

Combining Energy Saving Techniques in Data Centres using Model-Based Analysis

Björn F. Postema*
University of Twente
Enschede, the Netherlands
b.f.postema@utwente.nl

Tobias Van Damme*
University of Groningen
Groningen, the Netherlands
t.van.damme@rug.nl

Claudio De Persis
University of Groningen
Groningen, the Netherlands
c.de.persis@rug.nl

Pietro Tesi
University of Groningen
Groningen, the Netherlands
p.tesi@rug.nl

Boudewijn R. Haverkort
University of Twente
Enschede, the Netherlands
b.r.h.m.haverkort@utwente.nl

ABSTRACT

Advanced power management and cooling techniques for data centres often co-exist as separate entities in current-day operation of data centres. This paper proposes to combine these techniques to achieve greater power savings. To this end, an existing theoretical thermal-aware model is integrated in an extensive simulation framework for data centres using power and performance models, which allows for a detailed study in power, performance and thermal metrics. The paper compares four distinct cases for studying the effect on these metrics: a data centre with (i) basic functionality; (ii) advanced cooling; (iii) advanced power management; and (iv) a combination thereof. The combined case shows a significant reduction in the energy consumption compared to the other cases while performance *and* thermal demands are kept intact. The combination of these techniques shows improvements in energy savings and shows it is meaningful to investigate further into smart combined energy saving techniques.

CCS CONCEPTS

• **Computing methodologies** → *Simulation evaluation*; • **Hardware** → *Temperature simulation and estimation; Temperature control; Enterprise level and data centers power issues*;

KEYWORDS

data centres, power management, advanced cooling, model integration, thermodynamics, energy, performance

ACM Reference Format:

Björn F. Postema, Tobias Van Damme, Claudio De Persis, Pietro Tesi, and Boudewijn R. Haverkort. 2018. Combining Energy Saving Techniques in Data Centres using Model-Based Analysis. In *ICPE '18: ACM/SPEC International Conference*

*Both authors contributed equally to the paper

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICPE '18, April 9–13, 2018, Berlin, Germany

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-5629-9/18/04...\$15.00

<https://doi.org/10.1145/3185768.3186310>

on *Performance Engineering Companion*, April 9–13, 2018, Berlin, Germany. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3185768.3186310>

1 INTRODUCTION

From 2000 to 2006 the annual energy consumption of United States data centres increased from 28.5 TWh to 61.8 TWh, whereas in the years from 2006 to 2014, the annual energy consumption only increased to 69.8 TWh [8]. This small growth in energy consumption comes from efforts among data centre owners to push back the energy consumption of their data centres. According to [8], the three main energy-efficiency improvements that contribute to this flattening are (i) advanced cooling strategies, (ii) power proportionality, and (iii) server consolidation. The advanced cooling strategies focuses on techniques that increase the thermal efficiency of the data centre like hot aisle isolation, economizers and liquid cooling. Power proportionality is achieved with power management software and hardware, whereas server consolidation focuses on running current workload on as few servers as possible, in order to decrease the amount of hardware necessary in the data centre.

While these three areas separately show many improvements, we believe that more improvements can be gained by combining these areas, specifically combining the area of advanced cooling strategies with the area of power proportionality. In this paper, we investigate the cooperation between strategic power management control and strategic thermal control. Besides possible energy consumption benefits, this study allows us to show the general applicability of both these modelling approaches.

Recently, a simulation framework has been introduced to analyse models for both power and performance in data centres that use power management techniques to reduce its energy consumption [6, 7]. In this framework it is easy to study power and performance metrics of high-level models for any given data centre configuration and workload characteristic. Already these kind of analyses provide helpful insights in the design phase of data centres. Simultaneously a theoretical thermodynamical characterization of the cooling system in data centres has been performed [10]. This work proposes models for thermal-aware data centres that provide insights in the thermodynamics of the air flows in data centres. In that work, a control strategy was developed to dynamically steer the data centre to the optimal job distribution and cooling temperature to achieve the lowest energy consumption of the cooling system.

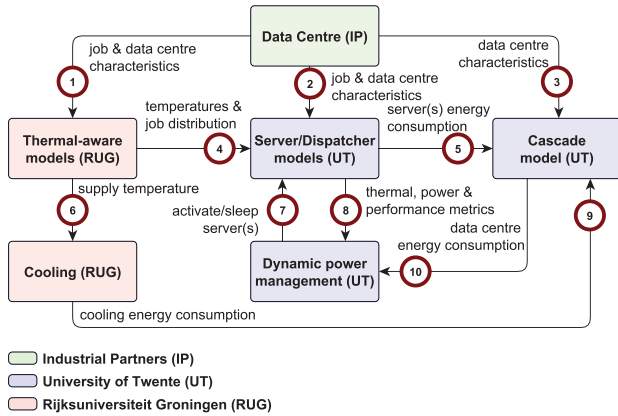


Figure 1: Detailed integration of thermal-, power- and performance-aware models for data centres.

In this paper, we focus on integrating the thermodynamical model in the existing simulation framework to study the interaction between the power management strategies from [7] and the thermal-aware controllers from [10]. Individually these areas have received much attention by researchers, e.g. [3] and references therein, however the combination of these two fields is much less studied [12]. This paper contributes to the existing state-of-the-art by providing an extensive simulation study that shows the viability of combining these two distinct control strategies and study the improvements that can be made by combining the two approaches.

The remainder of this paper is organised as follows. In Section 2 the models are introduced and their integration into the simulation framework is discussed. Next, the simulation configuration is given in Section 3 and different control scenarios are described in Section 4. Finally, the simulation results are studied in Section 5.

2 MODEL INTEGRATION

In this section we explain how the different models are integrated in the simulation framework as well as explaining the background of each of the models.

Overview. Figure 1 shows an overview of how the different models are connected and how they interact with each other. This overview shows that the models use characteristics from the data centre, and computed results from the other models as input. All models use job and data centre characteristics based on realistic data centre configurations, 1–3. The temperature of each unit and the optimal job distribution is communicated to the server models and to the job dispatcher, 4. If enabled, the dispatcher schedules jobs among the servers using this optimal distribution. Moreover, strategic power management related decisions are made based on the metrics available in the models, 7 and 8. The energy consumption of the computer room air conditioning (CRAC) is calculated using the thermodynamics and is communicated to the cascade model, 6 and 9; the energy consumption of the other infrastructural components remain linearly dependant on the energy consumed by IT equipment, 5. The total data centre energy consumption is sent to the power management module, 10.




Figure 2: Visualisation of the organisational hierarchy in data centres.

2.1 Data Centre Infrastructure

Data centres can be represented as a hierarchical system, see Figure 2. At the highest level, the *rack* level, data centres typically consist of aisles of racks. These racks typically consist of multiple subunits each containing a multitude of servers. This is the second level of the hierarchy, the *unit* level. On the lowest level, the *server* level, the servers consist of multiple computing *cores*.

2.2 Thermodynamical model

The aisles in a data centre are typically configured in cold and hot aisles. The cold air from the CRAC enters the data centre via the cold aisles and goes through the racks where it extracts some of the heat which is produced in the racks. The heated air then exits the racks into the hot aisles from where it is led to the CRAC to be cooled down to the desired cooling temperature. Since there are inefficiencies in the system, not all of the exhaust air is returned to the CRAC, and some of it is recirculated back to the cold aisles. There it mixes with the supplied air from the CRAC, heating the cool air slightly and reducing the effectivity of the CRAC. This inefficiency is the basis of the thermal models in [10]. This section briefly restates the models but leaves the details to the cited work. The models presented below model the temperature dynamics in the data centre and relate the temperature of the units to the energy required to cool the data centre.

Power consumption of units. The power consumption of the units in the racks is modelled with a linear function, i.e. it consists of a load-dependent and load-independent part. The load-independent part depends on the current global power state of the servers in the unit, which is further elaborated on in Section 2.3. The load-dependent part depends on the number of processors that are actively processing jobs. The power consumption of the units is

$$P(t) = V + WD(t), \quad (1)$$

where

$$\begin{aligned} P(t) &:= (P_1(t) \ P_2(t) \ \cdots \ P_n(t))^T, \\ V &:= (v_1 \ v_2 \ \cdots \ v_n)^T, \\ W &:= \text{diag}\{w_1, \ w_2, \ \cdots, \ w_n\}, \\ D(t) &:= (D_1(t) \ D_2(t) \ \cdots \ D_n(t))^T. \end{aligned}$$

Here $P_i(t)$ is the power consumption of unit i at time t , v_i [Watts] is the power consumption of unit i related to the current power state of its servers, w_i [Watts CPU⁻¹] is the power consumption per CPU in use in unit i , and $D_i(t)$ the number of CPU's that are actively processing work in unit i at time t .

Thermodynamics. The temperature change of each unit in the data centre can be expressed by a differential equation which captures the relation between current temperature, supply temperature and power consumption of the units

$$\frac{d}{dt}T_{\text{out}}(t) = A(T_{\text{out}}(t) - \mathbb{1}T_{\text{sup}}(t)) + M^{-1}P(t). \quad (2)$$

Here

$$T_{\text{out}}(t) := (T_{\text{out}}^1(t) \quad T_{\text{out}}^2(t) \quad \dots \quad T_{\text{out}}^n(t))^T,$$

and T_{out}^i [°C] is the temperature of the exhaust air at unit i , and T_{sup} [°C] is the temperature of the cool air supplied by the CRAC. Furthermore

$$A := \rho c_p M^{-1}(\Gamma^T - I_n)F,$$

$$F := \text{diag}\{f_1, f_2, \dots, f_n\},$$

$$M := \text{diag}\{c_p m_1, c_p m_2, \dots, c_p m_n\},$$

$$\Gamma := [\gamma_{ij}]_{n \times n}.$$

Here ρ [kg m⁻³] is the density of the air, c_p [J °C⁻¹ kg⁻¹] is the specific heat capacity of air, m_i [kg] is the mass of the air inside the unit, f_i [m³ s⁻¹] is the velocity of the air flow through unit i , and Γ is the matrix containing all the recirculation parameters.

Power consumption of CRAC. The power consumption of the CRAC is determined by the amount of energy which has to be extracted from the air. This in turn depends on the temperature of the air flows in the data centre, and the efficiency factor of the CRAC unit. The amount of energy to be extracted, $Q_{\text{rem}}(t)$, can be written in terms defined in the thermodynamical model. The efficiency of the CRAC is determined using the coefficient of performance, $\text{COP}(T_{\text{sup}}(t))$, first defined by Moore et al. [4]. The COP is a function of the target supply temperature and is defined such that a higher value denotes a more efficient CRAC unit. The heat removed and the corresponding CRAC energy consumption is given by

$$Q_{\text{rem}}(t) = -\mathbb{1}^T M A (T_{\text{out}}(t) - \mathbb{1}T_{\text{sup}}(t)), \quad (3)$$

$$P_{AC}(T_{\text{out}}(t), T_{\text{sup}}(t)) = \frac{Q_{\text{rem}}(t)}{\text{COP}(T_{\text{sup}}(t))}. \quad (4)$$

In [4] the COP for a water-chilled CRAC unit in the HP Utility Data Center is characterized as a quadratic, increasing function in the range of operation for T_{sup} .

2.3 Power and Performance Models

The models for power and performance are based on earlier work from [6]. Here we explain how each model is adapted to fit in the framework of this paper.

Performance. The performance models are extended with a two-level scheduling algorithm, see Figure 3. A central dispatcher distributes jobs to one of the n units using a scheduling algorithm of choice. Then, jobs are scheduled in round-robin fashion to servers 1 to N inside the unit. As in the original work, each server comprises a $G|G|1|\infty|\infty$ queue with a FIFO buffer.

Power consumption of IT equipment. The power consumption at time t for each of these servers is equal to the predefined reward $R(k)$ for each power state k as can be seen in Figure 4. Each state has a fixed power consumption with the exception of the processing state. As each server can have multiple computing cores, the power consumption of the processing state is also dependent on




Figure 3: Our extended dispatcher schedules jobs to the queues of the servers 1 to N via the hierarchy of the units 1 to n using a two-level scheduling algorithms in *Thermal-Aware (TA) and Round-Robin (RR)* fashion.

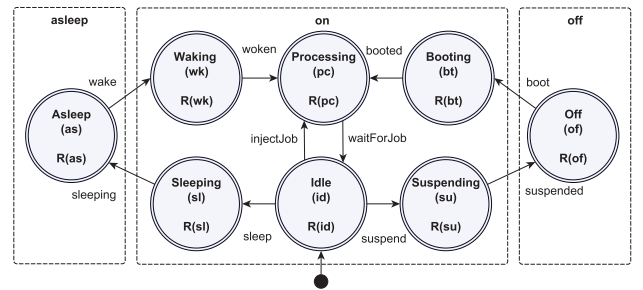


Figure 4: The power model for switching between three global power states *Asleep*, *On* and *Off* denoted by $R(k)$ where k denotes the current power state.

the number of active cores. The reward for the processing state is therefore given by $R[\text{pc}] = R[\text{id}] + w_i D_s(t)$, where D_s is the number of active cores in server s , and w_i is, as in Section 2.2, the power consumption per active core for the unit the server resides in. The power consumption of unit i is then given by the sum of the power consumption of the servers inside the unit.

The main power management feature is the ability to switch between global power states. This allows to adapt power consumption levels at the cost of time spent switching between global power states and therefore decreased performance.

Power consumption of data centre. Based on the IT equipment an estimation of the power consumed by other necessary infrastructural components can be computed using simple linear functions, which is called the *cascade model*. The total data centre power consumption is calculated by the sum of the power consumption of the CRAC, P_{AC} from (4), and the power consumption of the IT equipment and the other infrastructural components, as calculated by the cascade model.

2.4 Advanced Cooling Control

The thermodynamics described in Section 2.2 are used by the authors in [10] to characterise an optimal operating point to minimise the power consumption of the cooling system. This optimal operating point involves both an optimal CRAC supply temperature and an optimal workload distribution. Furthermore, controllers are designed that can dynamically steer the data centre to this optimal

operating point. The controllers are given by

$$\frac{d}{dt}T_{\text{sup}}(t) = \mathbb{1}^T A^T Z(T_{\text{out}}(t) - T_{\text{safe}}), \quad (5)$$

$$\frac{d}{dt}D(t) = \left(\frac{\mathbb{1}\mathbb{1}^T}{n} - I_n\right)(M^{-1}W)^T Z(T_{\text{out}}(t) - T_{\text{safe}}), \quad (6)$$

where T_{safe} could be the equipments' safe temperature level set by the manufacturer or a safe human working condition set by the data centre owner, and Z is a positive definite matrix.

The controller aims at steering the temperature distribution to the safe temperature level. The values of the inputs will be adjusted as long there are deviations from the safe temperature level.

When applying these controllers in the simulation, the values obtained from the controllers are used as set points. The CRAC takes the set point for T_{sup} and steers the supply temperature to this value, and the set point for D is sent to the job dispatcher. The controllers work only for a specified range of workloads. The exact range depends on the values of the parameters of the data centre, e.g. power consumption of servers, and recirculation flow. For our set of parameters the controllers work up to workloads of 50% of the total data centre computing capacity. Above controllers are therefore disabled for higher workload levels. The nature of this restriction is explained further in [10].

2.5 Advanced Power Management

An advanced power management strategy strategically puts servers in lower global power states, e.g. the low power consumption *asleep* state and *off* power state, to reduce overall energy consumption while performance is kept intact. The power consumption for each power state is depicted in Figure 4 with rewards R that are server configuration specific. Each r seconds a predefined power management strategy determines if servers should switch global power states. Following the specification in [7], a power management strategy, Θ , is defined by the global-level constraints, Φ_S , and the server-level constraints, $\Phi_C(s)$, for each of the available global power states in G . Our strategy Θ is defined as follows

$$\Theta = (\begin{array}{l} G = (as, on, of), \\ \Phi_S = \left(\begin{array}{l} \phi_S^{on} := (RT(\text{eavg}) > 0.75 \cdot R_{\text{SLA}}) \\ \phi_S^{as} := (RT(\text{eavg}) \leq 0.75 \cdot R_{\text{SLA}}) \\ \quad \wedge (PU(\text{id}, \text{ins}) \leq 0.3) \\ \phi_S^{of} := (PU(\text{id}, \text{ins}) \leq 0.3) \\ \quad \wedge (PU(\text{of}, \text{ins}) \leq 0.3) \\ \quad \wedge (PU(\text{bt}, \text{ins}) \leq 0.05) \end{array} \right), \\ \Phi_C(s) = \left(\begin{array}{l} \phi_C^{as}(s) := (QS(s) = 0) \\ \phi_C^{on}(s) := (PS(s) = \text{sl} \vee PS(s) = \text{of}) \\ \quad \wedge (\neg(PU(\text{as}, \text{ins}) \geq 0.0)) \\ \quad \wedge (PS(s) = \text{as}) \\ \phi_C^{of}(s) := (TO(s, \text{as}) \geq 100.0) \end{array} \right) \end{array}), \quad (7)$$

Our strategy requires the data centre to be able to observe the (exponentially moving average) response times and the (current) utilisation. The strategy compares exponentially moving average

response times to a threshold as stated in the Service-Level Agreement with a safety bound of 25%. Additionally, the strategy limits the number of idle servers to 30%; this ensures that enough servers are active to process the current workload and ensure sufficient capacity to be able to (de)activate servers. The strategy allows servers that are asleep for a duration of 100 s to be shut down under several utilisation conditions to ensure good performance.

2.6 General overview of the DACSIM simulator

In [6], a simulation framework has been proposed that allows for analysing the trade-offs between power consumption and performance in data centres. The aim of this framework is understanding ways to save energy via power management using the power and performance models from Section 2.3. A copy of the source code of the *Data Centre Simulation Framework* (also known as DACSIM) is publicly accessible via a GITHUB repository [5]. High-level simulation models allow us to estimate data centre power consumption and performance. The framework is developed in ANVLOGIC and allows for easy implementation of combinations of discrete-event and agent-based models. The framework features an intuitive dashboard that actively controls and obtains insights during each simulation run. Transient and steady-state behaviour can be analysed for (i) *power-state utilisation*, (ii) *response times* and (iii) *power consumption*. At the end of each simulation run, relevant data is exported for optional post-processing and more extensive analysis.

For the purpose of integrating the thermal-aware models in DACSIM, the matrix library EJML [1] is included to handle the differential equations. A module is set up that allows for (i) all the computations related to the thermal-aware models, (ii) transient analysis of the computed values during a simulation run and (iii) full logs of all the computed values.

3 MODEL PARAMETERS AND OUTPUT

3.1 Job and Data Centre Characteristics

The data centre in the simulation consists of 30 Dell PowerEdge 1855 server racks, i.e. units in Figure 2. Each unit has 10 dual-processor blade servers, i.e. a total of 20 CPU cores per unit. The base power consumption of a server in an idle state is $R[\text{id}] = 172.8$ W. The power consumption of each active CPU core is $w_i = 145.5$ W [9]. The power consumption of server s in the sleep or off power states is respectively $R[\text{as}] = 14$ W and $R[\text{of}] = 0$ W [2]. All other power states $R[\text{wk}]$, $R[\text{sl}]$, $R[\text{bt}]$ and $R[\text{su}]$ for global power state switching are rewarded as if all CPUs in the server are in use. The global power state switching time is distributed deterministically with mean $1/\alpha_{\text{wk}} = 1/\alpha_{\text{sl}} = 0.1$ (10 s) and $1/\alpha_{\text{bt}} = 1/\alpha_{\text{su}} = 0.01$ (100 s). The coefficients of the cascade model are taken from [6].

The data centre parameters were obtained from measurements by Vasic et al. [11] at the IBM Zurich Research Laboratory. The safe temperature threshold for the units is set at 30 °C. The initial temperature distribution of the units is set to 27.5 °C for all units.

Jobs arriving at the data centre are characterised by HTTP requests. The inter-arrival times and service times distributions in the model are calibrated with two data sets of HTTP requests from a real data centre, with each set having a duration of about 21 days (about 27.2 million entries), using a fitting algorithm in cooperation with Better.be. These distributions are exponential with a rate λ

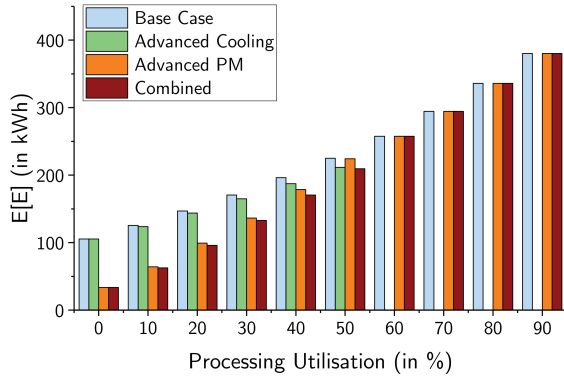


Figure 5: Total expected energy consumed by the data centre for the four scenarios with varying workloads from 0% to 90% of the total data centre capacity in increments of 10%.

Scenario	Scheduler	Cooling	PM Strategy
I: Base Case	RR-RR	Static	Always On
II: Advanced Cooling	TA-RR	Dynamic	Always On
III: Advanced PM	RR-RR	Static	Strategy Θ
IV: Combined	TA-RR	Dynamic	Strategy Θ

Table 1: Overview of the four scenarios

that is proportional to the desired workload in the case of the inter-arrival time, and a mixture of normal distributions with an average service time of about 107 ms in the case of the service times. The Service-Level Agreement (SLA) requires response times of HTTP requests to be below 1 s and an average response time of 300 ms.

3.2 Simulation Settings

Time units are set to seconds. The duration of the simulation was 3600 s. The warmup period for the system to adapt to the initial transient phase (e.g. sleeping the right number of servers) has been set to 1000 s. All simulations have been performed on a machine equipped with a 2.70 GHz INTEL[®] CORE[™] i7-4800MQ CPU, 8 GB of RAM and WINDOWS 7 64-bit with AnyLogic v8.1.0. The execution time of a single simulation run was between approximately 1 minute for the lowest workload and approximately 30 minutes for the highest workload. Results required a total of 40 simulation runs.

4 CASE STUDIES

This section focuses on four different control scenarios in a realistic data centre setting for the purpose of studying the impact of each control strategy on energy, performance and thermal measures. In Table 1 an overview of the different scenarios is given.

In the **base case** scenario (*Scenario I*), no advanced control mechanisms are applied, i.e., there is no feedback in control decisions. This scenario represents current day heuristics in many data centres. In this scenario, basic control of the data centre is applied at two levels, namely (i) *cooling*, (ii) *job scheduling*. The supply temperature of the cooled air of the CRAC is controlled in a way that keeps the temperature of the units below a certain safe threshold. If the maximum unit temperature is above the safe threshold, then the supply temperature will decrease, otherwise it will increase. Jobs

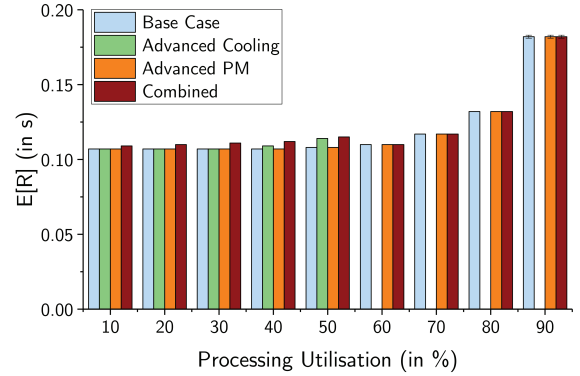


Figure 6: Mean response time for the four scenarios with varying workloads from 10% to 90% of the total data centre capacity in increments of 10%.

arriving at the data centre are scheduled in round-robin fashion to the units. The power management strategy is inactive, i.e., all servers are always turned on.

In the **advanced cooling strategy** scenario (*Scenario II*), only advanced cooling control is applied, there is no active power management, i.e. servers are always turned on. Controllers (5) and (6) are applied according to the steps described in Section 2.4. This control is tested up to and including a workload of 50% of the total data centre workload capacity.

In the **advanced power management strategy** scenario (*Scenario III*), cooling control and job distribution are the same as in the base case, whereas advanced power management strategy (7) is applied as specified in Section 2.5.

The **combined cooling and power management strategy** scenario (*Scenario IV*) allows for the investigation of a combination of both advanced power management strategies and advanced thermal-aware control. In this scenario, global power states are switched according to strategy (7) for energy-efficiency, and the job dispatcher follows the set point of the job distribution using controllers (5) and (6) for thermal-efficiency. Same as in Scenario II, advanced cooling control is applied up to workloads of 50% of the total data centre capacity. For workloads higher than 50%, this scenario operates according to Scenario III.

5 RESULTS

5.1 Energy

The total expected energy consumption $E[E]$ of the data centre for the full duration of the simulation is plotted in Figure 5 for all the scenarios, with different utilisation levels varying from 0% to 90% with increments of 10%. Note that with a processing utilisation of 100% the system would become unstable.

First, it is observed from Figure 5 that the higher the utilisation level becomes, the greater the energy reduction of the *advanced cooling* strategy becomes with respect to the *base case*. Secondly, a large energy reduction is observed at lower utilisation levels when only *advanced power management* is applied. However, the best energy savings, for all utilisation levels, are obtained when the two control approaches are combined as in Scenario IV. At higher

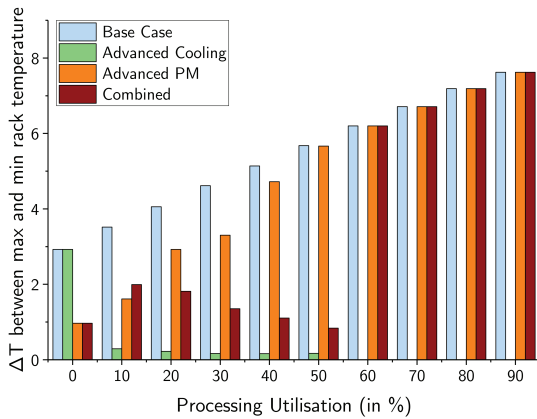


Figure 7: The temperature difference between the average maximum and minimum temperature for the four scenarios with varying workloads from 0% to 90% of the total data centre capacity in increments of 10%.

utilisation levels, our strategies have almost no room to control anything, and therefore no significant energy savings are observed.

5.2 Performance

For each of the simulation runs, the Service-Level-Agreement violations are recorded as a percentage of the overall number of jobs. The percentage of SLA violations for all processing utilisations has been 0% with an outlier of 0.011% at 90% processing utilisation due to the stochastic nature of the simulation.

Figure 6 shows the mean response times for the four scenarios with varying workload from 10% to 90%. The 0% case is skipped as there are no jobs arriving in the system in this case. The figure shows an increase of at most 100 ms in the average response times for all scenarios. It is seen that distributing jobs in a thermal-aware fashion gives rise to a slight increase in response times, with the biggest impact seen at 50% utilisation. The SLA requirements are still met however, because response times should be at most 1 s and the average response time should not exceed 300 ms. So, the overall performance is maintained while energy is being saved.

5.3 Thermodynamics

In order to plot the temperature data in an understandable way, the spread in temperatures among the units is studied. To do this, the difference between the average maximum and minimum unit temperature over the full simulation run is calculated for all simulation runs. This temperature difference is plotted in Figure 7.

Comparing the temperature differences of Scenario I with Scenario II, we see that the *advanced cooling* strategy results in a very balanced temperature profile among the units. This is the reason for the energy savings between the two scenarios, observed in Figure 5. When comparing the temperature differences between Scenario III and Scenario IV, we see again large improvements in favour of the *combined case*, where *advanced cooling* is applied. Same as before, this smaller spread results in less energy consumed.

Note that in the case of 0% workload, not much interesting can be done as there are no jobs available for redistribution. Also in the case of 10% workload it is seen that Scenario IV has an increased

spread compared to Scenario III. However when considering all units, less heat is generated overall, as can be deduced from Figure 5 from the lower energy consumption of Scenario IV in this case.

6 CONCLUSIONS

In order to analyse a potential power-, performance- and thermal-aware data centre, thermodynamical models have successfully been integrated in an existing extensive simulation framework with power and performance models. Moreover, advanced energy-aware control strategies are studied in a realistic simulation setting. Energy consumption, performance and thermodynamics are analysed in four scenarios where different control strategies are applied. From the simulation runs we see that combining thermal-aware control strategies with power- and performance-aware strategies yields the best energy savings without suffering any SLA violations. Furthermore it is seen that the thermal-aware controller successfully balances output temperatures of the units.

Future work includes studying the combined controllers for all workload levels and studying different ways of combining power- and performance-aware controllers with thermal-aware controllers. Also, current analysis can be extended by studying the transient phases as a consequence of fluctuating workload conditions.

ACKNOWLEDGMENTS

The work in this paper has been supported by the Dutch national NWO project Cooperative Networked Systems (CNS), as part of the program “Robust Design of Cyber-Physical Systems” (CPS), including the industrial partners Target Holding B.V. and Better.be.

REFERENCES

- [1] Peter Abeles. 2017. Efficient Java Matrix Library. <http://ejml.org>. (2017).
- [2] Anshul Gandhi, Sherwin Doroudi, Mor Harchol-Balter, and Alan Scheller-Wolf. 2013. Exact analysis of the M/M/k/setup class of Markov chains via recursive renewal reward. In *ACM SIGMETRICS Performance Evaluation Review*, Vol. 41. ACM, 153–166.
- [3] Abdul Hameed and et al. 2014. A survey and taxonomy on energy efficient resource allocation techniques for cloud computing systems. *Computing* (jun 2014), 1–24.
- [4] Justin Moore, Jeff Chase, Ranganathan Parthasarathy, and Sharma Ratnesh. 2005. Making scheduling ‘cool’ temperature-aware workload placement in data centers. In *USENIX Annual Technical Conference*. 61–74.
- [5] Björn F. Postema. 2013. DaCSim: A Data Centre Simulation Framework. <https://github.com/bjornpostema/DaCSim>. (2013).
- [6] Björn F. Postema and Boudewijn R. Haverkort. 2015. An AnyLogic Simulation Model for Power and Performance Analysis of Data Centres. In *Computer Performance Engineering (Lecture Notes in Computer Science)*, Vol. 9272. Springer International Publishing Switzerland, 258–272.
- [7] Björn F. Postema and Boudewijn R. Haverkort. 2017. Specification of Data Centre Power Management Strategies. In *Proc. of the 6th Int. Workshop on Energy-Efficient Data Centres (E2DC)*. Hong Kong.
- [8] Arman Shebabi and et al. 2016. *United States Data Center Energy Usage Report*. Technical Report. Lawrence Berkeley National Laboratory. <https://eta.lbl.gov/publications/united-states-data-center-energy>
- [9] Qinghui Tang, Sandeep K. S. Gupta, Daniel Stanzione, and Phil Cayton. 2006. Thermal-Aware Task Scheduling to Minimize Energy Usage of Blade Server Based Datacenters. In *2nd IEEE Int. Symp. on Dependable, Autonomic and Secure Computing*. IEEE, 195–202.
- [10] T. Van Damme, C De Persis, and P. Tesi. 2017. Optimized Thermal-Aware Job Scheduling and Control of Data Centres. *IFAC-PapersOnLine* 50, 1 (2017), 8244–8249. 20th IFAC World Congress.
- [11] Nedeljko Vasic, Thomas Scherer, and Wolfgang Schott. 2010. Thermal-Aware Workload Scheduling for Energy Efficient Data Centers. In *Proceedings of the 7th international conference on Autonomic computing*. 169–174.
- [12] Weiwen Zhang and et al. 2016. Towards Joint Optimization Over ICT and Cooling Systems in Data Centre: A Survey. *IEEE Communications Surveys & Tutorials* 18, 3 (2016), 1596–1616.