Modelling and Analysing the Power Consumption of Idle Servers

Robert Basmadjian
University of Passau
Innstrasse 43
Passau, Germany
basmadji@fim.uni-passau.de

Florian Niedermeier
University of Passau
Innstrasse 43
Passau, Germany
niederme@fim.uni-passau.de

Hermann De Meer
University of Passau
Innstrasse 43
Passau, Germany
demeer@fim.uni-passau.de

***Abstract*—To the best of our knowledge, there have been no efforts in devising power consumption prediction models for an idle server, where this latter contributes approximately 66% of the maximum power drain. In this paper, we propose power consumption prediction models for idle servers by taking into account their constituent components such as processor, memory, hard disk, fan and power supply unit. To this end, we identify the relevant energy-related attributes of each component necessary for the idle power consumption predictions. Furthermore, based on the proposed models, we provide an in-depth analysis by considering several types of servers (e.g. rackable, blade, etc) having different hardware characteristics and energy-aware mechanisms.**

***Keywords*-modelling; idle power consumption; servers; data centres;**

I. Introduction

Minimising the energy consumption of data centres has attracted a great deal of attention lately due to *ecological* (reduction of CO_2 emissions) and *economical* (increase of electricity costs) reasons. It was shown in [8] that energy usage of servers and data centres in U.S. doubled between 2000 and 2006 and it reached in the final year about 1.5% of the national electricity consumption. Moreover, it was shown that it was poised to double again by 2011.

This increase is correlated to the fact that servers consume about 70% of the total energy consumption of data centres [9]. Furthermore, since the workload of data centres fluctuates on a weekly basis[1], it is usually the case to over-provision computing resources to cope with peak utilisation. To this regard, Liu et al. [10] argued that only 30% of servers in data centres are fully utilised while keeping the other 70% in idle state. Moreover, idle servers consume between 60% and 66% of the peak load power consumption [11]. Hence, one of the motivations to reduce the energy consumption of data centres is to shut down unnecessary resources without violating the Service Level Agreements (SLA).

[1]Hourly basis in the case of cloud computing such as public clouds

Recently in [12], the authors proposed energy-aware optimisation algorithms for cloud computing by taking advantage of virtualisation and consolidation concepts. The basic idea of such optimisation algorithms is to consolidate workloads and to turn off unutilised resources (e.g. idle servers). However, in order that such algorithms can take the most suitable decisions, they are supposed to be guided by models that estimate the power consumption of servers.

In general, the overall power consumption of servers consists of two parts: idle and dynamic. The former is related to the case where the server is not active, whereas the latter is concerned with the situation where certain computations are being performed. To this end, several models [13], [14], [15], [16] have been proposed in the literature for estimating the total power consumption of servers. Basically such models assume that the idle power is constant and known in advance. This assumption does not hold true in the case of heterogenous environments where servers differ from each other in terms of installed hardware and implemented energy-saving mechanisms (e.g. Intel SpeedStep and AMD Cool'n'Quiet).

In this paper, we go one step further and propose a model to estimate the power consumption of idle servers. More precisely, we provide models for multi-core processors (by taking into account their energy-saving mechanisms), memory, hard disks, fans, and power supply units. For this purpose, we present a schema for idle servers in the form of Universal Modeling Language (UML) class diagrams, where we identify the relevant energy-related attributes. Those attributes, which build the basis for the power consumption prediction models, can be extracted from the manufacturers' data sheet, and hence makes the proposed models appropriate to any energy-aware optimisation algorithm. Furthermore, we analyse several types of servers having different hardware characteristics. The analysis is presented in the form of power consumption breakdown with respect to the servers' constituent components.

The rest of this paper is organised as follows: Section III presents the identified energy-related attributes in the

978-3-901882-46-3

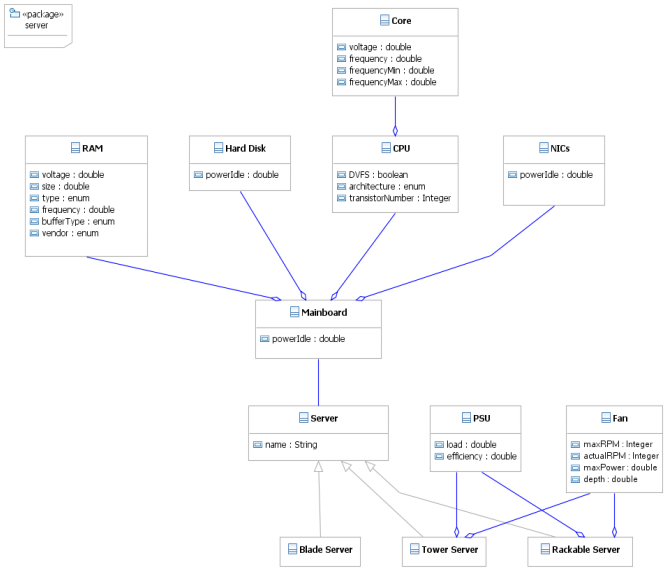

Figure 1. Idle Server UML Class Diagram

form of UML class diagrams for modelling idle servers. In Section IV, we introduce our proposed power consumption estimation models. The obtained evaluation results are given in Section V. The paper is concluded in Section VI.

II. Related Work

A variety of power consumption models both for single-core [24], [25], [26], [27], [28] and multi-core [29], [30] processors have been proposed in the literature. In our recent work [31], we provided a generic power consumption model for multi-core processors by circumventing the drawbacks of the above mentioned models. However, all the aforementioned models estimate only the dynamic power consumption and consider that the idle power of a processor is already known in advance. In [13], [14], [15] and [16], the authors proposed models for servers by taking into account processor's utilisation. Nevertheless, the proposed models also assume that the idle power of the servers is given. On the other hand, certain manufacturers provide the idle power consumption of the overall server. However, in most cases the provided numbers are rough estimates and hardware dependent. In this paper, we address this aspect and estimate the idle power consumption of servers by breaking it down into the power consumption of their constituent components.

III. Idle Server Schema

Figure 1 illustrates the UML class diagram of an idle server. The identified attributes of classes form the basis for the power consumption prediction models of Section IV.

The `Server` class presents an abstraction of a computing system and consists of different generalisations (e.g. `Blade`, `Tower`, and `Rackable`), each having its own physical form factor. Typically, a server consists of a `Mainboard` which is the central printed circuit board providing most of the core hardware components of the system, plus connectors to other peripheral devices external to the mainboard. Its *powerIdle* attribute denotes the power consumption of the mainboard when the server is in idle state. The main components attached to the `Mainboard` are: Central Processing Unit (`CPU`), Random Access Memory (`RAM`), Network Interface Card (`NIC`) and `Hard Disk`.

Since the advent of modern processors, a `CPU` commonly consists of multiple independent units (`Core`). The most relevant energy-related attributes of the `CPU` are:

1) *DVFS* (Dynamic Voltage and Frequency Scaling): is used to indicate whether corresponding energy-saving mechanisms are enabled or not;
2) *architecture* : indicates the processor's manufacturer (e.g. Intel, AMD, etc.); it's relevant due to different power consumption behaviours based on the manufacturer's corresponding energy-saving mechanisms (see Section IV-A);
3) *transistorNumber* : denotes the number of transistors in the order of millions.

Each `Core` operates at a certain *frequency* (GHz) and *voltage* (Volt) which have the major impact on the power consumption behaviour of multi-core processors. Hence, monitoring systems should keep these attributes up-to-date.

The `RAM` class has several attributes relevant to power consumption estimation:

1) *voltage*: presents the supply voltage under which the memory module operates; it is highly dependent on the memory *type* (e.g. DDR_2, DDR_3, etc.);
2) *size* (GiB) and *frequency* (MHz) denote respectively the memory module's size and operational frequency;
3) *bufferType*: is the type of buffer technology (e.g. fully buffered, buffered, registered, and unbuffered);
4) *vendor* presents the manufacturer (e.g. KINGSTON, HYNIX, etc.).

The `Hard Disk` and `NIC` both have the attribute *powerIdle* whose value can be found in the specifications' data sheet.

`Tower Servers` and `Rackable Servers` are equipped with their own Power Supply Units (`PSU`) and `Fans`. The most relevant energy-related attribute of a `PSU` is the *efficiency*, which indicates (in percent) the amount of loss of the power supplied to server components. It is

highly related to the value of the *load* attribute. Note that values of efficiency corresponding to given loads can be extracted from the manufacturer's data sheet. Inside the `Fan` class, *depth* denotes the depth (in meter) of a fan, whereas *maxRPM* and *maxPower* respectively indicate the maximum nominal values of rotations per minute and power consumption for the fan. *actualRPM* shows the actual instantaneous rotation speed of the fan; it's a dynamic attribute, whose value should be kept up-to-date by monitoring systems.

IV. Power Consumption Prediction Models

In this section, we present the idle power consumption models for different components (e.g. processor, memory, hard disk, etc.) of a server based on the UML class diagram of Figure 1.

A. Processor

With the advent of multi-core processors (e.g. dual-, quad-, hexa-core, etc.), several energy-saving mechanisms (e.g. C-states[2]) have been introduced which play a major role in reducing the idle power consumption of processors. Consequently, we first study the idle power consumption behaviour of multi-core processors with energy-saving mechanisms disabled, and then provide a generic model while enabling those mechanisms.

1) Without Energy-Saving Mechanisms: The power consumption of a processor can be determined by using the following formula derived from Joule's and Ohm's laws [1]:

$$P = I * V, \tag{1}$$

where P denotes the power (Watt), I represents the electric current (Ampere) and V indicates the voltage (Volt).

When a multi-core processor is in idle state, all of its constituent cores are respectively also inactive. As a matter of fact, we use Equation (1) at core level and assume that each core contributes equally to the total idle power consumption of a processor:

$$P_i = I_i * V_i, \tag{2}$$

where P_i, I_i, and V_i denote respectively the power, current and voltage of the corresponding core i. Since the power consumption of each core depends upon its number of transistors, we go one step further and compute the idle power of each core at transistor level. The higher the number of transistors a core has, the higher its idle power consumption. In order to compute the power consumption of transistors, we use Equation (1) at transistors level. For j^{th} transistor (in the order of millions) inside the i^{th} core, the power consumption is given by:

$$P_{ji} = I_{ji} * V_{ji}, \tag{3}$$

where I_{ji} and V_{ji} denote respectively the current and voltage of the j^{th} transistor of the core i. Note that the voltage V_{ji} is in general the same as the one V_i of the core i which on its turn is the same as the voltage of the processor.

By analysing the Current-Voltage characteristics of Figure 2, we notice that between 0 and 2V [3], the relationship is almost linear [7]. Given this fact, we adopt the curve-fitting methodology to model current leakage of a transistor using second order polynomial as:

$$I_{ji} = \alpha V_{ji}^2 - \beta V_{ji} + \gamma, \tag{4}$$

where $\alpha = 0.114312$, $\beta = 0.22835$ and $\gamma = 0.139204$ are the coefficients. Note that the values of α, β and γ are derived based on the current leakage I and corresponding voltage V obtained while analysing a Quad-core AMD processor. The results of this analysis are shown in Figure 3 such that "Act. I" denotes the real current leakage obtained from a power meter [5] and "Est. I" represents the estimated current leakage obtained by Equation (4).

Let t_i denote the total number of transistors (in the order of millions) of a core i, then its power consumption

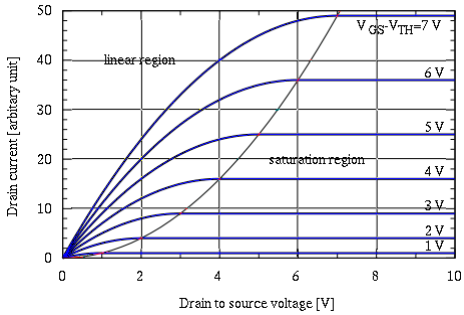

Figure 2. Current-Voltage Characteristics of MOSFET[2]

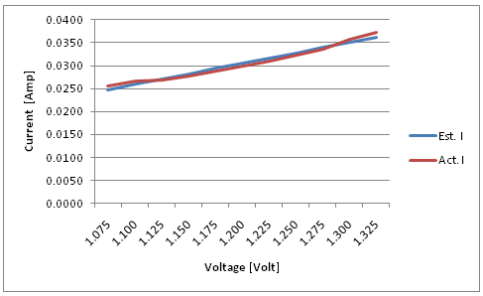

Figure 3. Estimated and Actual Voltage to Current Relationship per Million Transistors

[2]Is a technology enabling to configure for a processor one of a set of power related saving modes.

[3]Is the normal voltage operating range for processors

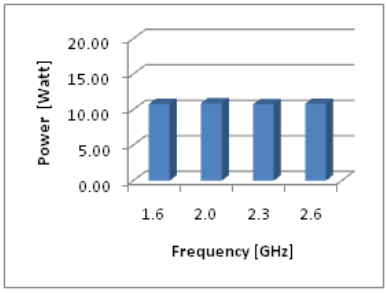

Figure 4. The Impact of Frequency on Intel SpeedStep for a Given Voltage

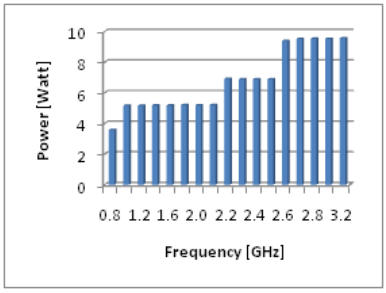

Figure 5. The Impact of Frequency on AMD Cool'n'Quiet for a Given Voltage

is given by:

$$P_i = \sum_{j=1}^{t_i} I_{ji} * V_{ji}, \tag{5}$$

where P_i is the power consumption of the i^{th} core. Hence, the idle power consumption of a multi-core processor having n cores is the sum of the power of each of its constituent cores:

$$P_{CPU} = \sum_{i=1}^{n} P_i \tag{6}$$

2) With Energy-Saving Mechanisms: With the emergence of energy-saving mechanisms (e.g. Intel SpeedStep [3] and AMD Cool'n'Quiet [4]), the idle power consumption of a core (processor) decreases. Basically, this is achieved by decreasing the voltage and/or frequency (DVFS) of a core (processor). In order to demonstrate such an impact, we propose the following model based on Equation (5):

$$P'_i = \delta_i * P_i, \tag{7}$$

where δ_i is the factor for reduction in the power consumption P_i of core i, whereas P'_i represents the reduced power consumption of a core i. Note that δ_i can vary depending upon the corresponding energy-saving mechanisms, where each of such mechanism has its own particular behaviour. To this end, we analysed Intel SpeedStep and AMD Cool'n'Quiet, and conclude that both have different behaviour as shown in Figures 4 and 5. Furthermore, as illustrated in Figure 6, frequency F (GHz) and voltage V are the main parameters playing the major role in decreasing the power consumption. Hence, we divide the modelling of δ_i into two phases: impact of voltage as well as frequency scaling.

Impact of Voltage Scaling: In order to investigate the impact of voltage scaling, we altered the voltage through BIOS and observed the idle power consumption as shown in Figure 6. Note that the four curves from bottom to up denote respectively the power consumption of a core with increasing order of frequencies (from 0.8 GHz to 3.2 GHz). Each experiment was carried out for

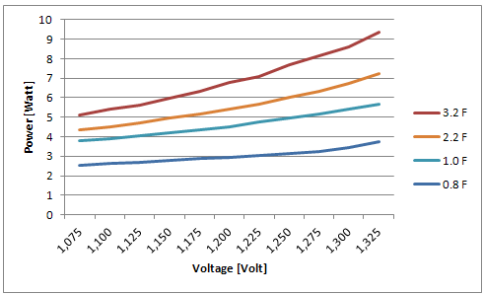

Figure 6. The Impact of Voltage and Frequency Scaling

50 minutes, where during that period the server was kept idle. We performed these observations for a long period of time in order to ensure that each component goes into idle state. The idle power consumption is obtained from a power meter [5], which takes 10 samples per second. The results of Figure 6 are obtained by observing an AMD Quad-core processor having minimum and maximum frequencies of 0.8 GHz and 3.2 GHz respectively.

It can be seen from Figure 6 that the impact of voltage is not linear. Due to this behaviour, we adopt curve-fitting methodology and propose a second order polynomial equation based on the voltage of the core i in the following manner:

$$\delta_i = \beta_{i1} V_i^2 + \beta_{i2} V_i + \beta_{i3}, \tag{8}$$

where β_{ij} are the coefficients whose values are given in Table I such that $i \in \{1, \ldots, n\}$ and $j \in \{1, \ldots, 3\}$ and V_i denotes the corresponding voltage of the core i, whereas n denotes the processor's total number of cores.

Impact of Frequency Scaling: To analyse the frequency scaling, the frequency[4] of each core is modified by editing the following file in Linux operating system: /sys/devices/system/cpu/cpui/cpufreq/scaling_min_frequency, where cpui denotes the corresponding core i. The experiments' setup (e.g. duration, power meter and number of samples per second) is the same as the

[4]Automatic voltage and frequency scaling were disabled during the observations

Table I
Coefficients of the Corresponding Frequencies

Frequency Range	β_{i1}	β_{i2}	β_{i3}
$f_{min} < f_i \leq f_{0.25}$	0.466201	−1.53706	1.60751
$f_{0.25} < f_i \leq f_{0.68}$	−1.13753	2.32643	−0.48657
$f_{0.68} < f_i \leq f_{0.81}$	−0.37296	0.64056	0.56800
$f_{0.81} < f_i \leq f_{max}$	−0.80187	2.00811	−0.26095

Table II
Power Reduction Factors δ_i for Intel Processors

Processor Type	δ_i
Dual-core	0.942
Quad-core	0.728
Hexa-core	0.316

one performed for voltage scaling. The analysis of both energy saving mechanisms (e.g. Intel SpeedStep including Hexa-core's C-states and AMD Cool'n'Quiet) for a given voltage and the corresponding frequencies are illustrated in Figures 4 and 5. It can be observed in Figure 4 that Intel SpeedStep including C-states has very small (negligible) impact on the idle power consumption by changing the frequency (i.e. the idle power remains nearly the same). On the other hand, the idle power consumption changes with altering the frequency for AMD Cool'n'Quiet as illustrated in Figure 5. Furthermore, we notice that during specific frequency ranges (e.g. $f_i = 0.8$, $0.8 < f_i < 2.2$, $2.2 \leq f_i < 2.6$ and $2.6 \leq f_i \leq 3.2$), the idle power consumption of a core remains the same. Consequently, we divide the frequency spectrum (from minimum to maximum) into four different frequency ranges (in GHz): $f_i \leq f_{0.25}$, $f_{0.25} < f_i \leq f_{0.68}$, $f_{0.68} < f_i \leq f_{0.81}$, and $f_{0.81} < f_i \leq f_{max}$, such that $f_{0.x}$ denotes the x% of the maximum frequency f_{max}. Note that AMD Cool'n'Quiet has partially similar behaviour as Intel SpeedStep within a certain frequency range where the change in frequency has no impact on the idle power consumption. Consequently, the four curves of Figure 6 were used to derive the coefficients of Equation (8) which are given in Table I.

In order to provide an overview of reduction factor, in Table II we give values of δ_i for multi-core Intel processors. Certain processors, for instance, Intel dual- and quad-cores do not possess advanced C-states (e.g. C6). Hence, the energy reduction factor for processors (e.g. hexa-core) having such states is significantly different (more savings) from the others.

Given a processor composed of n cores with a specific energy-saving mechanism, then its idle power consumption is given by:

$$P_{CPU} = \sum_{i=1}^{n} P'_i \tag{9}$$

where P'_i is introduced in Equation (7).

B. Memory

A Random Access Memory (RAM) consumes power in its idle[5] state by refreshing those ranks that hold certain stored data. In this paper, we focus on Synchronous Dynamic RAM. Furthermore, since DDR is obsolete in the sense that we can rarely find such technology in modern servers, in the rest of this paper, we provide models for DDR_2 and DDR_3 technologies.

The power consumption of a memory can be determined by using the following formula [1]:

$$P = I * V, \tag{10}$$

where parameters of Equation (10) are introduced in (1). As mentioned in Section IV-A, it was shown in [7] that there is a linear relationship between the current I and voltage V when the supplied voltage is between 0 and 2V (which is typically the case for DDR_2 and DDR_3 memory technologies). Hence, the current can be expressed in the following manner:

$$I = c * V, \tag{11}$$

where c is a constant introduced later in this section. Taking Equations (10) and (11) into account, the idle power consumption of a memory module for a given frequency f (MHz) and size s (GiB) can be rewritten in the following way:

$$P(f, s) = c * V^2. \tag{12}$$

Consequently, in order to reflect the impact of frequency f on the idle power consumption, Equation (12) can be expressed as:

$$P(s) = f * c * V^2. \tag{13}$$

Furthermore, in order to show the influence of size s (GiB) on the idle power consumption, Equation (13) can be given as:

$$P = s * f * c * V^2. \tag{14}$$

Given a set of n memory modules, then their idle power consumption is expressed as:

$$P_{RAM} = \sum_{i=1}^{n} s_i * f_i * c * V_i^2, \tag{15}$$

where s_i, f_i and V_i denote respectively the *size* (GiB), frequency (MHz) and the voltage (Volt) of a specific memory module i, whereas c takes a value of 0.00043 and 0.00013 for DDR_2 and DDR_3 respectively.

[5]Neither read nor write operations take place

C. Hard Disk

Typically, the hard disk is in *idle mode* when no activity (read or write operations) is carried out. Based on our observations performed on different families of hard disks, we noticed that the idle mode power consumption can be further split into three states[6]: *idle*, *standby* and *sleep*. Moreover, we noticed that the power consumption due to standby and sleep states is quite identical and it is in average 10% of the idle state power consumption. This is due to the fact that during standby and sleep states, the disk's mechanical parts are significantly shutdown. Then, the *idle mode* power consumption of the hard disk is given by:

$$P_{HDD} = P_{idle}(\alpha + 0.1 * \beta + 0.1 * \gamma), \tag{16}$$

such that P_{idle} is the idle state power consumption provided by the manufacturer's data sheet, whereas $\alpha, \beta, \gamma \in [0, 1]$ indicate respectively the probability that the disk is in idle, standby and sleep states. The values of α, β, and γ are chosen based on the operating system's configuration such that $\alpha + \beta + \gamma = 1$. For instance, in our case, we set $\alpha = 0.1$ and $\beta = \gamma = 0.45$, since the operating system is configured to put the hard disk into standby and sleep modes after 2 minutes of inactivity. Note that whenever it is possible to detect the exact state (e.g. idle, standby, sleep) of the hard disk, then the parameters $\alpha, \beta, \gamma \in [0, 1]$ can be configured appropriately such that always two of such parameters have a value of zero.

D. Mainboard

The idle power consumption of the mainboard together with its constituent components is given by:

$$P_{Mainboard} = \sum_{i=1}^{l} P_{CPU} + P_{RAM} + \sum_{j=1}^{m} P_{NIC} + \sum_{k=1}^{n} P_{HDD} + c, \tag{17}$$

where P_{CPU}, P_{RAM}, and P_{HDD} are introduced respectively in Equations (9), (15), and (16), whereas c is constant related to the mainboard's own power consumption. Technically, it is challenging to compute the power consumption of the mainboard. Hence, statistical values for c can be derived based on observations and server types (see Section V). Thus, the value of c can be configured through *powerIdle* attribute of the `Mainboard` class (see Section III). Finally, P_{NIC} denotes the idle power consumption of the network interface cards whose value can be found in the specifications data sheet.

[6]The hard disk changes its state sequentially from idle to standby and then to sleep.

E. Fan

The power consumption of a fan has a tight correlation with its Revolutions Per Minute (RPM): the higher the RPM is, the more power it consumes. Consequently, a model is derived based on the following formula for the power consumption of fans [23]:

$$P = d_p * q, \tag{18}$$

where P denotes the power consumption (Watt), d_p indicates the total pressure increase in the fan (Pa or N/m^2), and q represents the air volume flow delivered by the fan (m^3/s). Hence, replacing d_p by $\frac{F}{A}$ and q by $\frac{V}{t}$ in Equation (18), we obtain:

$$P = \frac{F}{A} * \frac{V}{t}, \tag{19}$$

where F, A, V and t denote respectively the force (N), area (m^2), volume (m^3) and time (seconds). By a simple simplification of volume V and area A, we obtain the following equation:

$$P = \frac{F * d}{t}, \tag{20}$$

where d indicates the depth of the fan (meter). Based on our observations performed on a set of fans and by adjusting the speed of their RPM appropriately, we found out that F is proportional to the square of the RPM:

$$F = c_{fan} * RPM^2. \tag{21}$$

By taking into account Equations (20) and (21), the power consumption model for the fan is given by:

$$P_{Fan} = \frac{c_{fan} * RPM^2 * d}{3600}, \tag{22}$$

where RPM denotes the actual instantaneous revolution per minute of the fan (*actualRPM* in Section III) whose value should be kept up-to-date through the monitoring system. Note that for a given fan, the value of c_{fan} remains constant. As a matter of fact, we compute the value of c_{fan} based on Equation (22):

$$c_{fan} = \frac{3600 * P_{max}}{RPM_{max}^2 * d}, \tag{23}$$

where P_{max} and RPM_{max} denote respectively the maximum power consumption and rotations per minute of the fan whose values can be extracted, in addition to the depth d, from the manufacturer's data sheet.

F. Power Supply Unit

Basically, the power consumed by the PSU itself (loss) is highly dependent on its efficiency: the higher PSU's efficiency is, the less power it consumes under the same load. To this regard, the PSU manufacturers provide the efficiency range with respect to a given PSU load. Hence,

we compute the power consumption of a PSU having an *efficiency* of e, in the following manner:

$$P_{PSU} = (\frac{P_{Mainboard} + P_{Fan}}{n * e}) * 100 - (\frac{P_{Mainboard} + P_{Fan}}{n}), \quad (24)$$

such that $P_{Mainboard}$ and P_{Fan} are introduced in Equations (17) and (22) respectively, whereas n denotes the number of PSUs and e (in percent) their efficiency (assuming that its identical for all the installed PSUs).

G. Total Server's Idle Power

Given a server composed of a mainboard, several fans and power supply units as illustrated in Figure 1, then its power consumption is given by:

1) For *Blade* type servers, the power consumption is:

$$P_{Blade} = P_{Mainboard}. \quad (25)$$

2) For *Tower* or *Rackable* type servers, the power consumption is given by the following equation:

$$P_{Tower\_Rackable} = P_{Mainboard} + \sum_{i=1}^{l} P_{Fan} + \sum_{j=1}^{m} P_{PSU}, \quad (26)$$

such that $P_{Mainboard}$, P_{Fan} and P_{PSU} are respectively given by Equations (17), (22), and (24).

V. Experimental analysis

In this section, we present the experimental results obtained by analysing different types (e.g. rackable, blade, etc.) of servers. For this purpose, we set up machines having various hardware characteristics (e.g. dual- / quad-/ hexa-core processors, DDR_2 / DDR_3 memory modules, etc.) as well as energy-saving mechanisms (e.g. Intel SpeedStep, AMD Cool'n'Quiet).

A. Setup Configurations

Table III gives the hardware characteristics of the analysed servers. Server 1 (rackable) is equipped with two dual-core AMD processors and power related information is monitored through a dedicated power meter. On the other hand, Servers 2 and 3 belong to HP ProLiant BL460c G6 series blade servers [17] such that the former is equipped with two quad-core Intel processors whereas the latter with two hexa-core Intel processors. In order to obtain monitoring information with respect to power consumption, HP integrated Lights-out (iLO) [18] is used for both servers.

It is worthwhile to note that this analysis was carried out within the context of [19] such that Servers 1, 2 and 3 belong respectively to the data centres of [20], [21] and [22]. Hence the results presented in Table VIII are obtained from power meters directly connected to Servers 1, 2 and 3 while keeping those machines in idle state for 15 minutes interval.

Table IV
Attributes of Processor

Attribute	Server 1	Server 2	Server 3
DVFS	TRUE	TRUE	TRUE
architecture	AMD	Intel	Intel
transistorNumber	243	840	1170
voltage	1.2	1.1	1.1
frequency	1.0	2.27	2.66
frequencyMin	0.8	2.26	2.66
frequencyMax	2.4	2.53	3.06

Table V
Attributes of Memory

Attribute	Server 1	Server 2	Server 3
voltage	1.8	1.54	1.54
size	1	4	4
type	DDR_2	DDR_3	DDR_3
frequency	667	1333	1333
bufferType	Buffered	Buffered	Buffered
vendor	KINGSTON	KINGSTON	KINGSTON

B. Configured Attributes

Tables IV, V, VI and VII provide respectively the configured attributes for the power consumption predictions of processor, memory, fan, mainboard, hard disk, and power supply unit presented in Section IV.

C. Obtained Results

Table VIII illustrates the idle power consumption breakdown of the corresponding servers. As mentioned above, the presented results are obtained by performing real experiments on Servers 1, 2 and 3 belonging to different testbeds within the context of [19].

Among the three servers, we notice that Server 3 has the least power consumption for processors in terms of the number of cores (e.g. 1.785 W/core). This is due to the fact that such a server is equipped with Intel hexa-core processors which possess advanced C-states that reduce the idle power consumption significantly

Table VI
Attributes of Fan

Attribute	Server 1
maxRPM	20000
actualRPM	9000
maxPower	13.2
depth	0.056

Table VII
Other Attributes of Servers

Attribute	Server 1	Server 2	Server 3
powerIdle (Mainboard)	45	70	70
powerIdle (Hard Disk)	7.1	6.0	5.96
efficiency (PSU)	83%	88%	88%

Table III
HARDWARE CHARACTERISTICS OF SERVERS

Component	Server 1	Server 2 [17]	Server 3 [17]
Processor	2× AMD Dual-core Opteron 2216	2× Intel Quad-core Xeon E5540	2× Intel Xeon Hexa-core X5670
Memory	8× 1 GB DDR_2 667 MHz	6× 4 GB DDR_3 1333 MHz	12× 4 GB DDR_3 1333 MHz
Hard Disk	Seagate ST3808110AS	2× 300GB 6G SAS 10K	2× 146GB 6G SAS 15K
PSU	ACBEL API3FS43 (efficiency 83%)	6× HP Common-Slot(efficiency 90%)	6× HP Common-Slot(efficiency 90%)
Fan	10× Sanyo Denki Ace 9CR0412S510	6× HP Active Cool 100	6× HP Active Cool 100

Table VIII
IDLE POWER CONSUMPTION BREAKDOWN OF SERVERS

Component	Server 1	Server 2	Server 3
Processors	11.8 W	35.5 W	21.42 W
Memories	7.6 W	9.8 W	19.72 W
Hard Disks	2.7 W	2.3 W	1.13 W
Mainboard	45 W	70 W	70 W
Fans	27 W	-	-
PSU	19 W	-	-
Total	**113 W**	**118 W**	**112 W**

by shutting down unused components of a processor. Furthermore, Servers 2 and 3 have the least power consumption for memory in terms of storage capacity (e.g. 0.44 W/GiB), since those servers are equipped with modern technology DDR_3 memory modules. Note that the idle power consumption of hard disks for the three servers are much smaller than the ones found in Table VII. The main reason for this is the configured parameters of Equation 16 where we set $\alpha = 0.1$ and $\beta = \gamma = 0.45$ (see Section IV-C for more details). It is worthwhile to note that since blade servers (e.g. Servers 2 and 3) share power supply units (e.g. PSU) and fans at enclosure level, hence those components were excluded from the total idle power consumption predictions.

VI. Conclusion

Lately, a great deal of research has been devoted to devising energy-aware optimisation algorithms that have the purpose of reducing the energy consumption of data centres. However, these optimisation algorithms should be guided by power consumption prediction models for the sake of taking the most appropriate energy-aware decisions. Traditionally, the models that estimate the power consumption of servers assume that the idle power is constant and known in advance. Consequently, the approaches proposed in the literature provide only the dynamic power consumption.

In this paper, we proposed models for idle power consumption estimations of processors, memories, hard disks, fans, and power supply units. To this end, we provided in-depth analysis of the idle power consumption of servers by breaking it down to its constituent components. Consequently, we believe that the provided analysis builds the foundation for in-depth scrutiny of idle servers' power consumption. As a proof of concept, the proposed models were implemented in [19] and it was shown that, with the help of energy-aware optimisation algorithms of [12], it is possible to save between 18-30% of energy consumption of data centres.

It is worthwhile to note that the proposed models are devised by taking into account the power characteristics of current technology. As technology evolves over the time, we expect that the power consumption behaviour of each component changes as well. For processors, due to absence of deep sleep states or dynamic voltage and frequency scaling, earlier processor generations show significant different power consumption values at low utilisation or in idle state. Additionally, the transistor size of integrated circuits is steadily decreasing, allowing for operation at lower voltages and hence power consumption. This does not only affect processors, but components like memory as well. The paper explicitly focuses on DDR_2 and DDR_3 memory modules. To model future memory standards like DDR_4, the derived equations may need to be revised. Regarding hard disk power, the most significant change can be observed through the advent of solid state disks, which eliminate hard disk's idle power nearly completely. Mainboard power consumption mostly depends on the number of integrated components; therefore, if the trend to integrate more and more components goes on, an increase can be expected. Regarding power supply units, in the past years major improvements have been reached with respect to their efficiency.

Acknowledgment

The research leading to these results was supported by the European Community's 7^{th} Framework Programme in the context of the FIT4Green and ALL4Green projects.

References

[1] R. Meade and R. Diffenderfer, *Foundations of electronics, circuits and devices*, 4th ed. Thomson Delmar Learning, 2003.

[2] http://en.wikipedia.org/wiki/MOSFET

[3] V. Pallipadi, *Enhanced Intel SpeedStep Technology and Demand-based Switching on Linux* http://software.intel.com/en-us/articles/enhanced-intel-speedstepr-technology-and-demand-based-switching-on-linux/, 2009.

[4] *AMD Cool'n'Quiet Technology* http://www.amd.com/us/products/technologies/cool-n-quiet/Pages/cool-n-quiet.aspx.

[5] *LMG500, Zimmer Electro System* http://www.zes.com/english/products/one-to-eight-channel-precision-power-analyzer-lmg500.html.

[6] http://www.kingston.com/us/memory/.

[7] H. J. van der Bijl, *Theory and Operating Characteristics of the Themionic Amplifier*, Proceedings of the IRE (Institute of Radio Engineers), pp. 97–126, 1919.

[8] US Environmental Protection Agency, *Report to Congress on Server and Data Center Energy Efficiency*, Response to Public Law 109–431, 2007.

[9] D. Kliazovich, Y. Audzevich, and S.U. Khan *GreenCloud: A Packet-level Simulator of Energy-aware Cloud Computing Data Centers*, GLOBECOM, pages 1–5, 2010.

[10] J. Liu, F. Zhao, X. Liu, and W. He *Challenges Towards Elastic Power Management in Internet Data Centers* , Workshop on Cyber-Physical Systems (WCPS), 2009.

[11] G. Chen, W. He, J. Liu, S. Nath, L. Rigas, L. Xiao, and F. Zhao *Energy-aware server provisioning and load dispatching for connection-intensive internet services*, the 5th USENIX Symposium on Networked Systems Design and Implementation, Berkeley, 2008.

[12] C. Dupont, G. Giuliani, F. Hermenier, T. Schulze, and A. Somov. *An Energy Aware Framework for Virtual Machine Placement in Cloud Federated Data Centres*, In Proceedings of the 3rd International Conference on Energy-Efficient Computing and Networking (e-Energy'12), ACM, Madrid, Spain, May 7–9, 2012.

[13] D. Economou, S. Rivoire, and C. Kozyrakis. *Full-system power analysis and modeling for server environments*, In Workshop on Modeling Benchmarking and Simulation (MOBS), 2006.

[14] X. Fan, W.-D. Weber, and L. A. Barroso *Power provisioning for a warehouse-sized computer*, In Proceedings of the 34th annual international symposium on Computer architecture, pages 13–23, 2007.

[15] T. Heath, B. Diniz, E. V. Carrera, J. Wagner Meira, and R. Bianchini *Energy conservation in heterogeneous server clusters*, In Proceedings of the tenth ACM SIGPLAN symposium on Principles and practice of parallel programming, 2005.

[16] S. Rivoire, P. Ranganathan, and C. Kozyrakis *A comparison of high-level full-system power models*, In Proceedings of the 2008 conference on Power aware computing and systems (HotPower). USENIX Association, 2008.

[17] HP ProLiant BL460c Generation 6 (G6) Server Blade Specification. http://h18004.www1.hp.com/products/quickspecs/13238_na/13238_na.pdf

[18] HP iLO Management Engine. http://h18013.www1.hp.com/products/servers/management/remotemgmt.html

[19] R. Basmadjian , C. Bunse , V. Georgiadou, G. Giuliani, S. Klingert and M. Majanen *FIT4Green - Energy aware ICT Optimization Policies*, In Proceedings of the COST Action IC0804 on Energy Efficiency in Large Scale Distributed Systems - 1st Year, 2010.

[20] Forschungszentrum Jülich (FZJ) http://www.fz-juelich.de/portal/DE/Home/home_node.html

[21] HP Italy Innovation Centre - Milano http://www.hp.com/it

[22] Ente Nazionale Idrocarburi (ENI) http://www.eni.com/en_IT/home.html

[23] http://www.engineeringtoolbox.com/fans-efficiency-power-consumption-d_197.html

[24] D. Brooks, V. Tiwari, and M. Martonosi *Wattch: a framework for architectural-level power analysis and optimizations*, In Proceedings of the 27th Int'l Symposium. on Computer Architecture, pages 83–94, 2000.

[25] W. Ye, N. Vijaykrishnan, M. Kandemir and M.J. Irwin *The design and use of simplePower: a cycle-accurate energy estimation tool*, In Proceedings of Design Automation Conference, pages 340–345, 2000.

[26] C. Huang, B. Zhang, A.C Deng and B. Swirski *The design and implementation of PowerMill*, In Proceedings of the International Symposium on Low Power Design, pages 105–110, 1995.

[27] A.P. Chandrakasan and R.W. Brodersen *Minimizing power consumption in digital CMOS circuits*, In Proceedings of the IEEE, pages 498–523, 1995.

[28] C.T. Hsieh, Q. Wu, C.S. Ding and M. Pedram *Statistical sampling and regression analysis for RT-Level power evaluation*, In Proceedings of International Conference on Computer-Aided Design, pages 583–588, 1996.

[29] C.H. Hsu, J.J. Chen and S.L. Tsao *Evaluation and modeling of power consumption of a heterogeneous dual-core processor*, In Proceedings of International Conference on Parallel and Distributed Systems, pages 1–8, 2007.

[30] R. Bertran, M. Gonzalez, X. Martorell, N. Navarro and E. Ayguade *Decomposable and responsive power models for multicore processors using performance counters*, In Proceedings of 24th International Conference on Supercomputing, pages 147–158, 2010.

[31] R. Basmadjian and H. de Meer *Evaluating and Modeling Power Consumption of Multi-core Processors*, In Proceedings of 3rd International Conference on Future Energy Systems (e-Energy), to be appeared, 2012.