Intelligent thermal management in M2DC system

Wojciech Piatek
Poznan Supercomputing and Networking Center
Poznan, Poland
piatek@man.poznan.pl

Ariel Oleksiak
Poznan Supercomputing and Networking Center
Poznan, Poland
ariel@man.poznan.pl

Micha vor dem Berge
Christmann Informationstechnik & Medien
Ilsede, Germany
Micha.vordemBerge@christmann.info

Jens Hagemeyer
Bielefeld University
Bielefeld, Germany
jhagemey@cit-ec.uni-bielefeld.de

Emmanuel Senechal
ReFLEX CES
Evry Cedex, France
esenechal@reflexces.com

ABSTRACT

Thermal management and cooling are essential parts that have significant influence on the energy efficiency of data centers as cost of cooling can exceed 50% of the whole energy data center energy consumption. Optimized thermal management of data center also affects reliability and availability of a data center due to prevention of creation of the so called hot spots. In this paper, we present a model and optimisation method for thermal management a server platform, developed within M2DC project, equipped with a high number of heterogeneous hardware. We also show how the management of the individual servers and chassis influences efficiency of the whole data center. First, we present how this affects the commonly used PUE metric and how this approach can be misleading in evaluation of the data center effectiveness. Secondly, we show how intelligent fan management may influence energy used for cooling, change of IT systems energy consumption and the overall gain.

KEYWORDS

energy-efficiency; fans management; power and thermal models; data centers; power leakage; microservers

ACM Reference format:
Wojciech Piatek, Ariel Oleksiak, Micha vor dem Berge, Jens Hagemeyer, and Emmanuel Senechal . 2017. Intelligent thermal management in M2DC system. In *Proceedings of ACM e-Energy 2017, Hong Kong, May 2017 (E2DC)*, 7 pages.
DOI: 10.475/123_4

1 INTRODUCTION

Thermal management and cooling are critical elements of data center control. They affect significantly the energy efficiency of the whole center. For still common levels of PUE around 1.5-1.7 cost of cooling can exceed 50% of the whole energy data center energy consumption. Poor design of data center and use of inefficient cooling technology imposes limits on a total power usage, which may block development of a company. Recently, many attempts have been made to improve cooling and thermal management. Those include good practices in data center layouts such as hot and cold isles, closed cabinets, air ducts, etc. Various cooling technologies such as adiabatic cooling or direct liquid cooling have also been adopted to significantly decrease energy consumption and improve PUE values. Optimized thermal management of data center has influence not only on energy efficiency. It has also impact on reliability and availability of a data center. Inappropriate mixing hot air may cause a creation of so called hot spots and in consequence switching of servers or reduction of computing capacities.

It is important that several aspects affect thermal management of a data center: cooling technology, data center layout, and management of individual servers. In this paper, we present a model and optimisation method for thermal management of a server platform equipped with a high number of heterogeneous hardware. This platform is being developed within M2DC project [1] and as a result a number of appliances for various relevant classes of applications will be provided. As now the platform is under development (having first designs, chassis layout, and potential hardware modules) we present a model, optimisation method and results based on simulations. Verification of the approach and customisation of models and algorithms will be possible after the hardware is available (within 1 year). We also show how the management of the individual servers and chassis influences efficiency of the whole data center. First, we present how this affects the commonly used PUE metric and how this approach can be misleading in evaluation of the data center effectiveness. Secondly, we show how intelligent fan management may influence energy used for cooling, change of IT systems energy consumption and what is the overall gain. Finally, we discuss how this impact could be applied to manage changing environmental conditions, available cooling capacity, and load.

The remaining part of this paper is organized as follows. In Section 2 we give a brief overview of the current state of the art. Section 3 presents the architecture of the analyzed system while Section 4 contains our fans management approach and methodology aiming at the optimization of the system power usage.

E2DC, Hong Kong
123-4567-24-567/08/06. . . $15.00
DOI: 10.475/123_4

2 STATE OF THE ART

In recent years, studies concerning thermal -aware management of fans located within the servers have become an investigated research area. In [4], the authors presented a fan controller that utilizes thermal models to manipulate the operation of fans. The main motivation behind their solution is to avoid the situation where the operation of a multiplicity of fans is driven by a single hotspot. Thus, the controller considers all the fans of the server simultaneously. Taking into account the prediction of server temperatures, the presented solution adjusts the speed of particular fans proactively. The authors compare their results to the reactive fan controller, showing the reduction of fan energy consumption by up to 20%. Similarly, Kim et al [12] proposed the fan speed control scheme reducing the performance degradation and the power usage. They present Proportional-Integral-Derivative controllers, which are immune to non-ideal temperature effects. Their solution is based on the set of sensors located inside a server supported with the power and thermal models used to determine trends in their changes. Zapater et al. [3] studied the relationship between leakage and temperature of a server and provided the empirical model. To determine the model, the authors studied the behaviour of a server under varying power usage and different fan speeds. Based on it, they designed a controller that tunes the fan speed to minimize the energy consumption for a given workload. The proposed approach adjusts the fan speed at runtime taking into account different utilization values of the system. Another method [6] leverages both DVFS and fan management at the same time. Their idea is based on using the thermal resistance of a forced-convection heat-sink as a control variable. The proposed approach tracks the energy-optimal temperature as closely as possible with a given workload, making the best trade-off between cooling power and temperature-dependent leakage power. Huang [2] proposed thermal-aware power optimization techniques that can be applied both on data center and server level. At the server level, the authors trade off fan and circuit leakage power by dynamically adjusting the server's thermal setpoint. Their solution tracks the changes in the server power and temperature for different values of the fan rotation speed, adjusting its thermal threshold toward an optimal fan speed that minimizes the overall system power. In this manner, the method lets the system to heat up when this saves more fan power than it costs in terms of leakage power. Ayoub [8] noticed thermal dependencies between CPU and memory and non-linearity in cooling energy. Thus, he presented a holistic approach that integrates the energy, thermal and cooling management. The author designed a unified thermal and cooling model for CPU, memory and fan subsystems. Final solution consists of a controller, sensors and actuators (on CPU and memory) that are activated depending on the thermal and cooling state of the system.

This research aims to provide a holistic approach to the aforementioned issues. By combining detailed server monitoring together with the prediction of server temperatures (based on power and thermal models), it will enable proactive fan speeds adjustment. By mapping fan zones to particular components (by the means of the weights related to the pair node-fan), it prevents from hotspot driven management. It also aims at minimizing overall system power usage by finding trade-off between power leakage and cooling power of fans. Finally, it considers heterogeneous system and keeps the temperature of its components beyond the desired level.

3 SYSTEM DESCRIPTION

The following section contains the description of M2DC server, its estimated power usage and the cooling approach.

3.1 Server overview

The M2DC [1] microserver system combines a large number of microserver modules in a single chassis using dedicated baseboards. The design of the chassis enables hot-swapping and hot-plugging of microserver baseboards. Different versions of the chassis are available depending on the number of microserver modules. All versions are based on the same characteristics:

- Compatible to standard 19-inch racks, 482.6 mm wide
- Standard air flow direction from front to rear
- Retractable blade-style microserver baseboard, hot-swap
- Hot-swap power supplies
- The width of a microserver baseboard is 45.72 mm
- A high performance microserver baseboard hosts 3 high performance microservers (COM Express form factor)
- A low power microserver baseboard hosts 16 low power microservers (TX1 Jetson/Toradex Apalis form factor)

The following versions of the chassis will be available:

- Small chassis, 1 RU, 3 microserver baseboard, baseboards will be directly connected to backpanel (Figure 1)
- Mid-Range chassis, 3 RU, approximately 800 mm length, 9 microserver baseboards slots of 9 HP width (Figure 2)
- Scale-out chassis, 3 RU, approximately 1000 mm length, 15 microserver baseboards slots of 9 HP width, which can be plugged from front (9) and rear (6) with 4 fans (92 mm) at the middle of the chassis

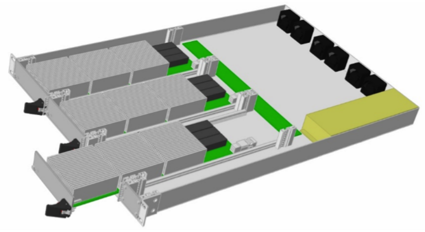

Figure 1: Small chassis with 3 microserver baseboard slots

The main advantage of the M2DC system, comparing to other server systems is it heterogeneity and high density. More details about that can be found in [9].

3.2 Power supply

The estimated total power consumption of the server system is detailed in Table 1. It shows the thermal design power (TDP) for the different parts of the M2DC server system.

According to these values, the power supply was chosen. Compact power supplies were selected to minimize the overall size of the M2DC Server.

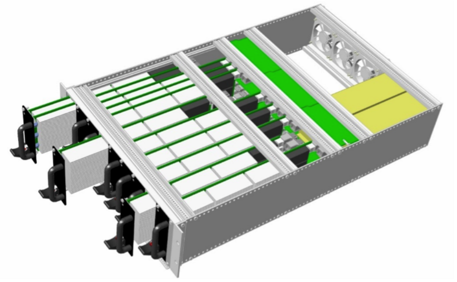

Figure 2: Standard chassis with 9 microserver baseboard slots

	TDP	Remark
Microserver Baseboard	280 W	16 LP microservers @ 15 W + 40 W infrastructure 3 HP microservers @ 80 W + 40 W infrastructure 250 W PCIe accelerator + 30 W infrastructure
Backplane	30 W / 50 W	For single/double sided chassis (Standard/Scale-out)
Backpanel	20 W	PHYs and Management
Fans	150 W / 200 W	92mm fans, EBM-PAPST 3212 JH4
PSU efficiency	0.9	4 redundant power supplies, 2000 W max. each
Standard chassis	3200 W	2870 W net. / 0.9 = 3189 W; 9 Baseboards
Scale-out chassis	5250 W	4720 W net. / 0.9 = 5245 W; 15 Baseboards

Table 1: Estimated power budget of the M2DC Server

3.3 Cooling concept

The cooling concept in M2DC box is based on conventional airflow cooling from front to back of the server. Special attention has to be paid to ensure sufficient cooling for all configurations and load conditions as the server is designed to provide a very high density of microservers resulting in high energy/heat density. Furthermore, since several microservers are placed in a row, the first servers will heat the air of the servers in the back. The following calculations indicate that it is possible to cool a full scale system under worst case conditions of an intake temperature of 45°C and an outlet temperature of 62°C as measured during the system proof of concept test:

$$V = \frac{Q}{c_p * \rho * \Delta t}, \rho = \frac{p}{R * (273 + T_{out})}, \Delta t = T_{out} - T_{in} \tag{1}$$

Justification of the variables in Equation 1 is presented in Table 2. Assuming a higher power consumption of 60 W per microserver, equipping a standard server chassis with 27 microservers would result in a total power consumption of 1620 W + 420 W (in the worst case for the infrastructure). With the formula above it is possible to calculate the air flow volume (V) needed to cool a given heat dissipation (Q). These values were taken for the calculation:

Description	Symbol	Value	Unit
Specific heat capacity of dry air	c_p	1.0053	kJ/(kg*K)
Air pressure	p	101325	Pa
Gas constant of dry air	R	287	J/kg*K
Intake temperature	t_{in}	45	°C
Outlet temperature	t_{out}	62	°C
Total heat dissipation	Q	Variable	W
Air flow volume	V	Target	m^3/s

Table 2: Values used for the evaluation

The result is a minimum air flow of 407.77 m^3/h. With three 92 mm fans that would be 135.92 m^3/h per fan. Based on the current assumptions, ebm-papst fans (model 3212 JH4) with a maximum air flow of 280 m^3/h and a maximum power consumption of 50 W have been selected for the M2DC Server. In systems with smaller power budget they can run at a lower speed to reduce energy consumption and noise. Additionally, this provides an additional level of safety if a fan fails or a microserver baseboard is being removed.

4 FANS MANAGEMENT CONCEPT

The aforementioned evaluation of the cooling requirements opens some possibilities in terms of fans management. According to the proposed system configuration, it would be possible to cool the nodes with some of the fans operating at half of their maximum speed or even with one of the fans disabled.

The main aim of our approach is to reduce the energy usage of the M2DC microserver system while keeping the temperature of the microserver below the desired level and to ensure the appropriate air flow level being able to remove all the heat. Presented methodology benefits from the possibility of managing each fan separately in a fine-grained manner (in general even with 1% speed accuracy). Moreover, due to a large number of sensors within the M2DC system, each fan in the box will have the outputs from multiple sensors mapped to it. This will allow providing cooling only where it is needed without wasting power in a situation where the operation of several fans is driven by a single hotspot. Apart from sensor readings, the proposed approach utilizes power and thermal models to determine trends in their changes. Taking into account the prediction of particular component temperatures, it will continuously adjust the speed of particular fans at runtime in a proactive way trying to maintain a particular set-point temperature and preventing the components from exceeding predefined thresholds. The algorithm will also consider the trade-offs between the power consumed by the fans and the change in power caused by the increase in the temperature (so-called power leakage). In this manner, the method will allow the system to heat up (without exceeding the given limit) when this leads to higher savings in fan power than it costs in terms of the power leakage.

In general, the speed of particular fans is derived based on the optimization criteria (as stated above - the energy usage), which is defined, initially, as:

$$min \int_{t}^{t+\Delta t} P(x) \mathrm{d}x \tag{2}$$

Figure 3: Concept of fans management

where power consumption of M2DC server is denoted as $P(x)$, t is the current time and Δt indicates the given time interval within which the energy usage will be optimized.

In our case the power consumption of the M2DC server can be simplified (with respect to Section 3.2) as:

$$P(t) = \sum_{i=1}^{n} P_{MS_i}(t) + \sum_{j=1}^{m} P_{fan_j}(t) \tag{3}$$

where P_{MS} is the power consumed by a given microserver, n indicates the number of microservers. Similarly, P_{fan} is the power consumed by a single fan, while m indicates their number.

Then, the power consumption of microservers and fans can be defined more particularly as ([10], [11]):

$$\begin{aligned} P_{MS}(t) = (&P_{MS_{idle}}(t) \\ &+ load * (P_{MS_{max}}(t) - P_{MS_{idle}}(t)) * g(T_{MS}(t)) \end{aligned} \tag{4}$$

where $P_{MS_{idle}}(t)$ is the power consumed by a microserver in an idle state, $P_{MS_{max}}(t)$ defines its maximum power drawn (while operating at full load), $load$ describes microserver's utlization and $g(T_{MS})$ is a function representing the power leakage of the microserver due to the increase in its temperature.

The power usage of the fan can be expressed using its dependency on the air flow:

$$P_{fan}(t) = k_f * V_{fan}(t)^3 \tag{5}$$

where V_{fan} defines the related airflow volume and k_f parameter that needs to be determined experimentally for specific hardware configuration (it covers fan efficiency and surrounding pressure drop).

In terms of thermal models we adopted the one presented in our recent studies [10]:

$$T_{MS}(t + \Delta t) = T^{\infty}_{MS} + (T_{MS}(t) - T^{\infty}_{MS})e^{-\frac{\Delta t}{RC}} \tag{6}$$

with $T^{\infty}_{MS} = P_{MS}R + T_{amb}$, $R = R_{cond} + R_{conv}$ and $R_{conv} = \frac{1}{k_n V^n}$

where $T_{MS}(t)$ defines a temperature at a given time t, Δt is a time step, T^{∞}_{MS} is a steady temperature for a microserver dissipating the given amount of heat, P_{MS} is the microserver power usage, T_{amb} is a temperature of ambient air, R defines thermal resistance and C is a thermal capacitance. Thermal resistance consists of conductive part R_{cond} and convective one R_{conv}, that can be defined by air flow volume V and k_v and n parameters that need to be determined experimentally as they are typical for a given equipment model.

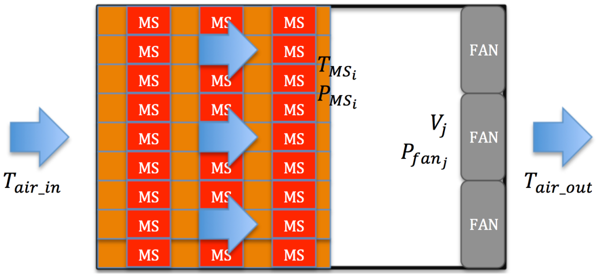

Figure 4: Thermal model of M2DC system

However, since, in general, the number of microservers does not correspond to the number of fans, the air flow affecting convective resistance is not equal to the one provided by the fan directly. That is because the air flow is shared by the neighbouring microservers and a fact that the air flows mixes with the ones provided by other fans as long as there are no dedicated ducts.

In order to define how much of the air blown by a single fan reaches the given microserver we introduce the weight function, defined as follows for each pair microserver-fan:

$$\forall i, 1 <= i <= n \forall j, 1 <= j <= m : w(MS_i, fan_j) \in [0, 1] \tag{7}$$

$$\sum_{i=1}^{n} \sum_{j=1}^{m} w(MS_i, fan_j) = m \tag{8}$$

Based on that we can recalculate the air flow value from Equation 6 as:

$$V = \sum_{j=1}^{m} V_{ij} \tag{9}$$

where V_{ij} is the air flow provided to microserver "i" as a result of operating state of fan "j" and with V_{ij} defined as:

$$V_{ij} = w(MS_i, fan_j) * V_{fan_j} \tag{10}$$

where V_{fan_j} is air flow volume related to fan "j".

This is also illustrated in Figure 5.

Hence, the overall optimization function is the following:

$$\begin{aligned} min \int_{t}^{t+\Delta t} (&\sum_{i=1}^{n} (P_{MS_{idle_i}}(t) \\ &+ load * (P_{MS_{max_i}}(t) - P_{MS_{idle_i}}(t)) * g(T_{MS_i}(t)) \\ &+ \sum_{j=1}^{m} k_{f_j} * V_{fan_j}(t)^3) \end{aligned} \tag{11}$$

with $T_{MS_i}(t)$ defined as in Equation 6 and corresponding value of $V = \sum_{j=1}^{m} w(MS_i, fan_j) * V_{fan_j}$

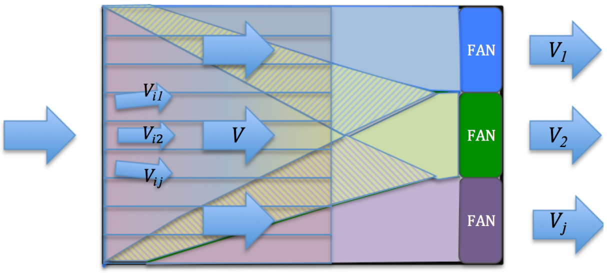

Figure 5: Air flow distribution model

Then, the speed of the fans is set by determining the values of corresponding fans air flows minimizing the function above (Equation 11).

To ensure system safety and reliability, the following constraints are applied (the second one is based on Equation 1):

$$\forall i, 1 <= i <= n : T_{MS_i} < T_{max_i} \tag{12}$$

$$\sum_{j=1}^{m} w(MS_i, fan_j) * V_{fan_j} >= \frac{Q_{MS_i}}{c_p * \rho * \Delta t},$$
$$\rho = \frac{p}{R * (273 + T_{out})}, \Delta t = T_{out} - T_{in} \tag{13}$$

They guarantee that temperature of each component of the system will not exceed the predefined threshold and that all the heat related to the given microserver will be removed.

In general the solution for the aforementioned problem has to be calculated using numerical methods. However, we can also easily solve the discretized version of the problem:

$$\sum_{x=t}^{t+\Delta t} (P_{MS_{idle_i}}(x) + load * (P_{MS_{max_i}}(x) - P_{MS_{idle_i}}(x)) * g(T_{MS_i}(x)) + \sum_{j=1}^{m} k_{f_j} * V_{fan_j}(x)^3) \tag{14}$$

Assuming that we consider 10 minutes time window within which we optimize the system energy consumption. Of course the aforementioned time frame does not affect the frequency of the fans management, which in general could be performed each second (our first simulation results shows that applying the proposed methodology each 20 seconds should be sufficient). Such a time window gives us insight into future system behaviour and allows analyzing the system power consumption until it reaches stability for the given state. We consider M2DC system equipped with low power servers (240 in total) where the speed of fans can be set with 5% speed accuracy (such level ensures appropriate granularity level and system reliability), we can easily determine the fan speed values even using the brute force approach. That gives us $1152 * 10^6$ power states/values that need to be calculated and 20^3 values (for each combination of fans operating speed) from which the minimal one has to be chosen. One should note that the branch and bound algorithm approach can be successfully applied to solve the aforementioned problem.

5 IMPACT ON DATA CENTER

This section discusses how the intelligent fan management (with respect to the presented approach) may affect the data center power consumption.

In general, power consumed by a data center consists of IT and cooling part. Power of the cooling equipment can be expressed as in [13] and [7]:

$$P_{cooling} = \frac{P_{IT}}{CoP(T_{sup})} \tag{15}$$

where P_{IT} is the total IT power and CoP is the coefficient of performance of the cooling facilities set to supply given T_{sup} temperature.

Based on Equation 1 we can obtain the relation between server's outlet temperature, heat dissipated by the server and the air flow used to cool it.

$$T_{out} = T_{in} + \frac{P_{server}}{V * c_p * \rho} \tag{16}$$

Following [13], the dependency between server's inlet temperature and the supply cold air can be expressed as follows:

$$T_{in} = T_{sup} + f(T_{out}) \tag{17}$$

where $f(T_{out})$ defined the impact of outlet temperature (heat) on the inlet one (due to the heat recirculation).

Table 3 shows assumed constant parameters and characteristics of the M2DC microserver system (as stated in Section 3):

Parameter	Value
c_p	$1,0053\ kJ/(kg * K)$
ρ	$1.19\ kg/m^3$
V	$840\ m^3/h$
P_{server}	$2040\ W$
PSU_{eff}	0.9
T_{in}	20°C
$f(T_{out})$	$0.1 * T_{out}$

Table 3: Characteristics of evaluated environment

Moreover, we adopted the performance function of cooling equipment as in [7] (it has been commonly applied to our previous and other researcher studies). We consider three fans management policies: normal mode (fans operate at their maximum speed), intelligent mode and min power mode (fans operate at their minimum, reliable speed). We perform our analysis for the data center consisting of 30 racks equipped with 10 M2DC standard chassis. Thus, the total IT load that needs to be removed is equal to 612 kW.

Table 4 contains the results of our calculations.

One should easily note that despite the increase in total IT equipment power consumption (due to the power leakage) and in the cooling facilities (due to required increase in supplied air temperature), applying intelligent fans management allows to reduce the overall power consumption by the substantial reduction in power

	Normal mode	Intelligent mode	Min power mode
T_{out}	22.03°C	24.06°C	26.76°C
T_{sup}	17.97°C	17.59°C	17.32°C
CoP	2.668238	2.576047	2.511736
P_{IT}	680 kW	689 kW	700 kW
$P_{cooling}$	255 kW	267 kW	278 kW
P_{fans}	66.6 kW	8.3 kW	1.8 kW
P_{DC}	1001.6 kW	964.3 kW	980 kW

Table 4: Data center power usage

used by fans inside a servers. On the other hand, applying much more rigorous eco mode does not improve the corresponding power metrics. That is because the reduction in fans power consumption can not compensate the power increase on it and cooling devices.

The aforementioned values are also used to calculate Power Usage Effectiveness (PUE) metric, which is one of the most commonly metric used to evaluate the efficiency of data centers and is defined as a ratio between energy used by the whole data center to the energy used by the IT equipment. However, the IT equipment includes fans, which are responsible for server cooling. This may lead to the situation in which adoption of powerful fans improves the PUE as it increases the IT equipment contribution in a total energy consumption. Thus, in our previous works [5], [11] we introduced another level of the PUE metric and refer to it as PUE Level 4. We proposed to perform the measurement of IT part only, excluding fans and PSU. Let us recall the definition of PUE Level 4 metric:

$$PUE4 = \frac{P_{total}}{P_{IT} - P_{IT_{fans}} - P_{PSU}} \tag{18}$$

Table 5 shows the differences in values for the PUE-based metrics.

	Normal mode	Intelligent mode
P_{DC}	1001.6 kW	964.3 kW
PUE	1.34	1.38
PUE 4	1.63	1.55

Table 5: Energy efficiency metrics

Figure 6 shows the comparison of PUE and PUE Level 4 metric in terms of two evaluated approaches for fans management. Bluish colors indicate the partial participation of IT power in the total data center power consumption with respect to both metrics.

With respect to Figure 6 and Table 5, one should note that, although the energy costs of operating fans do not affect PUE much, they have considerable impact on PUE Level 4. Furthermore, PUE metric does not reflect the increase in total energy consumed by a data center.

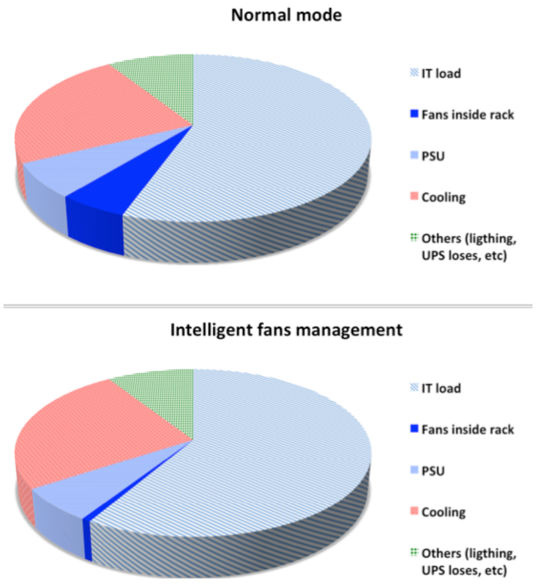

Figure 6: Comparison of PUE-based metrics

6 CONCLUSIONS AND FUTURE WORK

In this paper we presented a model and optimisation method for thermal management of the M2DC microserver platform equipped with a high number of heterogeneous hardware. According to the first simulation results, this approach may lead to savings of 7% energy consumption of the platform. We also presented how the management of the individual servers and chassis influences efficiency of the whole data center. First we show how this affects the commonly used PUE metric and how this approach can be misleading in evaluation of the data center effectiveness. We demonstrated that more efficient thermal management may lead to the worst PUE values and how to cope with this issue but the use of other metrics such as proposed PUE Level 4 or the overall energy to solution. We also presented the impact of the intelligent fan management on energy used for cooling, change of IT systems energy consumption and the overall gain. It turned out that, despite IT hardware power leakage and increased consumption of a cooling system it may bring 4% reduction of energy optimisation even for highly loaded systems. On the other hand, we showed that too aggressive fan control can end in no gain in overall energy consumption. This control can be applied to manage changing environmental conditions, available cooling capacity, and load. For example, if the there is problem with cooling capacity, e.g. due to high load and difficult environmental conditions the proper use of the central cooling system is of highest priority so the fans management should be reduced. On the other hand, if the heat is re-used in a data center and reduction of the overall energy consumption is a priority the intelligent fan management should be applied to bring additional benefits.

Currently, the M2DC microserver platform is under development (having first designs, chassis layout, and potential hardware modules). Verification of the approach and customisation of models and algorithms will be possible after the hardware is available (within 1 year) and is the main future work planned. The future work will also include research on integration of fan management with a control of computing nodes and allocation of workload providing additional energy reduction. Further study of the impact of thermal management algorithms on the whole data center efficiency and availability are also planned including feedback loop control to adjust management to the global state of data center.

7 ACKNOWLEDGMENTS

The results presented in this paper are partially funded from European Union's Horizon 2020 research and innovation programme under grant agreement No 688201 (M2DC).

REFERENCES

[1] 2017. Microserver Modular Data Centre (M2DC). (2017). http://m2dc.eu/

[2] M. Allen-Ware, J. Carter, E. Elnozahy, H. Hamann, T. Keller, C. Lefurgy, J. Li, K. Rajamani, J. Rubio, and W. Huang. 2011. Tapo: Thermal-aware power optimization techniques for servers and data centers. In *EEE International Green Computing Conference*. 1–8.

[3] J. L. Ayala, J. M. Moya, K. Vaidyanathan, K. C. Gross, A. K. Coskun, and M. Zapater. 2013. Leakage and temperature aware server control for improving energy efficiency in data centers. In *DATA*. 266–269.

[4] C. Bash, N. Tolia, M. Marwah, X. Zhu, P. Ranganathan, and Z. Wang. 2009. Optimal fan speed control for thermal management of servers. In *Proceedings of the ASME*.

[5] G. D. Costa, A. Oleksiak, W. Piatek, J. Salom, and L. Siso. 2014. Minimization of costs and energy consumption in a data center by a workload-based capacity management. In *3rd International Workshop on Energy-Ecient Data Centre*.

[6] J. Kim, N. Chang, J. Choi, S. W. Chung, E.-Y. Chung, and D. Shin. 2009. Energy-optimal dynamic thermal management for green computing. In *International Conference on Computer-Aided Design*. 652–657.

[7] J. Moore, J. Chase, P. Ranganathan, and R. Sharma. 2005. Making scheduling "cool": Temperature-aware resource assignment in data centers. In *Usenix Annual Technical Conference*.

[8] R. Nath, T. Rosing, and R.Z. Ayoub. 2012. JETC: joint energy thermal and cooling management for memory and CPU subsystems in servers. In *Proceedings of HPCA*. 299–310.

[9] Ariel Oleksiak, Michal Kierzynka, Wojciech Piatek, Micha Vor Dem Berge, Wolfgang Christmann, Stefan Krupop, Mario Porrmann, Jens Hagemeyer, René Griessl, Meysam Peykanu, Lennart Tigges, Sven Rosinger, Daniel Schlitt, Christian Pieper, Udo Janssen, Giovanni Agosta, Carlo Brandolese, William Fornaciari, Gerardo Pelosi, Mariano Cecowski, Robert Plestenjak, Justin Cinkelj, Loïc Cudennec, Thierry Goubier, Jean-Marc Philippe, Chris Adeniyi-Jones, Luca Ceva, and Holm Rauchfuss. 2017. M2DC – Modular Microserver DataCentre with Heterogeneous Hardware. In *Energyefficient Servers for Cloud and Edge Computing 2017 Workshop (ENeSCE 2017) co-located with HiPEAC 2017*.

[10] Wojciech Piatek, Ariel Oleksiak, and Georges Da Costa. 2015. Energy and Thermal Models for Simulation of Workload and Resource Management in Computing Systems. *Simulation Modelling Practice and Theory* (2015).

[11] Wojciech Piatek, Ariel Oleksiak, and Micha vor dem Berge. 2015. Modeling Impact of Power- and Thermal-Aware Fans Management on Data Center Energy Consumption. In *e-Energy 2015*. 253–258.

[12] M. M. Sabry, D. Atienza, K. Vaidyanathan, K. C. Gross, and J. Kim. 2014. Global fan speed control considering non-ideal temperature measurements in enterprise servers. In *DATE*. 1–6.

[13] Qinghui Tang, Sandeep Kumar S. Gupta, and Georgios Varsamopoulos. 2008. Energy-Efficient Thermal-Aware Task Scheduling for Homogeneous High-Performance Computing Data Centers: A Cyber-Physical Approach. *IEEE Transactions on Parallel and Distributed Systems* (2008).