A Cyber-Physical Approach to Combined HW-SW Monitoring for Improving Energy Efficiency in Data Centers

Josué Pagán , Marina Zapater , Oscar Cubo , Patricia Arroba , Vicente Martín , José M. Moya

Abstract—**High-Performance Computing, Cloud computing and next-generation applications such e-Health or Smart Cities have dramatically increased the computational demand of Data Centers. The huge energy consumption, increasing levels of CO_2 and the economic costs of these facilities represent a challenge for industry and researchers alike. Recent research trends propose the usage of holistic optimization techniques to jointly minimize Data Center computational and cooling costs from a multilevel perspective. This paper presents an analysis on the parameters needed to integrate the Data Center in a holistic optimization framework and leverages the usage of Cyber-Physical systems to gather workload, server and environmental data via software techniques and by deploying a non-intrusive Wireless Sensor Network (WSN). This solution tackles data sampling, retrieval and storage from a reconfigurable perspective, reducing the amount of data generated for optimization by a 68% without information loss, doubling the lifetime of the WSN nodes and allowing runtime energy minimization techniques in a real scenario.**

I. Introduction

The advent of Cloud computing and next-generation applications such as population monitoring, e-Health, Ambient Intelligence and Smart Cities is leading to a dramatic increase in the demand of computational capacity and data processing capabilities of computing facilities. Current Data Centers support traditional services such as Webmail, Web search, Databases, Social networking or distributed storage, as well as High Performance Computing (HPC) and Cloud Computing. Soon these infrastructures will also have to tackle the computational needs of the above mentioned next generation e-Science applications. These advances have already lead to a proliferation in the number of Data Centers. As next-generation applications continue to develop, Data Centers become more power-hungry, increasing their energy consumption at an unsustainable rate.

Electricity consumed by Data Centers today represents over a 1.3% of all electricity use in over the world, meaning 250 billion kWh consumption per year [1]. According to Barbagallo et. al. [2], these numbers imply that Data Centers produce tens of millions of metric tons of CO_2 annually, representing over 2% of total global emissions.

Energy efficiency is thus a major challenge that must be tackled to place Data Centers in a more scalable curve. For years, research has focused its efforts on minimizing either the cooling power of Data Centers, which usually represents around a 30% of the total infrastructure [3] of the power due to servers and computation (i.e. IT power).

However, recent research proposes that the energy problem has to be faced from a holistic perspective, focusing not only on the reduction of the cooling, but also on jointly minimizing IT and cooling power [4]. One of the most desirable features for this holistic energy optimization approach is to work on runtime and dynamically adapt to workload or environmental changes at the Data Center. Therefore, optimization frameworks exhibit one common requirement: the need to gather useful data at the appropriate rate from data room and servers. Modern enterprise servers are shipped with a wide range of sensors that prove huge amounts of data. Not all these data are useful, superfluous information has to be discarded. On the other hand, servers do not provide information on the data room, meaning that additional sensors have to be deployed to gather the appropriate environmental data.

This paper proposes the usage of a Cyber-Physical System (CPS) to gather, monitor and analyze both environmental and server information of a Data Center facility and makes the following contributions:

- Our work shows how environmental, server and workload information gathered via a CPS can be integrated in a holistic energy optimization framework, and describes the benefits of this approach.
- Data are analyzed and integrated to discover appropriate sampling rates, feeding the optimization framework only with relevant information. Redundant information sources and no relevant data are discarded.
- A low-power non-intrusive reconfigurable Wireless Sensor Network is deployed as a means to collect environmental data.
- Results show how runtime WSN reconfiguration allows the system to drastically reduce the amount of data collected and transmitted, by readjusting sampling rates depending on the workload parameters of the Data Center.

The remainder of this paper is organized as follows:

Section II describes the related work on this topic. Section III shows how this solution is integrated in a global optimization framework. Section IV and V analyze and detail the proposed solution respectively. Section VI shows an evaluation for a case study in a real scenario. Finally, the most important conclusions are drawn in Section VII.

II. Related work

Recent advances in the fields of technology and science have allowed the rapid development of new paradigms for population monitoring applications. Smart cities, e-Health solutions, Ambient Assisted Living, industrial applications and other e-Science applications produce a huge amount of data that has to be analyzed, processed and stored in order generate knowledge. To accomplish this, the usage of computing facilities is needed.

To ensure safe operation, most Data Centers set the Computer Room Air Conditioning (CRAC) air supply temperature to ensure the worst-case scenario of all servers. However, this solution is very inefficient in terms of energy consumption. From the data room perspective previous work addresses the power consumption problem by means of optimizing cooling costs at the resource manager level by assigning longer tasks to servers with lower inlet temperature [5]. In other cases Dynamic Voltage and Frequency Scaling (DVFS) techniques are applied by using heuristics, assuming ideal scenarios to achieve better thermal conditions but disregarding task assignment [6]. From the IT perspective, research has proposed solutions to reduce the computational power of servers by means of balancing workload across servers in the Data Center [7] via dynamic resource provisioning techniques [8]. Even though recent work performs runtime dynamic task allocation in heterogeneous Data Centers under large workloads [9], they do not implement their solutions in a real scenario with real applications.

Only very recently, the energy issue has started to be addressed from a joint cooling and IT perspective at the server level. Work by Xuefei et. al. [10] proposes a real-time fan controller based on a thermal model to reduce the energy consumption in the CPU. Research shows that energy efficiency can be improved by taking into account the leakage-temperature tradeoffs and reducing the overcooling of enterprise server [11].

At the Data Center level, industry has started to agree upon the importance of environmental room monitoring [12] to improve energy efficiency. Research by Gutpa et. al. [4] presents the Data Center as a distributed CPS in which both computational and physical parameters can be measured with the goal of minimizing energy consumption. As opposed to our work, they do not apply their solutions in a real scenario.

III. Optimization Framework Overview

This section shows how the proposed solution is integrated into the above mentioned optimization framework. The lower blocks of Figure 1 show the system architecture for the envisioned framework, which is deployed in a Cloud environment to guarantee redundancy and high availability. This framework takes as input all the information gathered from the application or service to be optimized. The system stores data and generates knowledge that can be used to run energy optimizations at different levels of abstraction (i.e. at the resource manager level, architecture, compiler, etc.). Optimizations use power models trained with real data obtained from the Data Center. Results of these optimizations are evaluated by a decision-making system, which finally proposes decisions to be executed at the Data Center.

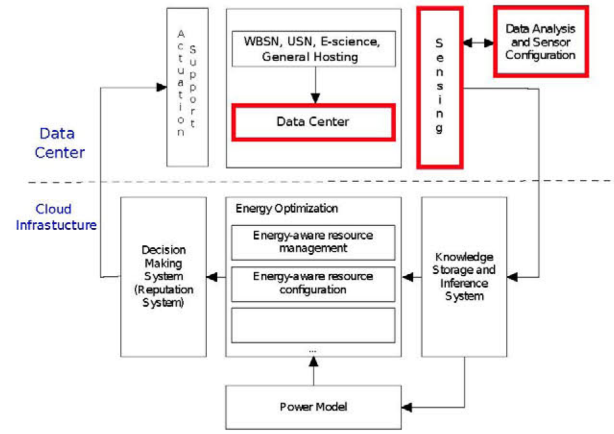

Fig. 1. Overview of the proposed energy analysis and optimization system in Data Centers

The upper blocks of Figure 1 show the application support network and the elements needed to gather useful data from the Data Center infrastructure. The sensing infrastructure represents the combination of two realities: a computational system and a physical network; this is commonly referred to as a Cyber-Physical System (CPS) and is the result of integrating the physical processes of the environment and software processes. In our case, the sensing infrastructure gathers computational parameters from the servers and the Resource Manager (RM), as well as the data room environmental information collected by means of additional sensors, deployed in the Data Center in a non-intrusive way.

Section IV analyzes the data needed to globally optimize cooling and IT at the Data Center, how it can be obtained and what is the optimum collection rate for each parameter.

IV. Data Analysis and Sensor Configuration

As mentioned before, one of the highest cost in a Data Center comes from the cooling needs of servers. In a raised-floor air-cooled Data Center, servers are mounted in racks on a raised floor. Racks are arranged in alternating cold/hot aisles, with the server inlets facing cold air and the outlets creating hot aisles. The CRAC units supply air at a certain temperature $T_{set,CRAC}$ and air flow rate f_{CRAC} to the Data Center through the floor plenum. The floor has some perforates tiles through which the blown air comes out. Cold air refrigerates servers and heated exhaust air is returned to the CRAC units via the ceiling, as shown in Figure 3.

The limiting factor for the data room cooling is the maximum temperature that can be achieved by the servers CPU. Temperatures higher than 85°C can cause permanent reliability failures (e.g. electromigration) that reduce the server Mean Time To Failure (MTTF) [13]. At temperatures above 95°C, servers usually turn off as a security measure. ASHRAE's Guidelines [14] recommend that inlet temperature of servers should not go above 27°C in order to ensure safe operation.

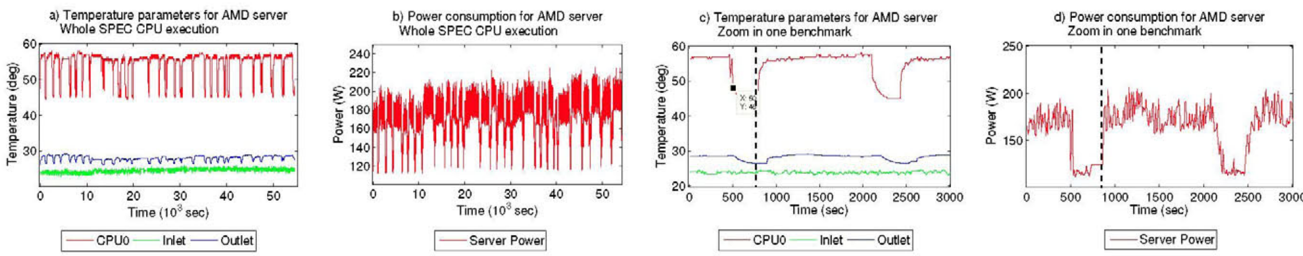

Fig. 2. Temperature and power values for AMD server under SPEC CPU 2006 Workload

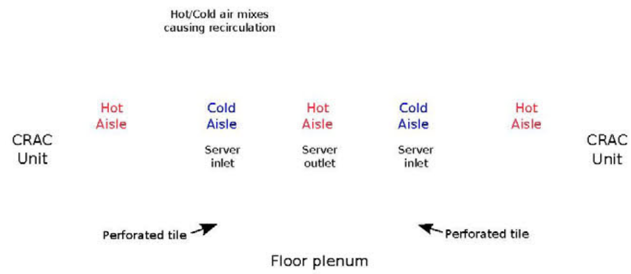

Fig. 3. Air-cooled Data Center layout

The temperature of the cooling air available at the inlet of IT equipment depends on the airflow dynamics between the perforated tiles and the equipment inlet. The maximum equipment power density that can be deployed in the Data Center is thus limited by the perforated tile airflow. Because the plenum is usually obstructed (e.g. blocked with cables in some areas), a non-uniform airflow distribution is generated and tiles exhibit different pressure drops ΔP, as described in Equation 1:

$$\Delta P = 1/2 f_{tile} \cdot \rho \cdot V^2 \tag{1}$$

where f_{tile} is the perforated tile loss coefficient, ρ is the air density and V is the air velocity.

To predict the inlet temperature of the servers, the pressure drop across tiles ΔP can be measured by means of differential pressure sensors. The inlet temperature, the interference of other servers and the power consumption of each server at a particular instant can be used to describe and predict the CPU temperature $T^{i,j}_{CPU,k+1}$ of a server at position (i,j) at the time instant $k+1$. This relationship can be described as in Equation 2:

$$T^{i,j}_{CPU,k+1} = f(T^{i,j}_{CPU,k}, P^{i,j}_{server,k}, T^{i,j}_{inlet,k}, T^{i,j}_{outlet,k}, \Delta P^{i,j}_{k}, FS^{i,j}_{k}, T_{CPU,k}\{i,j\}) \tag{2}$$

where $P^{i,j}_{server,k}$ represents the server power consumption, $T^{i,j}_{inlet,k}$ and $T^{i,j}_{outlet,k}$ the inlet and outlet server air temperature respectively, $\Delta P^{i,j}_{k}$ is the above mentioned tile pressure drop, $FS^{i,j}_{k}$ is the servers fan speed and $T_{CPU,k}\{i,j\}$ represents the interference of the temperature of adjacent servers, mainly due to outlet airflow recirculation.

Obtaining a temperature model for the servers can be used to predict their power consumption. The overall power consumption of an enterprise server can be described as in Equation 3:

$$P_{server} = P_{dynamic} + P_{static} + P_{fan} \tag{3}$$

$$P_{dynamic} = f(IPC, RW_{mem}) \tag{4}$$

$$P_{static} = I_{leak,m} \cdot V_{DD} = \quad \beta \cdot T^2_{CPU} \cdot e^{\frac{V_{GS}-V_{TH}}{nkT/q}} \tag{5}$$

Dynamic power is due only to the workload being executed in the server. Instructions per Cycle (IPC) is generally a good indicator of CPU power consumption [15], whereas the number of memory accesses (RW_{mem}), account for the memory power (see Equation 4). Dynamic consumption has historically dominated the power budget. But for technologies below the $100nm$ boundary, static consumption becomes much more significant, being around 30-50% [16] of the total power under nominal conditions. Leakage power is technology-dependant and primarily the result of unwanted subthreshold current in the transistor channel. Rewriting Equation 5 according to Rabaey [17] et. al. it can be seen that leakage is strongly dependant on CPU temperature.

Apart from obtaining all the above mentioned parameters, it is important to analyze the dynamics of the system to obtain the appropriate sampling rate. To do so, experiments are run on two different enterprise servers: (i) a quad-core Intel Xeon server with 16GB of RAM and (ii) a 2 CPU dual-core AMD Opteron server with 4GB of RAM. All the above mentioned parameters are monitored while executing tasks of the SPEC CPU 2006 benchmark.

Figure 2 shows the power and temperature traces for the AMD server while executing different benchmarks. Very similar dynamics can be observed in the Intel Xeon server. Figure 2*a* presents the CPU, inlet and outlet temperature while executing the whole benchmark suite, whereas 2*b* shows the overall server power consumption. Zooming in one particular execution (2*c* and 2*d*), we can get an insight into the dynamics of the system. The dashed black line is drawn on both plots at the same time instant, showing the beginning of a task. As can be seen, the server idle power oscillates around 120-130W. When a new workload begins power consumption raises above 160W, staying below 220W. At that same instant, as a result of the sudden power increase, CPU temperature raises around 6°C in just 10 seconds. However, it takes around 2 minutes for the temperature to reach the steady state. These two different dynamics have also been observed in previous work with different architectures [11]. Inlet temperature is steady at 24°C during all execution, as the air conditioning setpoint temperature is not changed. Outlet temperature, however, follows the trend of CPU temperature, but its reaction is slower.

As can be seen, power consumption has the fastest change

TABLE I. SUMMARY OF MAIN PARAMETERS TO BE MONITORED FOR OPTIMIZATION PURPOSES

Parameter	Dependency	Sampling rate
$T_{set,CRAC}$, f_{CRAC}	-	1day
ΔP_k	$T_{set,CRAC}$, f_{CRAC}	1 day
T_{inlet}	$T_{set,CRAC}$, ΔP_k	10 min
P_{server}	γ	1min
T_{CPU}	P_{server}, T_{inlet}	1-10 min
T_{outlet}	T_{inlet}, T_{CPU}	10 min
FS_k	P_{server}, T_{CPU}	10-30 min

frequency and defines the minimum sampling rate for the remaining parameters. CPU temperature exhibits fast changes when power gradients are large, and keeps steady otherwise. Outlet temperature follows this trend with some delay. The above results show the importance of considering different sampling rates depending on the magnitude to be measured. Moreover, we need to gather information about the allocation of each task in each particular time, as well as the begin time of tasks and their execution time in order to have adjustable and reconfigurable sampling rates. Equation 6 shows the main parameters that can be used to describe a job γ:

$$\gamma = f\left(t_{begin}, t_{reserved}, t_{end}, \alpha[p_k]\right) \tag{6}$$

where t_{begin} and t_{end} are the job begin and end time respectively, $t_{reserved}$ is the reservation time for that task and $\alpha[p_k]$ is a binary variable representing whether job γ is assigned or not to processor p in server k.

Table I summarizes the main parameters that have to be monitored to jointly optimize cooling and IT at the Data Center, as well as the events that generate changes on them and an estimation of their average sampling rate.

V. DATA CENTER SENSING ARCHITECTURE

In order to sense the entire Data Center, this paper proposes a multilevel star topology architecture built as shown in Figure 4, made of the following items: (i) a reconfigurable WSN based on a hardware deployment, to monitor the data room environmental conditions, (ii) server data information gathering tools, (iii) the Resource Manager (RM) that provides information about job allocation and (iv) the CRAC units.

All collected data is sent to a *gateway* that stores information and forwards it to the optimization platform. The *gateway* is also responsible for adding a common timestamp to the data or reconfiguring the data sampling rate. It also receives actuation information from the optimization platform and shows the proposed actions.

The next subsections further describe the different components of the sensing architecture.

A. Reconfigurable WSN

The WSN describes a star topology where peripheral nodes (*sensors*) send information to the central node (*gateway*). The WSN is composed of the following types of sensors:

- *Temperature nodes* measure the inlet and outlet air temperature of the servers at three heights in each rack. Nodes are placed on top of the racks, held with magnets and are equipped with six jack-plugged extensible thermistor-based temperature sensors. Three sensors are placed in the front of the rack and three at the back at different heights, also held with magnets, being non-intrusive to the normal operation of the Data Center.
- Humidity sensors are incorporated into temperature nodes; building up *Temperature-Humidity nodes*. There are fewer Temperature-Humidity nodes than Temperature nodes, in a relation of 2 to 1, as humidity gradients vary slowly in the room.
- *Pressure nodes* are built with a differential pressure sensor from which two probe-tubes come out. These nodes are placed under the raised floor, measuring the differential air pressure above and beneath the perforated tiles.
- *Power nodes* measure the current consumption drawn by the power supplies of servers. Each node can monitor up to 6 supplies per rack. Measures are made with current clamps based on hall effect. Current clamps are jack-plugged to the nodes, which are also placed on top of the racks.

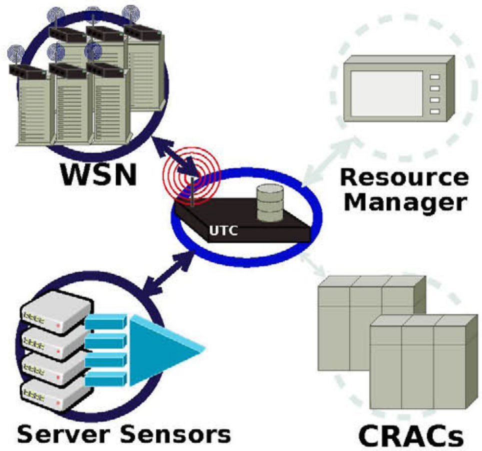

Fig. 4. Multilevel star topology for Data Center sensing architecture

All nodes are wireless, low-power and supplied via a $9V$ $450mAh$ battery. Their average lifetime is of 1.6 years for a worst-case scenario of a 1 minute sampling and 1 minute transmission rate. Nodes transmit the information collected by their sensors via radio at $868MHz$, by using the *Tulio*®module, a low-power radio device based on the CC1110 microcontroller and developed by AWD[1]. The *Tulio* offers an out-of-the-box radio interface solution with an indoor coverage of more than $20m$ for a transmission power of $-10dBm$, tested in a real Data Center environment. The *gateway* has a USB-pluggable Tulio receiver that collects the data sent by nodes and reconfigures the sampling rate of each node independently, increasing or decreasing the rate according to previous data analysis.

B. Server Sensors

Enterprise servers are shipped with a large number (from tens to hundreds) of internal sensors to monitor the server current status. These sensors can be polled via *IPMI*, *SNMP*, the server service processor or the proprietary tools from the manufacturers of each server. Large Data Centers usually

[1]http://www.awdynamics.com/productos/nodo_radio

deploy some kind of administration toolkit to retrieve, store and visualize data from all servers. Plenty of data are available but for the purpose of energy optimization perspective, the useful data are (shown in Section IV):

- CPU, memory and server ambient temperature.
- CPU frequency and voltage (i.e. DVFS mode), voltages of power supply units (PSU) and motherboard.
- Fan speed of server internal fans and PSUs.
- Power consumption of the overall server, chassis or rack.

Also, some parameters about the workload under execution can be gathered via the server hardware counters or OS utilities for profiling purposes. This is the case for: CPU and/or memory usage, Instructions per Cycle (IPC) or number of memory accesses. Knowing this together with the task allocation allows us to profile and predict the energy behaviour of tasks in the different servers.

C. Resource Manager

The Resource Manager (RM) is used to allocate in a spatio-temporal way the workload to be executed in the Data Center. The RM selects the resources that are going to execute a certain task, and has information on the reservation time for that task, and its real begin and end time. When a server completes the execution of a task before the end of the reservation time, the idle CPUs can be identified. Our system uses this information to decrease the sampling of sensors monitoring idle CPUs or servers, increasing it again when a new task is allocated. The RM is both part from the sensing and actuation network, as the optimization platform can perform actions that change the allocation decisions of the RM.

D. CRACs

CRACs are also both part of the sensing and actuation network. CRACs provide data about the current supply and exhaust temperature, humidity and airflow, and can also be used to change the cooling of the data room.

E. Gateway

The *gateway* is the central node in the multilevel star topology architecture designed (see Figure 4). It is a fanless 1U server based on an ARM1176 and managed by a light OS, that gathers all received data, stores it in an independent storage unit, analyzes and pre-processes the data and forwards it to the optimization platform at configurable intervals.

The main analysis and processing tasks carried out by the *gateway* before sending the information are:

- Establishing a common timestamp in Universal Time Coordinate (UTC) to match all data from the different sources, creating a multidimensional array that allows to correlate measurements. To accomplish this, all servers must have a synchronized clock and the WSN must have a known maximum delay which, in our case, is negligible.
- Sensors exhibit smaller or larger time constants depending on the nature of the measured magnitude, e.g. humidity and air pressure exhibit slower changes than CPU temperature, and CPU temperature varies slower than power. The *gateway* adjusts the sampling rate of these sensors, discarding superfluous data. Moreover, by runtime analyzing the gathered data, the *gateway* tracks changes in the workload and adapts the sampling rate, e.g. reducing the sampling rate of idle servers.
- Each sensor of the architecture allows independent and reconfigurable sampling rates. The *gateway* is responsible for reconfiguring each node at the appropriate rate depending on the workload conditions. This implies Over-The-Air reconfiguration of the WSN and changing the data collection rate of the RM, CRACs and servers.

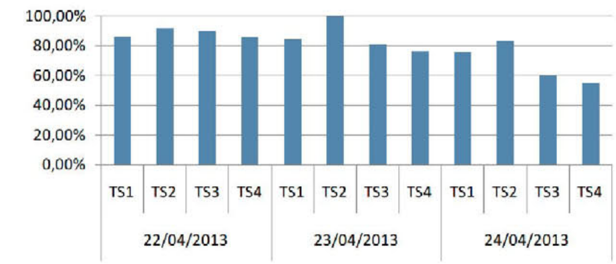

(a) Average load per Time-Slot (TS) and day for 3 days

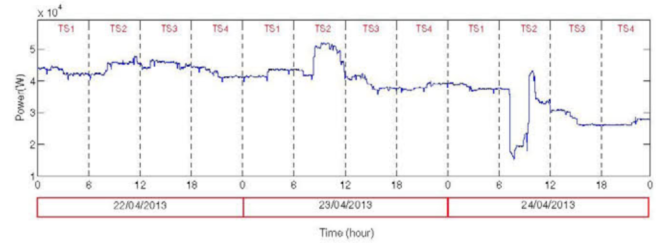

(b) IT power consumption 3 days

Fig. 5. Real load and power consumption traces for CeSViMa Data Center

VI. Results

This work develops techniques to allow energy optimization in real environments. To show the benefits of the Cyber-Physical approach, this paper presents a case study for a real High-Performance Computing Data Center at the Madrid Supercomputing and Visualization Center (CeSViMa)[2]. CeSViMa hosts the *Magerit* Supercomputer, a cluster consisting of 260 computer nodes in 9 racks, providing 4,160 processors and nearly 200 TB of storage. 245 of the 260 nodes are IBM PS702 2S with 16 Power7 processors running at 3.3GHz (26.4 GFlops) and 32GB of RAM. The remaining 15 nodes are HS22 with 8 Intel Xeon at 2.5GHz (10.2GFlops) processors and 96GB RAM. *Magerit* executes High-Performance jobs on demand with a maximum node reservation for high priority jobs of 1024 processors during 72 hours.

Figure 5 shows a real load and IT power trace of the *Magerit* supercomputer for three consecutive days, divided in 4 time slots of 6 hours each, starting at 12am. As can be seen, resource usage is not constant throughout the day, experiencing some drops and spikes. During those three days, the Data Center was working with 62% of its computer nodes active. The average load of the Data Center, considering only active nodes is kept around 80% most of the time.

[2]http://www.cesvima.upm.es/

TABLE II. SUMMARY OF MAIN PARAMETERS TO BE MONITORED FOR OPTIMIZATION PURPOSES

Cases	Avg. Sampling rate	Data gathered per day	Avg. WSN TX rate	WSN Battery Life
Baseline	1min	204.6 MB	1min	1.6 years
Naive	5min	129.6 MB	18min	3.3 years
CPS-based	10min	64.6 MB	5min	3.5 years

The environmental monitoring consists on the deployment of a WSN of 9 nodes on top of the racks: 5 *Humidity-Temperature nodes* and 4 *Temperature nodes*. In order to measure differential pressure, another 4 *Pressure nodes* are placed under the perforated tiles. To minimize the impact of the WSN deployment, the *gateway* is placed inside a rack in a non-intrusive way. *Magerit* uses *xCAT*[3] to retrieve information from servers and *SLURM*[4] Resource Manager to schedule and allocate incoming jobs. Also, gathered data is fed to the optimization platform.

Because workload is not stable over time, the proposed solution monitors the Data Center infrastructure at a reconfigurable sampling rate, reducing the amount of data collected when no relevant events take place in a server. On the other hand, variations in power consumption increase sampling rate to track associated events. Table II shows a comparison of data gathering per day and battery life for the case study under three conditions: (i) a baseline case in which all available sensors are sampled every minute (ensuring no relevant data losses) and sent to the optimization framework, (ii) a naive solution based on monitoring the relevant parameters for active servers at the appropriate data rate, but disregarding workload evolution and never reconfiguring sensors and (iii) our CPS-based reconfigurable monitoring solution with variable sampling rates.

Proposed solution is the one presenting the lowest power consumption for the WSN providing the lowest transmission time. When compared to both the baseline and the naive case, the amount of data gathered per day is drastically reduced, up to a 63% on average.

VII. CONCLUSIONS

Energy efficiency in Data Centers is nowadays a major challenge that has to be faced from a holistic perspective. This paper proposes a Cyber-Physical approach to integrate the Data Center into an energy optimization platform by performing runtime monitoring of environmental, server and workload parameters. To accomplish this, first the monitoring requirements are analyzed. Then, software techniques are used to gather workload and server data, and a Wireless Sensor Network is deployed to retrieve environmental data. A case study for a real Data Center scenario with variable workload is presented and we show how the monitoring solution can be reconfigured to adapt to workload changes. Results show how our solution improves monitoring, reducing the amount of data forwarded to the optimization framework by a 68%, speeding up optimization without information loss and maximizing the lifetime of WSN nodes.

[3]http://xcat.sourceforge.net/
[4]http://slurm.net

ACKNOWLEDGEMENT

Research by Josué Pagán has been funded by a research fellowship of the Spanish Ministry of Education for the collaboration of undergraduate students in research groups. Research by Marina Zapater has been partly supported by a PICATA predoctoral fellowship of the Moncloa Campus of International Excellence (UCM-UPM). This work has been partially supported by the Spanish Ministry of Economy and Competitiveness, under contracts TEC2012-33892 and IPT-2012-1041-430000, and INCOTEC. The authors thankfully acknowledge the computer resources, technical expertise and assistance provided by the Centro de Supercomputación y Visualización de Madrid (CeSViMa).

REFERENCES

[1] J. Koomey, "Growth in data center electricity use 2005 to 2010," Analytics Press, Oakland, CA, Tech. Rep., 2011.

[2] D. Barbagallo, E. Di Nitto, D. J. Dubois, and R. Mirandola, "A bio-inspired algorithm for energy optimization in a self-organizing data center," in *SOAR'09*. Springer-Verlag, 2010, pp. 127–151.

[3] N. Rasmussen, "Calculating total cooling requirements for data centers," American Power Conversion, White Paper 25, 2007.

[4] Z. Abbasi, M. Jonas, A. Banerjee, S. Gupta, and G. Varsamopoulos, "Evolutionary green computing solutions for distributed cyber physical systems," in *Evolutionary Based Solutions for Green Computing*. Springer Berlin Heidelberg, 2013, pp. 1–28.

[5] C. Bash and G. Forman, "Cool Job Allocation: Measuring the Power Savings of Placing Jobs at Cooling-Efficient Locations in the Data Center," USENIX'07.

[6] Q. Tang, S. K. S. Gupta, and G. Varsamopoulos, "Energy-efficient thermal-aware task scheduling for homogeneous high-performance computing data centers: A cyber-physical approach," *IEEE Trans. Parallel Distrib. Syst.*, vol. 19, pp. 1458–1472, November 2008.

[7] V. Cardellini, M. Colajanni, and P. S. Yu, "Dynamic Load Balancing on Web-Server Systems," *IEEE Internet Computing*, vol. 3, pp. 28–39, May 1999.

[8] E. Pakbaznia, M. Ghasemazar, and M. Pedram, "Temperature-aware dynamic resource provisioning in a power-optimized datacenter," ser. DATE '10, 2010, pp. 124–129.

[9] M. Zapater, J. L. Ayala, and J. M. Moya, "Leveraging heterogeneity for energy minimization in data centers," ser. CCGRID '12, Washington, DC, USA, 2012, pp. 752–757.

[10] X. Han and Y. Joshi, "Energy reduction in server cooling via real time thermal control," in *SEMI-THERM*, 2012, pp. 20 –27.

[11] M. Zapater, J. L. Ayala, J. M. Moya, K. Vaidyanathan, K. Gross, and A. K. Coskun, "Leakage and temperature aware server control for improving energy efficiency in data centers," ser. DATE '13, 2013.

[12] G. C. Bell, "Wireless Sensors Improve Data Center Energy Efficiency," DOE/Technical Case Study Bulletin. 1-877-EERE-INFO (1-877-337-3463). www.eere.energy.gov/informationcenter.

[13] D. Atienza, G. De Micheli, L. Benini, J. Ayala, P. Valle, M. DeBole, and V. Narayanan, "Reliability-aware design for nanometer-scale devices," in *ASP-DAC'08*, Jan. 2008, pp. 549–554.

[14] ASHRAE Technical Commitee (TC) 9.9, "2011 Thermal Guidelines for Data Processing Environments," American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., Tech. Rep., 2011.

[15] C. Hankendi and A. Coskun, "Adaptive energy-efficient resource sharing for multi-threaded workloads in virtualized systems," in *CHANGE'12*, 2012.

[16] S. Narendra and A. Chandrakasan, *Leakage in Nanometer CMOS Technologies*, ser. Integrated Circuits and Systems. Springer, 2010.

[17] J. Rabaey, *Low Power Design Essentials*, ser. Engineering (Springer-11647). Springer, 2009.