

Implementation and demonstration of a building simulation based testbed for assessment of data centre multi-domain control strategies

Vojtěch Zavřel¹, J. Ignacio Torrens¹, Jan L.M. Hensen¹

¹ Unit Building Physics and Services, Department of the Built Environment
Eindhoven University of Technology, Eindhoven, The Netherlands

Abstract

The traditional data centre (DC) infrastructure is being significantly extended by modern information technology (IT) trends on one side, and lasting calling for DC sustainability on the other. A holistic DC management will be necessary to coordinate different DC processes and to dock the DC environment into modern cities and district infrastructure. A development of such a complex management requires comprehensive testing possibilities. The testing is hardly possible on the real DC infrastructure due to the mission critical nature. Building energy modelling methods offer a suitable platform for the development of a safe and reliable testing environment. This paper deals with new application of Building Energy Simulation (BES) method and introduces a workflow for virtual closed-loop testing of enhanced multi-domain operation for data centres

Introduction

The traditional arrangement of data centres has turned from computer rooms within buildings to computer warehouses: separate buildings requiring all necessary facilities (Seaton 2015). The DC infrastructure is being significantly extended. The rise of “cloud computing” changes the conventional business scheme of hardware collocation, when the DC owner or operator rents the physical space in racks and guarantees “only” uninterrupted power supply and thermal conditioning of the DC interior space. In order to run cloud services, the DC owner or operator has to provide a complete IT infrastructure and rents computational capacity – virtual space via internet. The cloud computing requires widening of the traditional DC infrastructure about IT hardware and software. In addition, current standards aiming at carbon neutral data processing and sustainability are going to put more pressure on integration of renewables into the DC infrastructure.

The vision for future generation of DCs relies on higher integration of on-site renewable energy sources, meaningful harvesting of waste heat and integration into modern cities infrastructure. In order to successfully dock DC’s infrastructures to this modern city district concept, DC’s will have to act as one whole. The main three processes: data processing, thermal conditioning and powering need to be coordinated and optimized at system level. However, such high-level management must primarily ensure the reliability and availability of DC services.

In fact, the DC reliability and availability of data processing services is a key aspect to take into further consideration. It is important to understand that the downtime cost is not proportionally related to failure detection, diagnostics and repair, which are relatively minor expenses. Downtime mainly relates to business disruption, lost revenue and end-user productivity (Emerson 2011).

Therefore, the requirements for a trouble-free DC operation are very strict (ASHRAE TC 9.9 2009, chap.14). Uptime Institute, as a well-known advisory organization on improving the performance, efficiency and reliability of business critical infrastructure, defines four tier level certification standard for DC reliability and availability (Uptime Institute 2009). According to this standard, DC services must be available between 99.671% - 99.995% of the year with various levels of DC component (e.g. cooling unit) redundancy.

To conclude, the next generation DCs will require solutions going beyond the component level. Holistic operation, which coordinates all processes with respect to reliability of IT services, will be key for satisfying sustainable goals in the near future.

However, holistic operation can be poorly tested at real DC facilities and potential of this approach has not been really proved yet. The survey conducted by the Uptime institute in 2015 (Stansberry 2015) reveals that one of the main barriers is the need for extensive and comprehensive training, which is in conflict with DC reliability and availability requirements. Thus, the testing and commissioning of any enhanced control strategy are often slowed down if not even discarded due to safety reasons related with the mission critical nature of a DC environment. Alternatively, BES supported virtual testing offers a trustworthy and safe testing environment for new multi-domain operational strategies, accelerating their implementation in physical DCs.

Virtual closed-loop testing

Virtual testing is proposed as a complementary solution to current practice of separated stand-alone testing of individual algorithms. Stand-alone testing and training is not sufficient for enhanced operational platforms and leads to long and expensive commissioning processes. The typical constraints for commissioning of enhanced building operations, identified by (Salsbury & Ashish 2003) are (i) difficulties to test performance at extreme plant conditions and (ii) requirements of very broad

knowledge of multiple systems and controller types to carry out tests and interpret results. The virtual DC testbed can provide an online and safe testing environment for comprehensive experimentation. However, the multi-domain character of the tested algorithms requires co-operation of numerous experts from different fields. The virtual testbed needs to be integrated into a wider simulation-tool chain such as the one depicted in Figure 1




Figure 1 Schema of the wider simulation tool-chain

This DC virtual testing was developed and implemented in the project GENiC (Consortium Genic 2013) as proof of concept. A virtual DC testbed was set up based on the DC demo-site of the project described in the following section. After clarifying input and output specification from all partners, the virtual DC testbed was integrated into the internet communication framework of the project. Then, the virtual DC testbed was made available to various developers from industry and academy. Each algorithm is housed in geographically dispersed servers. Such a decentralized data exchange structure allows contribution from different partners, using different programming language and tools, while satisfying their privacy policy related with intellectual property. Virtual testing aims to reduce the commissioning time and experimentation at real facility. Such testing can minimize the risk of an infrastructure failure during platform training and preserve the DC reliability. Practically, a comprehensive functionality testing of control algorithms is done virtually and this activity is moved from the commissioning phase to the platform development phase. The tested control algorithm is connected to the virtual actuators and interactively receives the virtual feedback – processed signal from the virtual DC testbed. The closed-loop testing reveals the realistic performance of the control algorithms under dynamic conditions. The closed-loop testing considers a given DC design and local control of individual devices from several domains.

The basic schema of the typical closed-loop arrangement is depicted in Figure 2 where virtual DC testbed is shown instead of controlled DC plant.




Figure 2 Schema of the virtual closed-loop testing

Besides the risk reduction, the main advantages of virtual testing are (i) repeatability of testing with same boundary conditions, (ii) possibility to speed up the whole testing process using virtual time synchronization and (iii) testing of multiple control algorithms from different domains at the same time. The proposed testing workflow can be briefly summarized into the following steps:

- (0) virtual testbed development and validation
- (1) stand-alone (open-loop) testing of a new operation strategy
- (2) virtual (closed-loop) functionality testing of individual algorithms
- (3) virtual (closed-loop) functionality testing of overall platform (combination of partner algorithms)

Each of these four steps is described in detail hereunder.

(0) Virtual testbed development

As stated before, the intention behind the development of a virtual DC testbed is to realistically quantify the impact of multi-domain control algorithms and prepare the ground for their smooth application to a real DC infrastructure. The virtual DC testbed simulates the energy and thermal behaviour of a DC infrastructure such as IT equipment, the DC space, HVAC devices and Power Supply devices including renewable energy sources (RES). The virtual DC testbed mimics the feedback from the real controlled system in terms of energy, temperature and mass flow metering.

Physics based models (white-box models), and especially BES tools, fit well for the DC testbed application. Moreover sufficient support for co-simulation or integration of user models is available (Clarke 2001; Hensen & Lamberts 2011). The physics based models allow configuration based on technical specifications. Measured data, which are rarely available, are mainly used for fine calibration

Although there appears to be a potential for DC application, this potential is rarely used due to lack of awareness across the DC community. Typically, the DC energy modelling was focused on DC air-distribution using high-resolution models such as Computation Fluid Dynamics (Rambo & Joshi 2007). The requirement of understanding of the DC infrastructure behaviour at

system-level gives the opportunity for an interesting application of BES tools. This trend can be observed in recently published works. For example Phan and Lin (Phan & Lin 2014) investigated a multi-zone representation of a DC with cold and hot aisle arrangement in Energy Plus. Similarly, Salom, Oro et al. (Salom et al. 2015), (Oró et al. 2015) introduced dynamic modelling of data center whitespace and also a DC modelling review for support of renewables integration in DC using Transient System Simulation Tool (TRNSYS). In our case, the virtual data centre testbed is developed in TRNSYS Simulation Tool using mostly existing models of cooling and power supply devices including renewables. TRNSYS offers various embedded libraries with models of Building zones, HVAC, Power or controller components called types. More information about the TRNSYS tool can be found in (the Solar Energy Laboratory 2012), (TRANSSOLAR Energietechnik GmbH 2012). The virtual DC testbed development follows the configuration of the real DC infrastructure described below. The testbed includes also RES system, which is out of the scope of this study. The main challenges are (i) wide multi-domain scope, which requires, high conceptual modelling skills, (ii) integration of the virtual DC testbed into the larger simulation tool-chain enabling the online closed-loop testing.

In terms of implementation of the model, the BES tool lacks a thermal representation of DC space at required spatial resolution. Especially challenging is the representation of used air recirculation back to IT equipment related with local hot spots and supply air bypass related with reduction of cooling efficiency. In our model, these phenomena have been addressed by multi-nodal airflow network developed for the virtual DC testbed. The multi-nodal network method is presented in previous work (Zavřel et al. 2015; Zavřel et al. 2016)

(1) Stand-alone (open-loop) testing

The first step of the virtual testing workflow is done by the algorithm developers. The stand-alone (open-loop) testing is a current practice of control algorithm and remains as a starting point of the workflow. The open loop testing requires training data, which can be gathered from previously measured data. The virtual DC testbed can support the stand-alone testing by providing a set of training data according to the developer needs (e.g. data for “extreme” conditions, behaviour of the system for step response of control signal etc.).

In open-loop testing, the control algorithm is not connected to any actuators (real nor virtual). The performance of the individual control algorithm is assessed based on a quality of the control signal. This is mostly sufficient for simple controllers. However, enhanced control platform with multiple inputs and outputs from several domains can be hardly assessed based on individual control signal traces. Nevertheless, the stand-alone testing needs to be executed. In this phase, initial information regarding control signal variables such as expected variable range, data format, etc. are gathered for the next phases.

(2) Virtual testing of an individual algorithm

In this step, the tested algorithm is connected together with the virtual DC testbed to the virtual actuators. The virtual actuators receive the control signal and the numerical model of the virtual DC provides an estimate of processed signals of the controlled system. Seemingly, the algorithm performance is assessed based on the virtual processed signal. The assessment takes into consideration dynamic behaviour of the controlled system including local control of individual devices (usually given by the manufacturer). The closed-loop functionality testing reveals the algorithm performance in complex system-level perspective and indicates consequences to other domains. This assessment may detect potential limits or errors of the particular tested algorithm. Another interesting use-case of the virtual functionality testing, which is comparison analysis of “competitive algorithms”. “Competitive algorithms” are algorithms targeting the same objective in a single domain but using a different approach to reach the objective. The comparison analysis of their performance can support a decision-making process.

(3) Virtual testing of an overall platform

The complex, high-level control platform usually consists of several individual algorithms, which optimize individual domains. The last step of the workflow is the closed-loop testing of multiple “partner algorithms” against the virtual testbed. The “partner algorithms” are algorithms that co-operate together and actuate the multi-domain system. Each algorithm targets its own objective and optimizes its domain of interest (e.g. IT workload or thermal management).

The holistic control platform acts at several domains and actuate various devices. The evaluation based on processed signal (e.g. energy demand, temperature traces) is not sufficient at this level of complexity. In order to assess such a complex actuation, the development team has to agree about key performance indicators addressing the overall system aspects of interest. The multi-criteria assessment reveals possible conflicts between “partner algorithms” actuation and assesses the operation performance of the control platform as whole.

In addition, the interface of the virtual DC testbed follows the exact I/O specification of a real monitoring system of the demo DC. Therefore, the tested platform can be connected to the real monitoring system in plugin fashion after all testing procedures. The idea behind mimicking the demonstration site is to train the algorithms under similar conditions as they would face in real operation and thus accelerate the commissioning process. As a result of the extensive virtual testing, the commissioning process may focus solely on the platform installation and not on the algorithm functionality testing and debugging. The virtual testing is promising concept to accelerate the expensive commissioning process of the DC control platforms.

Data centre case study description for virtual testbed development

As stated previously, the virtual DC testbed is configured based on the demonstration site of the Genic project. The case study DC is a small-size university DC located in Cork (Ireland) (Figure 3). This DC has been chosen due to its testing and validation availability. Although the design of the DC might be considered as outdated, this demonstration site is sufficient as proof of concept.

The DC has approximately 40 m² of floor area. The DC space is arranged into hot and cold aisle zones. The nominal load of IT devices, housed in eight racks, is 30 kW.



Figure 3 Illustrative picture of the demo DC site in Cork (Ireland)

Under floor air distribution is applied in the DC space. The cooling unit supplies conditioned air to the space below the raised double floor. The air enters to the cold aisle zone through perforated floor tiles. From there, the IT equipment takes the conditioned air. The air passes through the electronics and removes the dissipated heat caused by computation processing. The warm air goes below the ceiling and then it is taken back in to the cooling unit. The layout of the demonstration DC is shown in Figure 4




Figure 4 Layout of the case study DC and placement of devices and space temperature sensors (red dots)

Every rack houses different number and types of servers except of the racks A2 and A4. The racks A2 and A4 are occupied by routes and switchers with minimal power

consumption and negligible heat dissipation. Other than A2 and A4 racks, all racks are equipped by temperature sensors at the inlet (cold aisle) and at the outlet (hot aisle) sides, at heights 0.3, 0.9, 1.5m.

The computer room air conditioning (CRAC) unit circulates the indoor air. The CRAC unit contains a direct expansion (DX) unit and a parallel water economizer that bypasses the DX unit in periods when the outside air temperature is lower than the inside air temperature. The dissipated heat from the DC space is removed by CRAC unit. Then the heat is transported via the water circuit that connects the CRAC unit and the roof mounted dry cooler. Finally, the dry cooler releases the heat to the outdoor environment. The nominal cooling capacity of the system is at 40kW. Figure 5 shows a schema of the cooling system.




Figure 5 Schema of the main cooling system of the case study DC

The current local control of the CRAC unit is an on-off control of cooling source (the DX unit or the economizer). The return air temperature set point is constant at 21°C. The airflow set point is also fixed at 10000 kg h⁻¹ to ensure high air changes in the room. The described control strategy is taken as baseline HVAC control for the simulation.

Validation of virtual DC testbed

The virtual DC testbed has been calibrated and validated based on the real measurements of the described demonstration site. Particularly, a detailed validation at the “rack/room” level focusing on DC space temperatures was performed. The DC space model could reproduce 1 month measurements of the rack inlet temperatures (18 nodes) and CRAC return temperature with the normalized root mean square error (NMRSE) at 6% - 11% (+/- 1.5°C). The input variables for the “rack/room” validation are power per 1rd of rack and CRAC supply temperature monitoring. Furthermore, a validation at the “building” level was carried out, which aims at DC energy breakdown into individual components. Full validation results of the testbed are presented in the framework of the project Genic (Consortium Genic 2016a).

The algorithm development team from industry and academic field has approved the validity of the virtual DC testbed and allow the testing of their algorithms in the virtual environment.

Simulation-based assessment of multi domain algorithms

This section introduces the experiment definition, the tested algorithm and the simulation results. The simulation-based assessment has been performed for all algorithms within the Genic platform. All tests have been done for a period of two weeks for characteristic winter and summer season periods.

This paper presents the results only from the testing of workload and thermal managements during the summer season. The main focus lays on demonstrating of the virtual DC testbed concept and its capabilities. The full set of results from these experiments are available in public deliverables of project Genic (Consortium Genic 2016).

Experiment definition

The simulation period has been set for two weeks with simulation time step of 5 minutes. The time step was selected based on agreed execution time of the platform. Since most of the algorithms can find an optimal solution in shorter time, the virtual testing has been synchronized with virtual time given by the virtual DC testbed for all components in the wider simulation tool-chain. Therefore, the time required for the virtual testing process can be reduced. In our case, the closed-loop simulation of two virtual weeks takes around 34 hours of computational time.

All tests have been done under identical boundary conditions given by weather data, IT workload traces and grid data. Specifically, the boundary condition has been defined by Typical Meteorological Year (TMY) weather file for Ireland (Kilkenny) (Marion & Ken 1995), the historical traces of requested IT tasks from the Wikipedia DCs (MediaWiki 2015) and the historical grid data provided by SEMO service (SEMO 2015). The Experiments presented in this paper are specified in Table 1.

Table 1 Simulated experiment specification

	Name	Period	Weather	IT tasks	WL mgmt. (migration limit)	TH mgmt. (temp. set point)
Exp.:0	Baseline	21 Jul - 4 Aug	Ireland, Kilkenny	Wiki media	0	const. 21°C
Exp.:1	optimized WL				10	const. 21°C
Exp.:2	optimized TH				0	Variable
Exp.:3	combination				10	Variable
Exp.:4	synchronized comb.				10	Variable

The results of the algorithm testing are always compared with the simulated baseline given by current control strategy in the DC case study, where no modern IT workload management is present. The cooling system works with constant DC room set points at 21°C.

Key performance indicators definition

This simulation based assessment carries out an evaluation of IT workload, thermal and combined management. Therefore, the key performance indicators have been selected with respect to IT productivity, DC infrastructure efficiency and thermal condition in DC.

IT productivity

The definition of IT productivity is given by Equation 1 as ratio of useful work to IT power demand. In our case, useful work is defined as number of CPU cycles for processing of given IT tasks.

$$\text{IT productivity} = \frac{\text{Useful work (CPU cycles)}}{\text{IT demand (kWh)}} \quad (1)$$

The issue of this performance indicator in practice is the non-uniform definition of useful work, which highly depends on character of an IT task (e.g. storage, web-service, high performance computation). The presented CPU cycles may not be always representative useful work for each IT task and sometimes additional computational parameters need to be taken into account to characterize the useful work (e.g. RAM or data storage access). In our case of the university DC, the CPU cycles is sufficient.

DC energy efficiency

One of the popular indicators of DC energy efficiency is Power Usage Efficiency (PUE) or its analogy known as Data Centre Infrastructure Efficiency (DCiE) introduced by green grid (Belady et al. 2008). The DC energy efficiency is defined by Equation 2 as ratio of total DC power demand to IT power demand.

$$\text{DCiE} = \frac{1}{\text{PUE}} = \frac{\text{IT demand (kWh)}}{\text{Total DC demand (kWh)}} \quad (2)$$

Temperature violation

The ASHRAE TC9.9 introduces the recommended operational conditions for standardized classes A1-A4 for compute and storage servers in tightly controlled environments. The recommended environmental range is defined by intake dry-bulb air temperature, humidity or maximum dew point or maximum temperature rate of change. The recommended environmental range can be summarized by dry-bulb temperature range at 18 to 27 °C and relative humidity range at 30% to 60% (ASHRAE TC 9.9 2011).

The violation of the recommended temperature range increases the probability of IT failure. The intake temperature can exceed the recommended temperature for short time to allowable operational range. Nevertheless, the allowable range is not statement of long-term reliability but definition of extremes of operational conditions.

The temperature violation is defined as deviation of intake server temperatures from lower or higher bound of recommended range. Temperature violation can be expressed by Equation 3

$$\begin{aligned} \text{Temp Violation}_{HI} &= t_{\text{server in},i} > 27 - 27 \\ \text{Temp Violation}_{LO} &= t_{\text{server in},i} < 18 - 18 \end{aligned} \quad (3)$$

Description of tested algorithms

It is important to understand, that we as testers and building energy simulation specialists, have often very limited understanding of the algorithm insides of tested algorithm due to lack of knowledge in a particular field or developer's privacy policy. However, no internal information is usually required. The I/O specification is sufficient for any testing. The results should be always analysed together with the algorithm developer. However, for the sake of results comprehensibility, the algorithm concept and received control signal are described below.

IT workload management

First, it is important to understand that the IT workload management can significantly influence the final IT power outcome. Even, all tests are done for identical workload demand (requested IT tasks), the IT power outcome varies based on the applied IT workload management strategy. The modern IT management using virtualization and consolidation of IT workload allows decoupling of IT hardware from software applications. The IT tasks can be numerically represented by virtual machines, which contains an estimate of the computational capacity. Thus, incoming IT tasks can be analysed or even predicted before their allocation to the hardware. In other words, the modern IT workload management allows virtual machine (IT task) migration and consolidation to individual hardware and reach the optimal workload allocation within the virtualized cluster. The virtualization provides operation flexibility and leads to higher utilization of the hardware (The Greaves Group 2007). The operation flexibility together with accurate workload prediction is beneficial also in terms of energy savings.

It is worth mentioning that, it is very common that the server efficiency and related idle energy is very heterogeneous over the DC space with respect of server type or year of release.

When the server efficiency mapping is given to a workload allocation algorithm, the algorithm may consolidate the IT workload to the more efficient servers first. Then non-utilized servers can be set to the standby mode. The idle energy of individual servers, which states between 20-30% of the maximum server power demand (Barroso & Hölzle 2009, chap.5), is saved. The accurate IT workload prediction is necessary to activate the servers in advance of a peak and to avoid any SLA violation.

The IT power demand is inherently related with heat dissipation. The IT workload can be also allocated based on thermal priorities given by thermal management. This second allocation mode can be beneficial for reducing the risk of hot spots and to even distribution of internal heat gain from IT equipment across the DC room. The virtual DC testbed receives a power demand of the IT equipment per 1rd of rack. The difference between baseline and two modes of the optimized IT management in terms of total IT power is shown in Figure 6.




Figure 6 Total IT power demand for baseline (marked black), optimized workload management (marked grey) and optimized workload management synchronized with thermal priorities (marked green), for 2 weeks July-Aug

Thermal management

The thermal management development has been mainly focused on the DC space conditioning based on high-resolution sensor network in DC room. The temperature sensors are located at each one third of a rack.

The first mode of the thermal management varies the return temperature set point of the CRAC unit based on the inlet IT temperatures in order to, firstly, reach the recommended temperature conditions and secondly minimize HVAC power consumption.

In the second mode, the high-resolution monitoring is used for generating of the thermal priorities for the workload management. This concept so-called thermal aware computation has been already introduced in literature (Tang et al. 2007), (Banerjee et al. 2011). However, it has been poorly tested in the closed loop due to its technical complexity.

The baseline and received temperature set point for the two modes of the thermal management are shown in Figure 7. The figure shows also simulated processed signal – CRAC return air temperature for all cases.




Figure 7 Temperature setpoints (control signal) and CRAC return temperature (representative processed signal) of baseline (black), optimized thermal management (blue), optimized thermal management with synchronization (green) for half day in July

As mentioned, the CRAC unit contains embedded local control, usually given by the manufacture. In this case, the compressor circuit has an embedded on-off local control. The replacement of the local control or access of

component level by supervisory platforms is not realistic in practice. The local control must be considered as a design aspect in the simulation based assessment.

Simulation results

The presented virtual testing workflow describes several phases of testing. During our work, we experienced all phases: Virtual DC development, providing training data for stand-alone (open-loop) testing, interconnection with tested algorithms and closed-loop testing.

Firstly, the testing was focused on individual algorithm testing. The control signal and the virtual processed signal feedback was evaluated several times. It was an iterative process of testing, evaluation and debugging or detail tuning, where several versions of the same algorithm were compared with each other to select the correct settings.

After finishing all individual tests, the group of “partner algorithms” was tested together.

As stated, the virtual processed signal (e.g. IT power or room temperature) is not sufficient for evaluation of enhanced control platform. The testing of multiple algorithms or overall platform require evaluation based on performance indicators, which will address the performance at wider scope of each domain.

The key performance indicators defined above have been evaluated for period of two weeks. In order to capture the KPI over the operational range, the results are presented as function of DC utilization. Thus, all plots show key performance indicator versus DC utilization. The primary information is a mean part-load efficiency curve. The mean part-load efficiency curve offers a general overlook of the algorithm performance over the whole operational range. It serves a quantitative comparison of the algorithms performance. The performance spread is also shown in the plot as secondary information for detail analysis. The size of each dot of the performance spread indicates occurrence of the efficiency during the time at particular level of DC utilization. This visualization offers global overlook of the operational performance for the algorithm developers.

Figure 8 shows the part-load IT productivity for baseline, optimized IT workload management, and synchronized IT workload management with thermal priorities.




Figure 8 IT productivity versus DC utilization, mean part-load efficiency and performance spread

The thermal management has minor influence on the IT productivity indicator and therefore only mean part-load efficiency curve is plotted.

Figure 9 shows part-load DC infrastructure efficiency for baseline, optimized IT workload management, Thermal management and combination of both the algorithms with and without coordination




Figure 9 DCiE versus DC utilization, mean part-load efficiency and performance spread

Similarly, Figure 10 shows the recommended temperature range violation of all experiments. The allowable temperature bounds for IT equipment classes according to ASHRAE TC9.9 are indicated in the figure.




Figure 10 Temperature violation versus DC utilization, mean part-load efficiency and performance spread

As a summary, Table 2 shows all KPIs for simulation period of two weeks. In addition, the improvement or diminishment of efficiency in comparison with baseline is indicated for each experiment

Table 2 KPIs summary for the simulation period

	Productivity	DCiE	PUE	Violation occurrence
	(CPU Cycles kWh ⁻¹) x10 ¹²	(-)	(-)	(hours)
Exp.: 0	33.93	0.453	2.21	288
Exp.: 1	▲ 58.52	▼ 0.403	▼ 2.48	▲ 182
Exp.: 2	▼ 33.92	▲ 0.464	▲ 2.15	▲ 117
Exp.: 3	▲ 59.25	▼ 0.412	▼ 2.43	▲ 146
Exp.: 4	▲ 53.23	▼ 0.426	▼ 2.35	▲ 137

Table 3 shows absolute values of energy demand over simulation period of two weeks for total DC. The energy demand can be broken down into individual divisions: IT equipment consumption, Cooling device consumption and other auxiliary power. The other auxiliary power demand is related with power supply losses during power delivery and lighting. Finally, Table 3 mentions the total energy savings potential for each experiment with respect of DC case study specification, baseline definition and simulation period.

Table 3 Electricity demand and total savings for the simulation period

	Electrocity demand for 2 weeks (kWh)				Total savings
	Total	IT	Cooling	Other	
Exp.: 0	13617.0	6173.0	5626.9	1817.1	-
Exp.: 1	9056.5	3796.6	4035.6	1224.3	-33.5%
Exp.: 2	13278.7	6172.9	5307.5	1804.6	-2.5%
Exp.: 3	8782.0	3771.0	3805.4	1205.5	-35.5%
Exp.: 4	9391.5	4101.7	4014.3	1275.5	-31.0%

Discussion of simulation results

The simulation results are further discussed in this section. Firstly, it should be kept in mind that these results are supposed to be always assessed in collaboration with the relevant algorithm developer or group of developers. Our role is to provide virtual DC testbed, where all tests can be freely executed and to guide through the testing procedure. The results presented in this paper were discussed with respect of individual developer's comments.

IT productivity

The virtual testing reveals that modern IT workload management has great potential for energy savings and increasing the total computational productivity of DC. The absolute value of the DC productivity shown in Figure 8 has limited usage, because it depends on the definition of useful work. In our case, there is a common definition of useful work and therefore IT productivity results can be compared to each other.

While, the baseline productivity descends considerably with respect to DC utilization, the modern workload management keeps the productivity relatively constant in range of 65 to 75 tera-CPU cycles per kWh over the whole utilization range. By increasing the IT productivity, the total DC energy demand is reduced about 33% on average, (38% energy savings at IT power consumption only). These significant savings are reached mainly by workload consolidation to the most efficient servers and saving the idle power of non-utilized servers.

Although it may conceptually seem like a very simple process, there are a number of technical difficulties for IT management to work in such an operation mode (e.g. accurate prediction of IT tasks). Nevertheless, the internal workload management architecture is not the aim of this paper. In this study, we are interested in algorithm comparison through KPIs and savings potential.

As can be seen, thermal management has minimal effects to IT productivity indicator. The main reason is that this indicator captures only IT equipment efficiency and it is

not designed to capture efficiency of cooling systems. The inlet temperature may influence the IT equipment productivity only via variable fan speed and related consumption of internal server fans. This effect has been studied in literature (Moss & Bean 2011) and its influence is negligible, if the temperature is kept below the higher bound of the recommended range as it is in our case.

Figure 8 shows the case, when thermal priorities for IT workload allocation are considered. IT allocation is restricted by considering additional constraints given by thermal management. These constrains lead to a reduction of IT productivity. A saving potential of 31% in total energy consumption is estimated, which is 2% less than in the allocation based on IT equipment efficiency. We can observe a conflicting situation. Thermal management is trying to advice preferable locations for data processing and related heat dissipation from a thermal perspective. However, these locations are occupied by servers with lower computational efficiency. As a result, the thermal preferences lead to a reduction of the IT productivity and higher heat dissipation in total. It is worth to note that the reduction of IT productivity is very case-dependent and the heterogeneous configuration of housed servers is a very important aspect in this case. In theory, if the DC hosts IT equipment with similar efficiency, this restriction would be negligible.

Data Centre Infrastructure Efficiency (DCiE)

Figure 9 shows DCiE versus DC utilization. Firstly, mean part-load efficiency is discussed. The baseline scenario reached DCiE at 0.45 (PUE 2.21) on average with utilization range of 55% to 70% of total IT capacity. When the IT management is applied, the total DC utilization is reduced in the range of 20% to 50% with the same amount of IT tasks. Although there is significant reduction of the total DC electricity demand, the DCiE drops to 0.4 (PUE 2.48). It shows that the part load efficiency of the cooling system is not proportional to the IT utilization.

It should be underlined that the single criteria assessment may lead to a very misleading conclusion. This may be an issue especially for PUE. This indicator is broadly adopted but often misused by industry. In the simulation based assessment, we can observe that PUE and analogically DCiE have their limitations of use and the indicators are not able to assess IT workload management. IT workload management requires some complementary indicator (e.g. IT productivity) to be able to correctly interpret the results. To avoid any misuse, DCiE and PUE indicators aim to mainly energy efficiency of cooling and power delivery systems by their definition. Therefore, mean part-load efficiency curves in Figure 9 for different IT workload management can be hardly compare with each other. The analysis is further focused on the influence of thermal management

Due to described limitations, the overall improvement of DCiE and PUE in comparison with baseline is reached only for stand-alone thermal management (Table 2). The other experiments consider the optimized workload management, which always leads to a misleading reduction of energy efficiency defined by DCiE and PUE.

The total energy savings related with the application of the tested thermal management are estimated around 2% to 2.5% regardless application of IT workload strategy. One of the reasons for relatively low energy savings is the limited actuation of the tested algorithm. The tested control algorithm actuates only temperature set points of the CRAC unit. The rest of components (e.g. fans, pumps) are not considered by the tested algorithm and the baseline control is used for these components. In order to reach higher savings, overall cooling system should be managed. In Figure 9, we can observe that the part load efficiency of DCiE is dominated by part load efficiency of compressor unit, which is related to the compressor unit design. For lower DC utilization, the DCiE decreases with a similar trendline for all experiments.

The performance spread is clearly clustered in two efficiency levels, which are related with on-off behaviour of local control of the compressor unit. In fact, design factors such as the on-off behaviour and oversizing of cooling leads to significant fluctuation of temperature and limits performance of the applied thermal algorithm in term of energy savings. The algorithm tries to balance these temperature fluctuations and to reach the recommended temperature range as first. (see in Figure 7). In this case, the satisfaction of temperature range has higher priority than energy savings.

Temperature violation

Figure 10 shows that the baseline control strategy provides the worst performance in terms of thermal conditioning. The return temperature set point at 21°C, on-off control and over-sized refrigerant unit leads to overcooling of the DC space. The baseline experiment results are below the lowest limit of the A1 class meant for most delicate hardware. This situation has also been observed in the real DC. It is common practice that DC facility managers select rather low setpoints, overcooling the DC to be on the safe side (Yogendra & Pramod 2012, p.202). The undercooling situation of the real DC has not been observed during regular operation.

The workload management reduces the IT demand and related heat dissipation. The compressor unit is triggered less often. The performance spread of temperature violation in Figure 10 is not as concentrated as in the baseline experiment. The lower utilization leads to minor improvement in terms of the mean temperature violation curve. In both cases, the recommended temperature range is violated most of the time. The DC space is overcooled, which may lead to condensation issues and inefficient operation of cooling unit.

The thermal management significantly improves the thermal conditions in the DC space and satisfied the recommended range given by ASHRAE for most of the time regardless of the application of IT workload management and level of utilization.

The synchronized co-operation of thermal and workload management performs slightly below of recommended range on average, which may be explained by a tendency to avoid any hot spots in DC.

In general, the performance spread reveals a relatively large fluctuation in the range of 8-10°C of temperature

over the simulation period. Such a large spread is related with mentioned design issues.

To summarize, the simulation-based assessment via virtual DC testbed provides complete information about the testing of several algorithms and their combination. The results demonstrate the importance of detail multi criterial evaluation, which helps to interpret the results.

Conclusion

This paper focuses on the simulation-based assessment of IT workload and thermal management and their combination for the bespoke DC. The simulation results and the discussion demonstrate the features of the proposed virtual testing concept.

In general, this paper presented a new use-case of BES for DC application. Building energy model has been used in wider simulation tool-chain as development and commissioning support of enhanced DC operation. In this case, the numerical model is not directly used in any optimization process but as a virtual testing facility. The numerical model can be understood as a virtual lab, which can be tailored for a given purpose and given case study. Moreover, the DC model outputs may deviate from the real measurements, the premise is that the numerical DC model is able to realistically represent the building dynamics in the whole operational range. Thus, the DC outputs can be considered as training monitoring signal with same uncertainties as real measurements. Contrary to typical training dataset, the virtual DC testbed reacts dynamically to the given actuation. The saving potential and algorithm performance can be relatively indicated but the results have to be always compared to the baseline simulated for the same boundary condition. The BES method provides a trustworthy simulation environment. The DC algorithm developers mainly appreciate the dynamic feedback of “virtual monitoring”; the capability to replicate tests under identical conditions, comparability of simulated experiments and acceleration of the testing procedure.

The key achievement is the integration of the BES DC model to the wider simulation tool-chain enabling virtual closed-loop testing. This approach is especially suitable for testing of multi-domain operational platforms, where the real testing is time-consuming or not feasible. The DC environment as mission critical environments is exactly the case. The proposed virtual testing workflow has been applied and its usability has been validated in collaboration with partners from industry and academia in the framework of the international project Genic. All the steps of the virtual testing workflow have been successfully executed. After the virtual testing, the platform developed by a wider consortium has been deployed at the DC case study.

Acknowledgement

We are thankful to our industry and academia partners: Atos for providing integration framework, IBM Research providing workload generator and prediction, United Technologies Research Centre providing thermal management, University College Cork (UCC) providing workload allocation algorithm, Cork Institute of

Technology (CIT) providing project coordination and integration framework, Acciona Infrastructure providing RES lab and power management. Especially, we would like to thank to Dirk Pesch, Thomas Scherer, Piotr Sobonski, Deepak Mehta, Diarmuid Grimes and Enric Pages Montanera



This project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement no. 608826
More information about project GENiC (no. 608826) available <http://projectgenic.eu/>

References

- ASHRAE TC 9.9, 2011. 2011 Thermal Guidelines for Data Processing Environments – Expanded Data Center Classes and Usage Guidance. *Data Processing*, pp.1–45.
- ASHRAE TC 9.9, 2009. *Design Considerations for Datacom Equipment Centers*, American Society of Heating, Refrigerating, and Air-Conditioning Engineers. Available at: <https://books.google.nl/books?id=n634PgAACAAJ>.
- Banerjee, A. et al., 2011. Integrating cooling awareness with thermal aware workload placement for HPC data centers. *Sustainable Computing: Informatics and Systems*, 1(2), pp.134–150.
- Barroso, L.A. & Hölzle, U., 2009. *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*,
- Belady, C. (Microsoft) et al., 2008. Green Grid Data Center Power Efficiency Metrics : Pue and Dcie Editors : *The Green Grid*, pp.1–9.
- Clarke, J., 2001. *Energy simulation in building design* 2nd ed., Glasgow: Butterworth Heinemann.
- Consortium Genic, 2016a. *Deliverable 5.5 Verification and validation of GENiC platform*, Available at: <http://projectgenic.eu/>.
- Consortium Genic, 2016b. *Deliverable 6.3 Energy optimization and control strategies demonstration results*, Available at: <http://projectgenic.eu/>.
- Consortium Genic, 2013. project GENiC. Available at: <http://projectgenic.eu/>.
- Emerson, 2011. Understanding the Cost of Data Center Downtime: An Analysis of the Financial Impact on Infrastructure Vulnerability.
- Hensen, J.L.M. & Lamberts, R., 2011. *Building performance simulation for design and operation* Abingdon., Spon Press.
- Marion, W. & Ken, U., 1995. *User's Manual for TMY2s*, MediaWiki, 2015. MediaWiki. Available at: <https://www.mediawiki.org/wiki/MediaWiki>.
- Moss, D. & Bean, J.H., 2011. *Energy Impact of Increased Server Inlet Temperature*,
- Oró, E. et al., 2015. Energy efficiency and renewable energy integration in data centres. Strategies and modelling
- view. *Renewable and Sustainable Energy Reviews*, 42, pp.429–445.
- Phan, L. & Lin, C.-X., 2014. A multi-zone building energy simulation of a data center model with hot and cold aisles. *Energy and Buildings*, 77, pp.364–376.
- Rambo, J. & Joshi, Y., 2007. Modeling of data center airflow and heat transfer: State of the art and future trends. *Distributed and Parallel Databases*, 21(2–3), pp.193–225.
- Salom, J., Oró, E. & Besòs, A.G.S.A. de, 2015. Dynamic Modelling of Data Centre Whitespaces, Validation with Collected Measurements. In *Building Simulation 2015, 14th International Conference of the International Building Performance Simulation Association*. Available at: <http://www.bs2015.in/proceedings.php>.
- Salsbury, T.I. & Ashish, S., 2003. Control System Commissioning for Enhanced Building Operations. In *Proceedings of the 3rd International Conference for Enhanced Building Operations, October 13-15, 1(414)*, pp.1–9.
- Seaton, I., 2015. Top Data Center Trends and Predictions to Watch for in 2016. Available at: <http://www.upsite.com/blog/top-data-center-trends-and-predictions-to-watch-for-in-2016/>.
- SEMO, 2015. Single Electricity Market Operator. <http://www.sem-o.com/>.
- Stansberry, M., 2015. *UPTIME INSTITUTE DATA CENTER INDUSTRY SURVEY 2015 Focus on budgets, planning, metrics, and life safety*,
- Tang, Q., Gupta, S.K.S. & Varsamopoulos, G., 2007. Thermal-aware task scheduling for data centers through minimizing heat recirculation. *Proceedings - IEEE International Conference on Cluster Computing, ICC*, pp.129–138.
- The Greaves Group, 2007. Virtualization in Education. *Virtualization View*, (October), pp.1–20. Available at: [http://www-07.ibm.com/solutions/in/education/download/Virtualization in Education.pdf](http://www-07.ibm.com/solutions/in/education/download/Virtualization%20in%20Education.pdf).
- the Solar Energy Laboratory, 2012. *Trnsys 17 - Volume 4 Mathematical Reference*, Wisconsin-Madison.
- TRANSSOLAR Energietechnik GmbH, 2012. *Trnsys 17 - Volume 5 Multizone Building modeling with Type56 and TRNBuild*, Stuttgart.
- Uptime Institute, 2009. *Data Center Site Infrastructure Tier Standard: Topology*,
- Yogendra, J. & Pramod, K., 2012. *Energy Efficient Thermal Management of Data Centres* J. Yogendra, ed., London: Springer.
- Zavřel, V. et al., 2015. Model Development for Simulation Based Global Optimisation of Energy Efficient Data Centres. In *Building Simulation 2015, 14th International Conference of the International Building Performance Simulation Association*. p. 8. Available at: <http://www.bs2015.in/proceedings.php>.
- Zavřel, V., Torrens, J.I. & Hensen, J.L.M., 2016. Simulation-based assessment of thermal aware computation of a bespoke data centre. In *12th REHVA world congress CLIMA 2016*. Allborg, pp. 22–25.