global pollution, with their CO2 emissions known to be in the range of thousands of tonnes per annum. The CO2 emissions are measured based on the actual method of power generation; the worst case of emission occurs when the electricity is generated with conventional coal combustion. It has been reported that there is a possibility of zero CO2 emissions [45] for the electricity being generated from renewable energy sources such as nuclear and hydro-sources, especially if the emissions related to plant construction, use and maintenance are neglected.

IX. APPLICATION OF THE WORK

Although the analysis presented in this paper is based on the publicly available Google Cloud trace logs, we expand the scope of the applicability of this analysis by using the trace logs as a baseline for the energy profile of various commercial servers being widely used in 2016. The analysis presented in this paper provides insightful observations and inferences for the target audience and for researchers promoting green computing. We believe that our analysis can also be applied in similar computing environments such as transparent computing, grid computing and other parallel and distributed systems. The analysis presented in this paper will find applications in the following ways:

- To achieve an effective PUE. The PUE figure is usually the calculated ratio of the total facility power to the IT equipment power. By reducing the presence of zombie servers, Cloud providers can ensure there is no substantial overhead or non-compute energy incurred. Thus achieving an effective PUE by transforming most of the power consumption into useful compute energy. With the world average PUE [46] is reported as 1.7, there is still a lot of scope for enhancing the PUE of datacentres.
- To achieve optimum resource management. Cloud providers are intent on maximising profits by minimising production costs. Spotting server zombieness and allocating tasks accordingly can help to reduce the presence of low-level server utilisations. Issues such as red spots can be avoided in the datacentre, these usually incur excess energy consumption resulting from the increased operating temperature levels of individual servers.
- Reducing under-utilised resources and allocating optimum level of resources within individual VMs without affecting the job requirements can consolidate more number of VMs onto the physical servers to achieve maximum utilisation from minimum number of running servers.
- Prediction analysis. The characterisation of the user requests can be used to predict their resource requirements. The prediction of anticipated workloads in the near future would help providers with not only achieving energy-efficient resource scaling but also ensuring the availability of the requested resource to achieve effective QoS.

X. CONCLUSION

This paper has extensively analysed the energy related implications of the zombie servers caused by the presence of idle resource times whilst executing user requested tasks in large-scale Cloud environments based on power profiles of current servers used in large-scale datacentres. Three such server profiles have been used as examples to determine power consumption incurred by the presence of idle CPU and memory resource times. Our empirical analysis demonstrates that CPU resources account for the highest proportion of idle resource time at 75.6%, with memory resources exhibiting 25.5% of idleness. We further presented the environmental implications of idle resource time in terms of CO2 emissions in order to highlight the demand for energy efficient computing. Idle resource time is the result of not fully realising the complete compute potentials of IT resources; the idle resource time has the scope of being utilised to process a greater number of useful workloads in the same time period. Important inferences obtained in the analysis presented in this paper include:

- Idle resource time results from over-allocated resources. The proportion of idle resource time is high in every individual task execution session. Over-estimated resources cause the providers to over-allocate the server resources, which in turn increases the proportions of resource idleness.
- Power consumption levels are proportional to the server capabilities. We observe that the higher-spec server consumed 3 times more power than the mid-range servers for the same level of idle resource time given the fact that their CPU compute capacities are 6 times higher than those of the mid-range servers. We note that, in some cases, the higher-spec server offers more effective consolidation of workloads than the mid-range servers. As a result, resource idleness can be greatly reduced when the datacentre composition includes several higher-spec servers capable of effectively consolidating the workloads onto a minimum number of servers which are operating at near maximum utilisation levels.
- CPU behaviours varies more abruptly than memory resource behaviours. We observe abrupt spikes in the presence of idle resource times for CPU resources on a day-wise analysis, with the memory counter-part being almost saturated over the period of observation. Predicting the future workloads is a potential way of achieving optimum resource scaling by reducing idle resource times. However, while it offers the greatest benefit, predicting CPU idleness is more complex than predicting the memory idleness trend.
- Cloud workloads are highly heterogeneous in terms of the CPU and memory consumption and the encompassed task length within a single job.

As a future work, we plan to develop a prediction model to estimate the user requests anticipated in the near future.

The objective is to achieve optimum resource scaling, ultimately to promote energy-efficient computing by reducing the server zombiness. Another extension of the analysis presented in this paper would be predicting the resource consumption levels of individual task requests, which would help allocating the optimum levels of server resources for processing. Finally, we plan to study the characteristics of other distributed processing environments such as grid, mobile cloud and transparent computing systems with the motivations of promoting green computing.

REFERENCES

[1] S.-Y. Jing, S. Ali, K. She, and Y. Zhong, "State-of-the-art research study for green cloud computing," *J. Supercomput.*, vol. 65, no. 1, pp. 445–468, Dec. 2011.
[2] J. Patel, V. Jindal, J. Xu, and P. Garraghan, "Workload estimation for improving resource management decisions in the cloud," in *Proc. IEEE Int. Symp. Auto. Decentralized Syst.*, Mar. 2015, pp. 25–32.
[3] A. Khan, X. Yan, S. Tao, and N. Anerousis, "Workload characterization and prediction in the cloud: A multiple time series approach," in *Proc. IEEE Netw. Oper. Manage. Symp.*, Apr. 2012, pp. 1287–1294.
[4] J. Whitney and P. Delforge, "Scaling up energy efficiency across the data center industry: Evaluating key drivers and barriers," NRDC, New York, NY, USA, Tech. Rep. IP:14-08-a, Apr. 2014.
[5] R. Buyya, A. Beloglazov, and J. Abawajy, "Energy-efficient management of data center resources for cloud computing: A vision, architectural elements, and open challenges," in *Proc. Int. Conf. Parallel Distrib. Process. Tech. Appl.*, Jul. 2010, pp. 6–20.
[6] Y. Jin, Y. Wen, Q. Chen, and Z. Zhu, "An empirical investigation of the impact of server virtualization on energy efficiency for green data center," *Comput. J.*, vol. 56, no. 8, pp. 977–990, Feb. 2013.
[7] Z. Lu, S. Takashige, Y. Sugita, T. Morimura, and Y. Kudo, "An analysis and comparison of cloud data center energy efficient resource management technology," *Int. J. Service Comput.*, vol. 2, no. 4, pp. 32–51, Dec. 2014.
[8] I. S. Moreno, P. Garraghan, P. Townend, and J. Xu, "An approach for characterizing workloads in Google cloud to derive realistic resource utilization models," in *Proc. IEEE 7th Int. Symp. Service-Oriented Syst. Eng.*, Mar. 2013, pp. 49–60.
[9] A. Beloglazov and R. Buyya, "Energy efficient resource management in virtualized cloud data centers," in *Proc. 10th IEEE/ACM Int. Conf. Cluster, Cloud, Grid Comput.*, May 2010, pp. 826–831.
[10] I. S. Moreno and J. Xu, "Customer-aware resource overallocation to improve energy efficiency in real-time cloud computing data centers," in *Proc. IEEE Int. Conf. Service-Oriented Comput. Appl.*, Dec. 2011, pp. 12–14.
[11] S. U. Sharma and Y. B. Gandole, "Virtualization approach to reduce network latency for thin client performance optimization in cloud computing environment," in *Proc. Int. Conf. Comput. Commun. Inform. (ICCCI)*, Jan. 2014, pp. 1–6.
[12] B. Ciciani *et al.*, "Automated workload characterization in cloud-based transactional data grids," in *Proc. IEEE 26th Int. Parallel Distrib. Process. Symp. Workshops Ph.D. Forum (IPDPSW)*, May 2012, pp. 1525–1533.
[13] P. Garraghan, X. Ouyang, P. Townend, and J. Xu, "Timely long tail identification through agent based monitoring and analytics," in *Proc. 18th Int. Symp. Real-Time Distrib. Comput.*, Apr. 2015, pp. 19–26.
[14] P. Garraghan, I. S. Moreno, P. Townend, and J. Xu, "An analysis of failure-related energy waste in a large-scale cloud environment," *IEEE Trans. Emerg. Topics Comput.*, vol. 2, no. 2, pp. 166–180, Jun. 2014.
[15] R. Yang, I. S. Moreno, J. Xu, and T. Wo, "An analysis of performance interference effects on energy-efficiency of virtualized cloud environments," in *Proc. 5th IEEE Int. Conf. Cloud Comput. Technol. Sci.*, Dec. 2013, pp. 112–119.
[16] T. Ropars, A. Guermouche, B. Uçar, E. Meneses, L. Kalé, and F. Cappello, "On the use of cluster-based partial message logging to improve fault tolerance for MPI HPC applications," in *Proc. Int. Conf. Parallel Process.*, vol. 1. 2011, pp. 567–578.
[17] J. Panneerselvam, L. Liu, N. Antonopoulos, and M. Trovati, "Latency-aware empirical analysis of the workloads for reducing excess energy consumptions at cloud datacentres," in *Proc. IEEE 11th Int. Symp. Service-Oriented Syst. Eng.*, Mar. 2016, pp. 62–70.
[18] F. Chen, J. G. Schneider, Y. Yang, J. Grundy, and Q. He, "An energy consumption model and analysis tool for cloud computing environments," in *Proc. 1st Int. Workshop Green Sustain. Softw. (GREENS)*, Jun. 2012, pp. 45–50.
[19] P. Garraghan, P. Townend, and J. Xu, "An empirical failure-analysis of a large-scale cloud computing environment," in *Proc. IEEE 15th Int. Symp. High-Assurance Syst. Eng. (HASE)*, Jan. 2014, pp. 113–120.
[20] B. Schroeder and G. Gibson, "A large-scale study of failures in high-performance computing systems," *IEEE Trans. Dependable Secure Comput.*, vol. 7, no. 4, pp. 337–350, Oct. 2010.
[21] T. Nguyen and W. Shi, "Improving resource efficiency in data centers using reputation-based resource selection," in *Proc. Int. Green Comput. Conf.*, Aug. 2010, pp. 389–396.
[22] J.-A. Quiané-Ruiz, C. Pinkel, J. Schad, and J. Dittrich, "RAFTing MapReduce: Fast recovery on the RAFT," in *Proc. IEEE 27th Int. Conf. Data Eng.*, Apr. 2011, pp. 589–600.
[23] Y. Liang, Y. Zhang, A. Sivasubramaniam, M. Jette, and R. Sahoo, "BlueGene/L failure analysis and prediction models," in *Proc. Int. Conf. Dependable Syst. Netw. (DSN)*, Jun. 2006, pp. 425–434.
[24] H. Li, D. Groep, L. Wolters, and J. Templon, "Job failure analysis and its implications in a large-scale production grid," in *Proc. 2nd Int. Conf. e-Sci. Grid Comput.*, Dec. 2006, pp. 27–35.
[25] Z. Wan, "Sub-millisecond level latency sensitive cloud computing infrastructure," in *Proc. Int. Congr. Ultra Modern Telecomm. Control Syst. Workshops*, Oct. 2010, pp. 1194–1197.
[26] M. S. Bali and S. Khurana, "Effect of latency on network and end user domains in Cloud computing," in *Proc. Int. Conf. Green Comput., Commun. Conservation Energy*, Dec. 2013, pp. 777–782.
[27] Z. Wan, "Cloud computing infrastructure for latency sensitive applications," in *Proc. 12th IEEE Int. Conf. Commun. Technol.*, Nov. 2010, pp. 1399–1402.
[28] Z. Wan, P. Wang, J. Liu, and W. Tang, "Power-aware cloud computing infrastructure for latency-sensitive Internet-of-Things services," in *Proc. 15th Int. Conf. Comput. Modelling Simulation (UKSim)*, Apr. 2013, pp. 617–621.
[29] I. S. Moreno, R. Yang, J. Xu, and T. Wo, "Improved energy-efficiency in cloud datacenters with interference-aware virtual machine placement," in *Proc. IEEE 11th Int. Symp. Auto. Decentralized Syst. (ISADS)*, Mar. 2013, pp. 1–8.
[30] J. Baliga, R. W. A. Ayre, K. Hinton, and R. S. Tucker, "Green cloud computing: Balancing energy in processing, storage, and transport," *Proc. IEEE*, vol. 99, no. 1, pp. 149–167, Dec. 2010.
[31] T. Lin, T. Alpcan, and K. Hinton, "A game-theoretic analysis of energy efficiency and performance for cloud computing in communication networks," *IEEE Syst. J.*, vol. 11, no. 2, pp. 649–660, Jun. 2015.
[32] S. Shaoling, Z. Jing, C. Moliang, R. Hui, C. Yi, and F. Xiaodong, "Network energy consumption analysis and dormancy mechanism based on ant colony algorithm in cloud computing environment for IOT service and real-time embedded industrial control system," in *Proc. 27th Chin. Control Decision Conf.*, May 2015, pp. 5568–5572.
[33] C. Reiss, A. Tumanov, G. R. Ganger, R. H. Katz, and M. A. Kozuch, "Towards understanding heterogeneous clouds at scale: Google trace analysis," Carnegie Mellon Univ., Pittsburgh, PA, USA, Tech. Rep. ISTC-CC-TR-12-101, Apr. 2012.
[34] J. Panneerselvam, L. Liu, N. Antonopoulos, and Y. Bo, "Workload analysis for the scope of user demand prediction model evaluations in cloud environments," in *Proc. IEEE/ACM 7th Int. Conf. Utility Cloud Comput. (UCC)*, Dec. 2014, pp. 883–889.
[35] Google. *Google Cluster Data V2*. Accessed on Jan. 10, 2016. [Online]. Available: https://github.com/google/clusterdata/blob/master/ClusterData2011_2.md
[36] C. Reiss and J. Wilkes, *Google Cluster-Usage Traces: Format + Schema*. Menlo Park, CA, USA: Google Inc, 2013.
[37] I. S. Moreno, P. Garraghan, P. Townend, and J. Xu, "Analysis, modeling and simulation of workload patterns in a large-scale utility cloud," *IEEE Trans. Cloud Comput.*, vol. 2, no. 2, pp. 208–221, Apr. 2014.
[38] P. Garraghan, P. Townend, and J. Xu, "An analysis of the server characteristics and resource utilization in Google cloud," in *Proc. IEEE Int. Conf. Cloud Eng.*, Mar. 2013, pp. 124–131.
[39] Google. *Efficiency: How We Do it*. Accessed on Feb. 15, 2016. [Online]. Available: https://www.google.com/about/datacenters/efficiency/internal/
[40] *Solutions for the Next Generation Data Center*. Accessed on Feb. 20, 2016. [Online]. Available: http://www.42u.com/measurement/pue-dcie.htm

[41] SPEC. (2016). *SPECpower_SSJ2008 Results SPEC*. [Online]. Available: https://www.spec.org/sfs2008/results/
[42] Vertatique. (2015). *Average Power Use Per Server*. [Online]. Available: http://www.vertatique.com/average-power-use-server
[43] Standford. (2013). *Data Centrers Can Slash CO2 Emissions 88% or More*. [Online]. Available: https://energy.stanford.edu/news/data-centers-can-slash-co2-emissions-88-or-more
[44] (2016). *Independent Statistics and Analysis*. [Online]. Available: http://www.eia.gov
[45] Infrarati. (2014). *Data Center CO2 Emissions*. [Online]. Available: https://infrarati.wordpress.com/2014/03/25/data-center-co2-emissions/
[46] Digitalisation World. (2015). *Adder Technology Helps Data Centre Provider Achieve Highest Environmental Efficiency*. [Online]. Available: https://digitalisationworld.com/article/46717/blogsall

JOHN PANNEERSELVAM is currently pursuing the Ph.D. degree with the University of Derby, U.K. He is also a Lecturer in computing with the University of Derby. He has published his recent research works in notable peer reviewed international conferences, journals and as book chapters. His current research is focused on energy efficient cloud systems. His research interests include cloud computing, big data analytics, opportunistic networking, and P2P Computing. He is an Active Member of the IEEE.

LU LIU (M'07) received the Ph.D. degree from the University of Surrey. He is currently the Head of the Department of Electronics, Computing and Mathematics, University of Derby, and an Adjunct Professor with the School of Computer Science and Communication Engineering, Jiangsu University. His research interests are in areas of cloud computing, social computing, service-oriented computing, and peer-to-peer computing. He is a Fellow of the British Computer Society.

JAMES HARDY is currently pursuing the Ph.D. degree with the University of Derby, U.K. He is also a Lecturer in computer networks with the University of Derby. In addition he holds manufacturer specific data networking qualifications. He has authored or co-authored papers on traffic control, Internet of People, virtualization, IPv6 addressing, software aging, and web services location. He has industrial experience, including aerospace control systems and global scale wide area data networking. His research interests include smart city transportation, autonomous vehicles, communication systems, green computing, HaaS, IaaS, control systems, simulation, and virtualization.

NICK ANTONOPOULOS received the Ph.D. degree in computer science from the University of Surrey in 2000. He is currently the Pro Vice-Chancellor of Research with the University of Derby and the University of Derby Technical Coordinator of the framework collaboration with CERN and the ALICE experiment. He has over 18 years of academic experience and has authored over 150 articles in fully refereed journals and international conferences. His research interests include cloud computing, P2P computing, software agent architectures, and security.