OPTIMIZING THE TOTAL ENERGY CONSUMPTION AND CO_2 EMISSIONS BY DISTRIBUTING COMPUTATIONAL WORKLOAD AMONG WORLDWIDE DISPERSED DATA CENTERS

Frank T. van Schie[1,2], Vojtěch Zavřel[1], J. Ignacio Torrens[1],
T.F.W. Hundertmark[2], Jan L.M. Hensen[1]

[1] Unit Building Physics and Services, Department of the Built Environment
Eindhoven University of Technology, Eindhoven, The Netherlands
[2] Unit Data Centers, Deerns Nederland B.V.

ABSTRACT

Major internet service providers have built and are currently building the world's largest data centres (DCs), which has already resulted in significant global energy consumption. Energy saving measures, from chip to building level, have been introduced gradually in recent decades. However, there is further potential for savings by assessing the performance of different DCs on a wider scale and evaluating information technology (IT) workload distribution strategies among these DCs. This paper proposes a methodology to optimize the electricity consumption and CO_2 emissions by distributing IT workload across multiple imaginary DCs. The DCs are modelled and controlled in a virtual test environment based on a building energy simulation (BES) tool (TRNSYS). A controller tool (Matlab) is used to support testing and tuning of the optimization algorithm. A case study, consisting of the distribution of IT workload across four different types of data centers in multiple locations with different climate conditions, is presented. The case study will illustrate the efficiency of the approach proposed in this paper.

INTRODUCTION

In 2013, the electricity demand of IT systems approached 10% of world electricity generation (Mills, 2013). The demand for IT workloads, e.g. storage, network, and computation, is increasing rapidly (Rao et al., 2012a). This continuous growth of IT workloads has resulted in larger, more complex and energy dense DCs to process the data requests of all customers (Oró et al., 2015). As these IT workloads become larger, DCs' electricity consumption shows corresponding increases. The DCs' electricity consumption and the energy sources used to generate this electricity greatly influence the carbon footprint of a DC (Oró et al., 2015). An even larger increase of the carbon footprint is expected in the future as the dependence on coal as a source of electricity rises (Mills, 2013). A decrease in the electricity consumption and carbon footprint could be obtained by equipping DCs with renewable energy sources (RES). However, energy generation from RES is unpredictable due to ever-changing weather conditions, while the IT requests of a DC must be processed at any time (Oró et al., 2015).

Electricity saving measures for DCs have been examined on all different levels of functional abstraction. For example, twenty-two DCs have been benchmarked for the 'best-practice' technologies to reduce the electricity costs (Greenberg et al., 2006). Another approach to reduce the electricity costs is to combine information from management systems in a DC to enhance the performance of various systems (Sharma et al., 2008; Mohsenian-Rad et al., 2010). Similar to these studies, most research in this area has focused on reducing the electricity costs per DC at one specific location. Little research has been performed to reduce the total electricity consumption at multiple locations (Rao et al., 2012). Moreover, there is no earlier research studying the optimization of multiple performance indicators at the same time, such as the combination of the total electricity consumption and CO_2 emissions.

The main objective of this research is to identify the potential reduction in electricity consumption and CO_2 emissions by distributing the IT computational workload among geographically dispersed DCs around the world. This principle will be referred to as 'Guiding the Cloud'. First, in order to answer this question, a literature review is performed. The next step is to define an evaluation method to assess the performance of 'Guiding the Cloud' Simulation-based assessments can support the testing of early-stage DC operation strategies, avoiding the safety and economic risks derived from real testing in physical DCs. Building Energy Simulation (BES) tools, based on a white-box modelling approach offer a suitable platform for the development of virtual test environments. Numerical models of different typologies of DCs (i.e. different geographical locations, cooling systems and on-site RES) are developed to achieve a sufficiently heterogeneous case study for the testing and tuning of the first prototype of the 'Guiding the Cloud' algorithm. Thus the potential of an optimal IT workload distribution at the wide world (cloud) level can be assessed.

GUIDING THE CLOUD

The original principle of 'Guiding the Cloud'(GtC) is based on a collaboration between the HVAC, power supply, and IT management systems to establish an optimum IT load distribution in an appropriate time (Deerns, 2012). Combining the measured performance data and weather data, enables a representation of the DC to be modelled. Then, using the weather forecast and predicted IT workload as an input, this representation can be used to estimate the DC's behaviour in time. Based on this information, the IT workload can be scheduled via the IT management system in such a manner that minimal energy consumption is needed. If the goal was to optimize other performance indicators than electricity consumption and CO_2 emissions, additional information from other sources could be used. For example, including the electricity contract enables us to minimize the electricity costs. Or, by predicting the availability of different systems, the maintenance and IT workload schedule can be configured to maximize the reliability of the DC.

Besides the collaboration between management systems, as described above, another important element of 'Guiding the Cloud' is the focus on multiple DCs at the same time. More specifically, DCs are increasingly operating worldwide with activities dispersed globally. To support these worldwide activities, multiple IT resources are often geographically scattered. The use of these IT resources would be more beneficial, if worldwide IT resources were globally integrated by virtualization techniques, whereby the surplus of the IT resources could be used by another location (Stanoevska-Slabeva et al., 2010). Considering multiple DCs, as such, would enable a global optimization of the DCs' performance.

The DCs' performance is a key element in the decision to distribute IT workload from one location to another. However, the performance of the data network used to transport the IT workload should also be included in decision-making (Taal et al., 2013). In a global scale data network the performance is dominated by the switching infrastructure rather than the transport infrastructure (Tucker, 2008). The switching infrastructure contains, for example, switches and routers while the transport infrastructure includes the line amplifiers, regenerators and optical switches.

Selecting a strategy to distribute IT workload among several DCs, including the DC's and infrastructure's performance, represents a challenge due to the different domains involved in the decision. For example each DC and infrastructure has its own characteristics considering HVAC, power supply systems, IT facilities and possible on-site RES. Besides, the characteristics associated with their location such as climate conditions and net electricity generation must also be considered. Many of the new technologies to reduce electricity consumption or costs have been developed to reduce a single objective; however, this paper focuses on simultaneously reducing multiple objectives.

ASSESSMENT METHOD

As previously stated, a virtual environment is developed to assess the GtC control algorithm. The method presented in this paper consists of two types of model:. *Prediction* and *simulation* models.

As part of the prediction model, the controller requires a fast response in order to evaluate a large number of alternatives for the optimization of the chosen objectives. Based on this requirement, simplified data-driven (black-box) models were selected.

The lack of documentation and measured data from real DCs, generally due to financial and security factors, often delays or even discards the development of novel control strategies.

The *simulation model* represents a virtual test environment for DCs that generates the needed performance indicators for an early stage assessment of the GtC control algorithm presented in this paper. It is a physics-based (white box) model developed by a BES tool (TRNSYS), and it is used firstly for an identification of prediction models and then for the final evaluation of the GtC control concept. Figure 1 shows the different models involved in the process as well as their function.

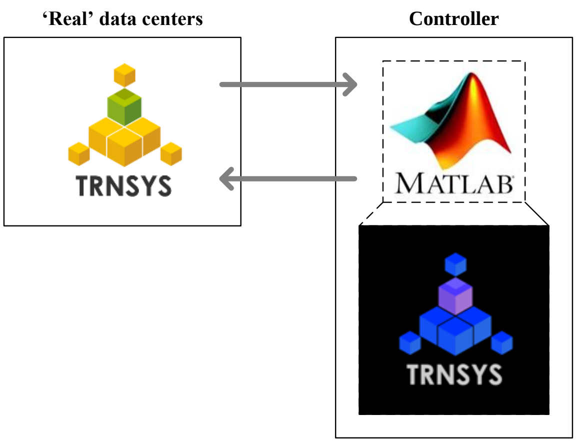

Figure 1: Schematic representation of the test environment consisting of the 'real' DC model (TRNSYS) and the controller (Matlab) during simulation run-time. The reference models resides inside the controller; communication inside the controller is based on Matlab scripts

CONTROL OPTIMIZATION OF GUIDING THE CLOUD

This section describes the process and the different elements that form the control optimization algorithm of the 'Guiding the Cloud' concept.

Schematic overview

Figure 2 presents the method of 'Guiding the Cloud' to optimize the performance indicators by varying the IT workload's distribution sequence over different DCs. The optimization process considers multiple solutions by altering the control sequence of the IT workload. When the Pareto Front is found, the decision maker selects the control sequence

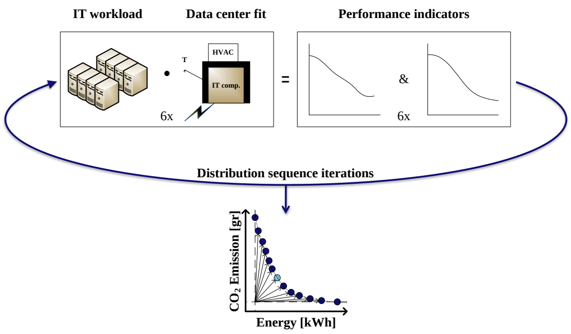

Figure 2: Schematic overview of the performance assessment methodology for 'Guiding the Cloud'.

Model predictive control

The algorithm is based on model predictive control (MPC) rather than on conventional rule-based control strategies. A key element of MPC is that it searches for the best control strategy using a data-driven, black-box model of the DC and its systems. This model can be implemented with non-linear and complex interactions in multiple-input-multiple-output systems (Gyalistras et al., 2010). Other peripheral matters in conjunction with the MPC of 'Guiding the Cloud' are described in more detail further on in this section. The MPC algorithm is based on the following steps:

1. Altering the control sequence
2. Predicting the performance of 'Guiding the Cloud' using the control sequence
3. Selecting an optimized control sequence
4. Shifting forward and updating the boundary conditions and starting again with step 1 until the total number of control sequence variants has been executed

The model developed to predict the DC performance influences the quality of the control sequence of the MPC. Also the optimization horizon and the control horizon length affect this performance. Further, the optimization method and objective function have an impact on the end result.

Multi-objective optimization

The core of the control algorithm is a multi-objective optimization. Generally, the optimization problem consists of two (or more) conflicting objectives, for example minimizing the energy demand while maximizing thermal comfort. It is impossible to find just one best design solution for these so-called multi-objective optimization problems. Instead, a set of 'trade-offs' or good compromise solutions are all Pareto optimal (i.e. an increase in one objective would simultaneously lead to a decrease of the other objective).

In this research, the multi-objective optimization is performed over two objectives, electricity consumption and CO_2 emissions. At first, these objectives appear complementary rather than conflicting, because a decrease in energy should result in a lower exhaust of CO_2 emissions. However, when IT workload is distributed to another DC due to a higher energy efficiency, it is also possible that the CO_2 emissions increases when the other DC produces relatively more CO_2 in generating the net electricity.

Control strategy

The algorithm optimizes the objective values by varying possible distribution sequences. The results of the algorithm leads to a Pareto front with optimal distribution sequences. The selection of the final distribution sequence depends on the weight that is given to the objectives (i.e. CO_2 emissions and electricity consumption). This final choice may vary depending on who is the stakeholder or final decision maker.

In the initial assessment presented, the minimum value of the first objective, the minimum value of the second objective and a trade-off solution are presented to examine the savings potential of 'Guiding the Cloud'.

The trade-off solution could be based on three types of indicators (Hoes, 2014), i.e. (1) robustness indicator, (2) robustness vector and (3) robustness balance. The robustness vector is used in this assessment, because it is independent of user preferences. The robustness vector depends on the distance from a utopian point to a solution. The length of the robustness vector is a quantified indicator of the robustness of the solution. The shortest distance is the most robust solution.

The genetic algorithm, from the Matlab Optimization Toolbox, was used for the generation of a reduced number of control sequence variants that are needed to obtain a Pareto front.

Prediction model

A controller, developed in Matlab, uses prediction models to optimize the IT computational workload distribution, in order to minimize the total electricity consumption and CO_2 emissions.

The prediction models can be identified as data driven black-box models (specifically artificial neural networks). As stated before, the lower computational time required was a decisive factor in selecting this modelling approach.

The prediction models are based on a state space characterization, which deals with each DC with several inputs and outputs (MIMO data). The inputs consist of (1) dry bulb temperature, (2) relative humidity, (3) wind speed, (4) total horizontal radiation, and (5) IT workload. The outputs are (1) electricity consumption and (2) emission, which are the objectives to be optimized.

VIRTUAL TEST ENVIRONMENT

This section presents the virtual test environment developed for testing. The virtual test environment for ‘Guiding the Cloud’ is created in TRNSYS and is used to simulate the energy consumption and CO_2 emissions of different typologies of DCs that constitute our case study. TRNSYS has been selected due to its extensive library with pre-defined components for building models, HVAC systems and RES. In addition, TRNSYS contains the functionality to directly embed other software tools (e.g. Matlab/Simulink).

The data simulated by TRNSYS are used to train the surrogate black-box models in Matlab and then the overall model-based controller is embedded into the TRNSYS environment so the influence of ‘Guiding the Cloud’ concept can be evaluated by comparing with a baseline scenario.

Case study description

The case study, that is described below, is used to examine the potential of ‘Guiding the Cloud’. It includes a description of the relevant DCs and IT characteristics, model assumptions, IT scenarios, and performance indicators.

Locations

Six imaginary DCs are located around the world as presented in figure 3.

Figure 3: Locations

The DCs are located in different climate and time zones. Table 1 provides an overview of the climate and time zones of each location.

Table 1
Climate and time zones for each DC location (Lord Toran)

LOCATION	CLIMATE	TIME ZONE (GMT)
Sacramento	Subtropical zone	-8
New York	Moderate zone	-5
Madrid	Moderate zone	+1
Bergen	Cold zone	+1
New Delhi	Tropical zone	+5*
Sydney	Moderate zone	+10

* The time zone has been rounded from 5.5 to 5 hours.

The time zone of Madrid and Bergen is selected as reference time zone (meaning that the data for Sacramento has been shifted by -9 hours). The performance of each DC is influenced differently due to the various climate circumstances and time zones.

Data center characteristics

The specific DC characteristics of each DC are assumed to be identical. Each DC has a capacity of 2000 kW of IT resources available. The main difference between the DCs is the type of HVAC systems they have for cooling the IT resources. Four typical configuration of HVAC systems for DCs have been considered, these are (1) chillers, (2) chillers in combination with dry coolers, (3) sea water absorption cooling (SWAC) and (4) indirect evaporation cooling unit (IECU). The HVAC systems are chosen based on typical solutions for the climate conditions of the different locations selected. Table 2 lists the type of HVAC systems for each location considered in the case study.

Table 2
HVAC systems assigned for each location

LOCATION	HVAC SYSTEM
Sacramento	Dry coolers
New York	IECU
Madrid	IECU
Bergen	SWAC
New Delhi	Chillers
Sydney	Chillers

Renewable energy systems (RES)

Based on the efforts to achieve net-zero and even positive energy buildings, DCs try to obtain as much energy as possible from on-site RES. The DCs in this case study are also equipped with different types of on-site RES. These are presented in table 3. The peak power generated by the on-site RES is also provided.

Table 3
RES systems assigned for each location

Location	RES system	Power [kW]
Sacramento	PV panels (4000m^2)	P_{peak} (462.7)

	Bio gas turbine	400
New York	-	-
Madrid	PV panels (2000m^2) Bio gas turbine	P_{peak} (230.4) 300
Bergen	Wind turbine(10pcs)	P_{nom}(2000)
New Delhi	Wind turbine(5pcs) PV panels (3000m^2) Bio gas turbine	P_{nom}(1500) P_{peak} (349.7) 600
Sydney	PV panels (2000m^2) Bio gas turbine	P_p (252.4) 1000

Internal and external data distribution network

The components of a data distribution network consume electricity while processing and distributing data.

Each DC has its own internal distribution network, called a local area network (LAN), which connects the computing infrastructures and storage capacities with the outside world. The LAN components contain a host (network interface), switches, firewalls, and routers (Taal et al., 2010).

The DCs are interconnected by an external dedicated distribution network, a light path network. This network is used to distribute the data from one location to another. When data is transported over the light path network it passes hops. Each hop contains two dense wavelength division multiplexing (DWDM) nodes. The number of hops between the DCs depends on the location. Three hops are assumed for transporting data in the same continent. An extra hop is added when data is transported to another continent. The contribution of other components, for example line amplifiers, regenerators or optical switches, are excluded from the case study, because the DWDM nodes have a significantly larger electricity consumption (Taal et al., 2010).

When IT workload crosses over to another data network, it takes time to arrive at the final destination. The transport time depends on the bandwidth and latency of the network. The bandwidth is the amount of data that can travel through a network at a time. The latency is the speed of sending IT workload from one location to another. To analyse the potential of 'Guiding the Cloud' the transport time for IT workload distribution is ignored in this research.

Total physical case study

Figure 4 illustrates the complete physical case study to examine the potential of 'Guiding the Cloud'.

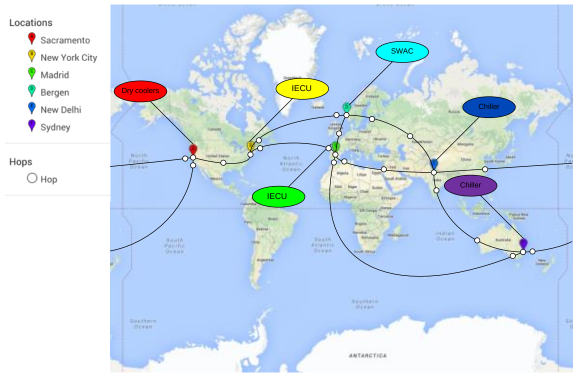

Figure 4 Case study

Weather data

The weather data used in the simulations represent a typical meteorological year, which is based on an 'average' from historical data for each location. The weather data contains hourly values of different climate parameters. The weather data is manually edited to correspond with the time zones of each DC.

IT workload

Three different IT scenarios have been defined processing (CPU-intensive), software (interactive) and data storage (hot and cold) (Taal et al., 2010).

In this research, the software (interactive) scenario has been selected to examine the potential of 'Guiding the Cloud'. Three input parameters characterize this scenario. The first is the amount of input data, the second is the CPU processing time, and the third is the output data. The CPU processing time and output data are assumed to be dependent on the input data by the following two formulas:

$$t_{CPU} = 0.001 \cdot InputData^2 + 0.17 \quad (2)$$

and

$$OutputData = 0.15 \cdot InputData + 0.10 \quad (3)$$

When IT workload is distributed from one DC to another, the amount of input data changes, which results in a change of the CPU processing time and output data.

The input data consists of a combination of fixed and variable IT workload. The fixed IT workload represents the IT requests that should be done locally, while the variable IT workload can be distributed among the other DCs. TRNSYS uses a predefined weekly profile of the IT workload. The IT workload of each DC is assumed to be identical. Yet, due to differences in time zones between the DCs, there are also differences in the IT workload between DCs at a certain time of a day. Figure 5 presents the IT workload daily profile for each DC.

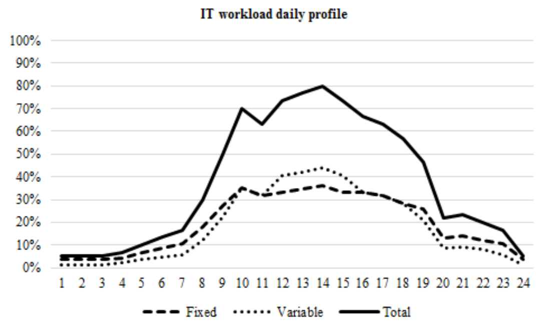

Figure 5 Predefined daily profile of the IT workload for each DC. The IT workload will shift in correspondence with time in the time zone.

The IT workload profile describes the internal gains of each DC. During the day the IT workload will cause a higher internal gain in the DC than at night.

Performance indicators

The multi objective optimization considers two performance indicators, energy consumption and CO_2 emissions, that are simulated in TRNSYS for every DC of the case study. The energy consumption in a DC is simulated by a single zone model. Temperatures in the zone are controlled by the respective HVAC system. The CO_2 emissions are based on the energy consumption multiplied by a conversion factor. The conversion factors from energy consumption to CO_2 emissions are presented in Table 4.

Table 4
Conversion factor from energy to emission per location (IEA, 2014)

LOCATION	EMISSION [gr. CO_2 / kWh_e]
Sacramento	488
New York	488
Madrid	305
Bergen	246
New Delhi	926
Sydney	798

PERFORMANCE COMPARISON

The performance comparison consists of the following four simulation steps:

1. Gain results from the DCs using the predefined daily IT workload profile
2. Characterize state space models based on the results in the previous step
3. Integrate the state space models in the MPC, search for optimized distribution sequences and select one distribution sequence
4. Gain results from the 'real DCs using the selected distribution sequence

In the following paragraphs, two alternatives are compared, namely the performance of the combination of DCs applying and not applying the 'Guiding the Cloud' concept. In addition, the maximum savings of electricity consumption and CO_2 emissions will be presented.

TRNSYS model versus black-box models

It is a challenging task to characterize a black-box model. In order to obtain a black-box model, usually a trial-and-error process precedes. The black-box model has been characterized by varying the number of inputs and outputs, training data and order of the state space model. The best results are obtained when different amounts of weeks in January are used for training and when various higher orders of the state space model for each DC are used. The results from the state space model are compared with the first two weeks of comparison data of the results in February from the TRNSYS model. The normalized root mean square error (NRMSE) fitness value provides an indication about the goodness of the fit. The fit is calculated (in percentage) using the following formula:

$$fit = 100\left(1 - \frac{\|y - \hat{y}\|}{\|y - \bar{y}\|}\right) \tag{1}$$

Where y is the validation data output (i.e. output of TRNSYS) and $\hat{y}$ is the output of the state space model (i.e. output of Matlab).

The goodness of the fit is determined by the NRMSE. Results are presented in figure 6.

As presented, most DCs show a good fit of approximately 90% or higher. However, the estimations of the DCs' CO_2 emissions in New York and Bergen differ significantly from an optimum fit.

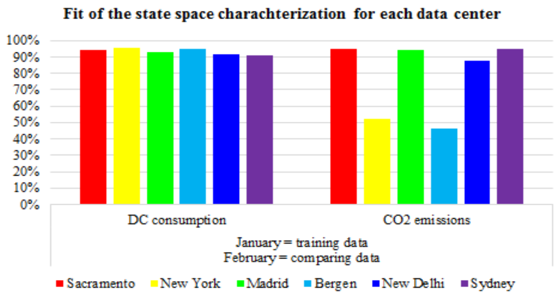

Figure 6: Best fit for each DC for a day in February

Day in February

The 8th of February has been selected as the day, to examine the potential of 'Guiding the Cloud'. Figure 7 presents the total electricity consumption based on the IT workload distribution sequence of the reference case (where 'Guiding the Cloud' is not applied), and the, optimization, minimum and maximum case (where 'Guiding the Cloud' is applied).

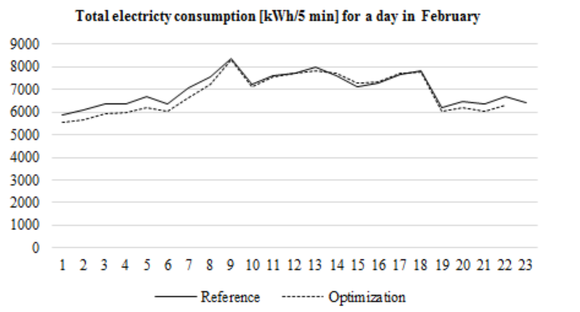

Figure 7: Results of the total energy consumption divided in the reference, fit, and controller case for a day in February

Figure 8 presents the total CO2 emissions based on the IT workload distribution sequence of the reference, optimization, minimum and maximum case.

As presented in figure 7 and figure 8, the optimization algorithm results in an optimization of the total electricity consumption and CO_2 emissions. Figure 8 provides an overview of the results that are obtained for the different cases compared to the reference case.

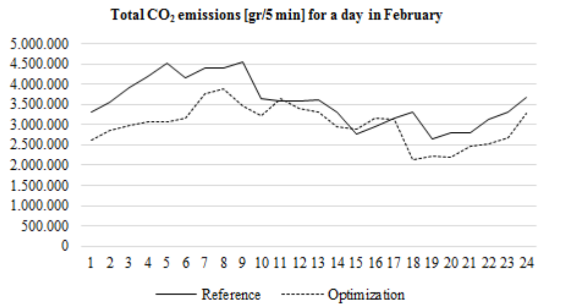

Figure 8: Results of the total CO_2 emissions divided in the reference, fit, and controller case for a day in February

As presented in Figure 9, using the decision maker (coloured as yellow) to select the distribution strategies results in electricity consumption and CO_2 emissions savings of 3.0% and 15.5% respectively. Selecting the minimum values for the electricity consumption from the Pareto solutions (coloured as blue), results in a decrease of 0.6% for the total electricity consumption and a decrease of 0.7% for the total CO_2 emissions. Selecting, on the other hand, the minimum values for the CO_2 emissions from the Pareto solutions (coloured as green), results in a decrease of 17.0% in CO_2 emissions and a 2.5% decrease in electricity consumption..

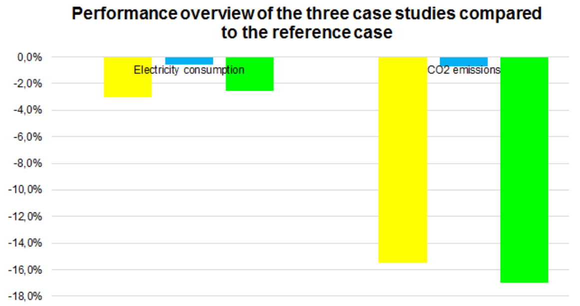

Figure 9: Performance overview of the results of the total electricity consumption and CO_2 emissions for 1 day.

It is observable that applying the minimum values for electricity consumption does not lead to the maximum reduction in electricity consumption. This can be explained by selecting a distribution sequence which will influence the future distribution sequences. In this case, the influence of the decision-making resulted in a reduction of both electricity consumption and CO_2 emissions.

Overall, the results show that by applying ‘Guiding the Cloud’ it is possible to reduce energy consumption and CO_2 emissions of the case study. The case study that is used in this research indicates a maximum reduction of 2.5% in total electricity consumption and 17.0% in total CO_2 emissions.

DISCUSSION

Data centers’ representation

The case study presented in this paper is based on physics-based models of six typical configurations of DC’s located in different climates. This case study was created to assess the potential of the ‘Guiding the Cloud’ concept. Real DCs and measurements to calibrate these models are necessary in the future to validate this concept.

In order to improve the reliability of the result of ‘Guiding the Cloud’, the representation of the DCs in the controller should be enhanced, because this will result in a more accurate prediction of the distribution sequences.

Multi-objective optimization

While conducting this research, a potential drawback of the genetic algorithm was detected. With this objective optimization technique, no weighting factors can be applied to emphasize more on one objective. The preferences of the user can only be implemented using the decision maker to select a distribution strategy.

In the case study, which was presented in this research, the ‘optimization’ case contained an unprejudiced control strategy for selecting the distribution sequence. This means that the decision maker has no clear preferences for the two objectives. However, it would be better if the decision maker considers the preferences of the user when selecting the distribution sequence.

Transport time of the IT workload distribution

The IT workload consists of a fixed and a variable IT workload. The fixed IT workload represents the just-in-time, just-in-place IT requests of customers while the variable IT workload could be distributed to another DC. However, transport time for distributing IT workload from one DC to another DC has been ignored. This research has focused on a simple setup of the problem without bandwidth or latency constraints. Future work should include the transport time.

CONCLUSION AND FUTURE RESEARCH

In this research, the potential of distributing IT workload among geographically dispersed DCs is investigated. The main objective is to examine its (possible) effect on reducing the total energy consumption and CO_2 emissions. The preliminary results indicate that a reduction in total energy consumption and CO_2 emissions is achievable. Savings of the total electricity consumption and CO_2 emissions can be up to 3% and 17.0% respectively.

In future research, the concept presented in this research could be extended with a time delay when distributing IT workload from one DC to another. Besides the time delay, other relevant aspects can also be included, such as maximizing or setting boundaries for the reliability or minimizing the energy costs using spot pricing markets of the DCs. Another interesting concept is to combine a local and a global optimization. The local optimization will internally schedule the IT workload while the global optimization distributes the IT workload among the geographically dispersed DCs.

ACKNOWLEDGEMENT

This research was conducted as collaboration between the Technical University of Eindhoven and Deerns Nederland B.V.

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

Deerns 2012. Deerns Sets the Standard for Chasing Energy Reduction in Cloud Data Centers. Inside Networks. February, 2012.

Gyalistras, D, The OptiControl Team. 2010. Final report, Use of Weather and Occupancy Forecasts for Optimal Buidling Climate Control (OptiControl). ETH Zurich, Systems Ecology, Institute of Integrative Biology, Switserland.

Greenberg, S., Mills, E., Tschudi, B., Rumsey, P., & Myatt, B. 2006. Best Practices for DCs: Lessons Learned from Benchmarking 22 DCs. ACEEE.

Hoes, P. 2014. Computational Performance Prediction of the Potential of Hybrid Adaptable Thermal Storage Concepts for Lightweight Low-Energy Houses.

IEA. 2014. CO_2 emissions from fuel combustion.

Mohsenian-Rad, A.H., & Leon-Garcia, A. 2010. Coordination of Cloud Computing and Smart Power Grids. IEEE Smart Grid Communications Conference. Gaithersburg, Maryland.

Mills, M.P. 2013. The Cloud Begins With Coal, Big Data, Big Networks, Big Infrastructure, and Big Power.

Oró, E., Depoorter, V., Garcia, A., Salom, J. 2015. Energy efficiency and renewable energy integration in data centres. Strategies and modelling review, Renewable and Sustainable Energy Reviews, Vol. 42, pp. 429-445.

Rao, L., Liu, X., Ilic, M. D., & Liu, J. 2012a. Distributed Coordination of Internet DCs Under Multiregional Electricity Markets. Proceedings of the IEEE, Vol. 100(1), pp. 269-282.

Rao, L., Liu, X., Xie, L., & Liu, W. 2012b. Coordinated Energy Cost Management of Distributed Internet DCs in Smart Grid. Smart Grid, IEEE Transactions on, Vol. 3(1), pp.

Sharma, R.K., Shih, R., Bash, C., Patel, C., Varghese, P., Mekanapurath, M., . . . Kumar, M. 2008. On building next generation DCs: energy flow in the information technology stack. COMPUTE '08 Proceedings of the 1st Bangalor.

Stanoeska-Slabeva, K. Wozniak, T. 2010. Grid and cloud computing , Introduction: Business and Technological Drivers of Grid Computing. Springer.

Taal, A., Grosso, P., Bomhof, T. 2013. Transporting Bits or Transporting Energy: Does it matter? SURF.

Tucker, R. 2008. Optical packet-switched WDM networks - A cost and energy perspective. Optical Fiber Communication Conference. San Diego, CA.

Wit de, M.S. 2001. Uncertainty in predictions of thermal comfort in buildings, PhD thesis, Delft University of Technology, the Netherlands