

# OPTIMISING EGOVERNMENT DATA CENTRE OPERATIONS TO MINIMISE ENERGY CONSUMPTION: A SIMULATION-BASED ANALYTICAL APPROACH

*Research in Progress*

## **Abstract**

*The energy consumption of data centres is increasing over and over. However, there is no or few decision support system introduced for data centre practitioners to use it for their daily operations including simulating their new server installation and forecasting power consumption for the target periods. We propose a simulation model based on CloudSim, which is widely used for data centre research. Our simulation model will be tested with datasets including IT workloads, cooling performance, and power consumption of servers and cooling devices. In the final stage, we provide a decision support system to monitor, forecast, and optimise the power consumption of a data centre easily on the web.*

*Keywords: data centre; optimisation; simulation; decision support system; greendc.*

# 1 Introduction

In recent years, the use of the information and communications technologies (ICT), comprising communication devices and/or applications, DCs, internet infrastructure, mobile devices, computer and network hardware and software and so on, has increased rapidly. Internet service providers such as Amazon, Google, Microsoft, etc., representing the most significant stakeholders in the IT sector, constructed a large number of geographically distributed internet data centres (IDC) to satisfy the growing demand and providing reliable low latency services to customers. These IDCs contain vast numbers of servers, large-scale storage units, networking equipment, infrastructure, etc., to distribute power and provide cooling. The number of IDCs owned by the leading IT companies is drastically increasing every year. As a result, the number of servers needed has reached an astonishing number. EPA (Environmental Protection Agency) reports that IDCs in the US use 61 billion kWh in 2006 (1.5% of US energy consumption), 91 billion kWh in 2013. Natural Resource Defence Council expects energy consumption of IDC will be increased up to 140 billion kWh annually by 2020 (Ni and Bai 2017).

Due to huge amount of energy implied and the related cost, IDC can make a significant contribution to the energy efficiency by reducing energy consumption and power management of IT ecosystems. It is why most researchers focus on reducing power consumption of IDCs. Those efforts include designing innovative architecture of DCs to minimise loss of cool airs, heats from IT devices, protecting from outside heats and so on. Also, computer scientists developed energy efficient workload reallocation algorithms to distribute the peaked workloads of servers to minimise energy consumption and heat generation.

However, efforts to reduce energy consumption via efficient DC operations are still scarce. The two significant energy consumptions in DCs are IT devices and cooling devices which consumers around 90% of total energy consumption in DCs. Specifically, many scholars reported that the power consumption of cooling devices are about 40 per cent (e.g., Rong et al. (2016)). It is believed that the energy consumption by two types of the device has a trade-off relationship as keeping lower temperatures in a DC will increase the energy consumption of cooling devices while decreases that of IT devices due to increased computing efficiency and vice versa. However, the literature has mixed results on the shape of the trade-off relationship via scientific experiments. Understanding the trade-off relationship will help DC managers to decide optimal temperature of their DCs as the increase of the temperature has a significant impact on the total energy consumption of cooling devices and the DCs. In other words, DC managers are yet to find answers to such questions as what are the optimal temperature of the DCs to minimize energy consumption without affecting the performance of IT devices and meeting the service-level-agreement? How many servers or virtual machines need to be on for the next 24 hours or a week considering expected workloads? What are the optimal schedules of servers and VMs for handling workloads which are various from time to time? Where are the best operational schedules of cooling devices within the DC to have a maximum cooling effect? To answer these questions, we develop a simulation model that can consider IT devices (workloads) as well as cooling devices in one model based on a real data collected from one of the largest data centres in Turkey that runs several commercial services on it. By utilising the real-time data from the data centre, we propose an accurate simulation model that can be utilised by practitioners to optimise their daily changing data centre configuration.

This paper aims at finding answers to those questions for DC managers via a multi-disciplinary study. It takes a more holistic view by considering the system including servers, cooling system, backup power and electrical distribution system. Notably, a decision support system (DSS) that integrates functionalities including workload and energy forecasting, generation of optimal operation scheduling of cooling and IT devices and simulation for impact analysis of DC operation strategies. Since the configuration of data centres is continuously changing over time, it is necessary for managers to pro-

vide forecasting (or simulation) tool for a new configuration. There are some suggestions and trials to develop such decision support systems (Gupta et al. 2014; Sankar et al. 2013) though there are no or few tools available for practitioners yet.

## 2 Literature review

Studies on energy consumption are primarily focused on understanding the engineering perspective of DC components (e.g., Fulpagare and Bhargav 2015). For example, Bhopte et al. (2006) analyse effects of underfloor blockages on data centre performance using computational flow dynamics (CFD) simulator and report that some blockages under floor impact larger on cooling performance. Iyengar et al. (2012) design server liquid cooling systems without a chiller. They report that they can reduce 90 per cent of cooling energy on a relatively hot New York summer day. There are similar works that introduce new technologies to reduce cooling energy on data centres (e.g., Habibi Khalaj and Halgamuge 2017; Zimmermann et al. 2012).

Zhang, Wei, and Zhang (2017) investigate the impact of several free cooling technologies (economiser) that are proper for a newly built data centres. They evaluate four free cooling methods that are currently available. Ham, Park, and Jeong (2015) offer optimum supply air temperature ranges for each economiser set-up and climate condition. The other approaches for the economiser studies are to increase energy efficiency by reusing the waste heat from data centres (e.g., Ebrahimi et al. 2014) rather than reducing the waste heat itself. Cho et al. (2017) analyse economic benefits for the seven options of data cooling strategies that include several kinds of economiser. Song et al. (2015) also calculate economic and operational benefits using chillers and economiser operating hours of the worldwide selected data centres. Cho and Kim (2016) provide conceptual models that comprise the major cooling technologies and show the simulated results for each cooling technologies including economisers and several energy sources such as renewable energy considering climate zone and energy-saving levels.

The other stream of data centre research focuses on optimisation studies to minimise power consumption without impact on service level agreement. In the early stage of workload optimisation, Bash and Foreman (2007) can increase cooling efficiency by reallocating heavy-loaded servers on the cold area in a server room. They use computational fluid dynamics (CFD) simulation tool to sense the cold spots and hot spots in server rooms; they migrate high workloads into servers on the cold spots. More recent works using virtual machine reallocation strategies for the server room optimisation. Liu et al. (2009) implement a framework that can monitor the energy consumption of a DC and migrate virtual machines to minimise the energy consumption into idle servers. They report that the new architecture can save up to 20 per cent of energy consumption. Chen et al. (2012) extend the prior research: their new algorithm can optimise virtual machines over the multiple data centres over the world considering daily/seasonal effects on each location. Chen and his colleagues assert that this approach reduces up to 40 per cent energy consumption in comparison with other scheduling policies. Unlike previous studies, Parolini et al. (2008) proposed an algorithm that controls both cooling management and workload reallocation using the constrained Markov decision process (CMDP) simultaneously. The authors report that the proposed control strategy is better than the traditional schemes that control computational and cooling subsystems separately. Banerjee et al. (2010) propose a similar approach to Parolini et al. (2008)'s and can achieve up to 15 per cent energy saving.

Other groups of scholars focus on optimising network traffic in data centres. Gao et al. (2012) introduce a flow optimisation based framework for request-routing and traffic engineering. By adopting this framework, they can reduce 10 per cent of carbon footprint without increasing traffic latency or electric bills. Another team adopts a similar approach to optimise network traffic and reports that they can reduce up to 50 per cent of network energy (Heller et al. 2010).

Even though the guideline of ASHRAE is broadened to higher inlet air temperature (refer the table 1 on Ham et al. 2015, p. 165), most practitioners try to keep their cooling environment conservative as before. But, recent studies show that increasing temperature do not harm IT devices in data centres as well as decrease energy consumption on cooling data centre (e.g., Mone (2012)). For example, El-Sayed et al. (2012) check failure rates of most of the server components when cooling air temperature in data centres is increased. The authors report that increased inlet air temperature does not influence the server and its component reliability until 40°C. They also find that increased inlet air temperature does not decrease the performance of the server and its components as well as there is no CPU or hard drives' throttling occurred up to 40°C inlet air temperature. Breen et al. (2010) also show that to increase the 1°C inlet air temperature makes 4-5 per cent of data centre energy costs without increasing equipment failure.

To show the impact of the increasing inlet temperature, we need a simulator that can calculate the difference between the inlet temperature change. Until now, thermodynamic fluid models are the mainstream of the cooling simulation that are used to calculate cooling energy consumption, visualise air flows within the data centre, and illustrate hot and cold spots in the server room. This approach is generally called computational fluid dynamics (CFD) models (Alkharabsheh et al. 2015; Choi et al. 2008; Ni et al. 2017; Rambo and Joshi 2007). CFD models can generate accurate results for given parameters. However, CFD models need high computational power (and computing time), professional knowledge to model the target and commercial tools in general.

Other groups of scholars generate black-box models for cooling devices to estimate power consumption of cooling devices in a data centre (Dayarathna et al. 2016; Ham et al. 2015; Liu et al. 2012; Mousavi et al. 2015; Pakbaznia and Pedram 2009; Patterson 2008). This black-box modelling approach has some advantages compared to computational fluid dynamics models. Firstly, the black-box model is generally more straightforward than the CFD model. This feature is a significant advantage when we need to calculate many objects in a simulation model. The simplicity also helps us to use the black-box model in the everyday working environment. Secondly, it is essential to have professional knowledge on thermodynamics as well as data collection for CFD modelling. On the other hand, the black-box models can be described by basic mathematics. It helps us to adopt widely in a general data centre working environment.

### **3 Simulation model**

We design a simulation model with a straightforward black-box model from prior studies.

#### **3.1 An example of a server room in a data centre**

Figure 1 depicts an example of a server room. There are two types of devices in a server room. One is IT devices that include server racks, servers, and network devices. The other type is non-IT devices including CRACs, chillers, power distribution units (PDUs) and UPS. The primary power consumption of a server room comes from servers and cooling devices (air conditioners and chillers). Air conditioners supply cold air at the constant air volume (Cho et al. 2017, p. 236). So, the cold air is supplied more in the middle of the server room compared to the other locations as described in figure 1.

Recently, data centre owners resell some space of a server room to other companies that can install their own servers. It is called as 'collocation.' The collocation servers are maintained by the data centre owner, but there is no right for the data centre to access the data on it. It is a very general situation for data centre owners to reduce operating costs by utilising unused resources (server racks). But, because of the collocation service, it is impossible to measure all data from a server room as well as real-

locate workloads or virtual machines in a server room. To overcome this limitation of data measurement, we propose a simple conceptual model of a server room in the next section. After the basic conceptual model development finished, we extend our conceptual model that can consider the collocation issue in a DC.




Figure 1. An example of a server room in a data centre (Ni and Bai 2017, p. 628)

### 3.2 A simple conceptual model of a server room in a data centre

In this section, firstly, we define a conceptual model including IT devices (servers) and non-IT devices (air conditioners) like figure 2. For simplicity, we assume that there is only an air conditioner (CRAC) in a server room on the basic model. Secondly, we describe how the conceptual model can be generalised and how we define a server room using the simple conceptual model.

From the many black-box models for cooling devices (Dayarathna et al. 2016; Ham et al. 2015; Liu et al. 2012; Mousavi et al. 2015; Pakbaznia and Pedram 2009; Patterson 2008), we choose the simplest power model for CRAC (Dayarathna et al. 2016, p. 770) and modify it slightly. As we know the maximum power consumption of the CRAC from the manufacturer's technical specification, we replace the summation of power consumption of servers. The modified power model for CRAC is

$$c_k = \frac{P_{max}}{\eta}$$

$$\eta = 0.0068t^2 + 0.0008t + 0.458$$

, where  $c_k$  is the power consumption of the  $k^{\text{th}}$  CRAC unit.  $P_{max}$  is the maximum power consumption of the CRAC.  $\eta$  is the coefficient of performance (CoP). The second equation is given for the water chilled commercial CRAC.  $t$  is the supply (inlet or target) air temperature of the CRAC unit. Our power model is only related to the target air temperature.

Before we describe the model, we assume the simplest case of a data centre: there are only one server rack and one air conditioner in a server room of a data centre. In the server rack, there are several servers installed. Each server also contains several virtual machines (VMs). Each VM handles the given workload from the simulator.




Figure 2. The simple conceptual model of a server rack/air conditioner model for a data centre

In the inside of the server rack, each server uses fans to flow the cold air into the inside of the server. The cold air that enters the server through fans of each server takes heat from the CPU, memory, and hard disk and flows out of the server. The heated air expelled from servers goes up to the ceiling. Under 35°C, a server consumes electricity according to CPU utilisation. Over 35°C, fans within the server start to blow air faster than the condition under 35°C and the server consumes more electricity (El-Sayed et al. 2012). The air conditioner expels heat out of the server room until the temperature of the heated air that is returned to the air conditioner is the same as the target temperature. During the operation, the air conditioner expels the collected heat out of the building by exchanging the heat with the chiller. Even if there is no heat removal, CRAC consumes the minimum electricity to run fans to flow air. The total power consumption of the data centre is the summation of power consumption of servers and the power consumption of CRAC.

### 3.3 Integrate the workload simulation model

CloudSim is a toolkit for modelling and simulation of cloud computing environments (Son et al. 2015; Goyal, Singh, and Agrawal 2012). Based on the authors, it was used by several researchers from organizations such as HP to investigate cloud resource provisioning and energy-efficient management of data centre resources (Calheiros et al. 2011). This toolkit is widely adopted and used until now when scholars investigate data centre with given workloads for virtual machines and physical servers. For example, Srivastana et al. (2017) use the CloudSim as a basis to develop, implement, and experiment the optimization algorithm for a cloud computing environment. It is well-designed and widely used toolkit for data centre research. However, CloudSim does not implement cooling devices for data centre, which consume the electronic energy of data centre up to 40~50% (Ohadi et al. 2012). This lack of implementation makes scholars hard to simulating the total power consumption of the target data centre using CloudSim. We implement this gap - the cooling devices with a simplified model of the datacentre.

One of the main shortcomings of CFD simulation is that CFD simulations assume the given workload as constant over time. However, the workload for each server fluctuate over time, and it affects cooling loads change in a DC. To consider the fluctuated workload, we adopt the CloudSim as a basis for the simulation and extend it to include the cooling devices on it.

## 4 Data collection and Expected results

There are two primary sources of data from the data centre in Turkey. One is IT device (server) data from Zabbix system. The other source is non-IT device measurement data from various sensors and metres installed. We chose one server room that is isolated to other server rooms. There are eight CRAC units, two external chillers connected to four CRACs respectively, two UPS, five rows of server racks. We chose two target server racks of which location is different from each other to collect a different set of data for each server rack.

There are more than 10 target servers in each target server rack. Each server is connected to the Zabbix system. Zabbix system creates logs of the physical and virtual machines' activities in real-time. For example, the logged data contains usage data of CPU, memory, hard drive, and network interfaces, number of processes in running states, and each server's specification. Zabbix system saves the information from the servers for every 5 minutes.

We installed energy metres, temperature sensors, smart PDU (Power Distribution Unit) to collect non-IT devices measurement data from various sources. We also connect non-IT devices that contain measurement functions for itself to a data centre infrastructure monitoring (DCIM) system. The installed devices are also connected to the DCIM system to collect the data. We install energy metres and temperature sensors on CRAC units. We install two SmartPDU to monitor a server's power consumption in the target racks. We also install three temperature sensors on the front and rear doors each to monitor inlet and outlet air temperature from the server racks. There are four temperature and humidity sensors under the floor and on the ceiling for each that can capture the supplied cold air temperature and hot air before consumed by CRAC units. The UPS and chillers provide their own functions by themselves, so we can receive the total power consumption in the target room from the UPS, and we can get the ambient air temperature and humidity outside of the data centre building. The DCIM system also records the data every 5 minutes. Data centre programmers provide RESTful API to receive both datasets at once.

Since the target data centre provides highly demanded public and commercial services, the implementation of a data collection system is in progress. It is expected to be done within the second quarter of 2019.

## 5 Conclusions

In this study, we develop a decision support system to support data centre operations. Specifically, we develop a simulation model that can count for the server workloads and cooling performance simultaneously to provide a simulation-based optimisation approach for the data centre. For example, the simulation results can show the optimal target temperature for the target periods by calculating the forecasting server workloads and the expected cooling capacities that can be defined by the server workloads and the outside temperature for the target periods. Practitioners also use the simulation model to determine their future server demands. In the future, we will extend our basic cooling model with more constraints that are given for the data centre operations.

## References

- Alkharabsheh, S., Fernandes, J., Gebrehiwot, B., Agonafer, D., Ghose, K., Ortega, A., Joshi, Y., and Sammakia, B. 2015. "A Brief Overview of Recent Developments in Thermal Management in Data Centers," *Journal of Electronic Packaging* (137:4), p. 040801.
- Banerjee, A., Mukherjee, T., Varsamopoulos, G., and Gupta, S. K. S. 2010. "Cooling-Aware and Thermal-Aware Workload Placement for Green HPC Data Centers," in *International Conference on Green Computing*, IEEE, August, pp. 245–256.
- Bhopte, S., Sammakia, B., Schmidt, R., Iyengar, M. K., and Agonafer, D. 2006. "Effect of Under Floor Blockages on Data Center Performance," in *ITHERM*, IEEE, pp. 426–433.
- Breen, T. J., Walsh, E. J., Punch, J., Shah, A. J., and Bash, C. E. 2010. "From Chip to Cooling Tower Data Center Modeling: Part I Influence of Server Inlet Temperature and Temperature Rise across Cabinet," in *ITHERM*, IEEE, June, pp. 1–10.
- Chen, C., He, B., and Tang, X. 2012. "Green-Aware Workload Scheduling in Geographically Distributed Data Centers," in *IEEE 4th International Conference on Cloud Computing Technology and Science Proceedings*, IEEE, December, pp. 82–89.
- Cho, J., and Kim, Y. 2016. "Improving Energy Efficiency of Dedicated Cooling System and Its Contribution towards Meeting an Energy-Optimized Data Center," *Applied Energy* (165), Elsevier Ltd, pp. 967–982.
- Cho, K., Chang, H., Jung, Y., and Yoon, Y. 2017. "Economic Analysis of Data Center Cooling Strategies," *Sustainable Cities and Society* (31), Elsevier B.V., pp. 234–243.
- Choi, J., Kim, Y., Sivasubramaniam, A., Srebric, J., Wang, Q., and Lee, J. 2008. "A CFD-Based Tool for Studying Temperature in Rack-Mounted Servers," *IEEE Transactions on Computers* (57:8), IEEE, pp. 1129–1142.
- Dayarathna, M., Wen, Y., and Fan, R. 2016. "Data Center Energy Consumption Modeling: A Survey," *IEEE Communications Surveys & Tutorials* (18:1), pp. 732–794.
- Ebrahimi, K., Jones, G. F., and Fleischer, A. S. 2014. "A Review of Data Center Cooling Technology, Operating Conditions and the Corresponding Low-Grade Waste Heat Recovery Opportunities," *Renewable and Sustainable Energy Reviews* (31), Elsevier, pp. 622–638.
- El-Sayed, N., Stefanovici, I. A., Amvrosiadis, G., Hwang, A. A., and Schroeder, B. 2012. "Temperature Management in Data Centers: Why Some (Might) like It Hot," in *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE joint international conference on Measurement and Modeling of Computer Systems (SIGMETRICS '12)*, pp. 163–174.
- Fulpagare, Y., and Bhargav, A. 2015. "Advances in Data Center Thermal Management," *Renewable and Sustainable Energy Reviews* (43), Elsevier, pp. 981–996.
- Gao, P. X., Curtis, A. R., Wong, B., and Keshav, S. 2012. "It's Not Easy Being Green," in *Proceedings of the ACM SIGCOMM 2012*, New York, New York, USA: ACM Press, p. 211.
- Gupta, S. K. S., Banerjee, A., Abbasi, Z., Varsamopoulos, G., Jonas, M., Ferguson, J., Gilbert, R. R., and Mukherjee, T. 2014. "GDCSim: A Simulator for Green Data Center Design and Analysis," *ACM Transactions on Modeling and Computer Simulation* (24:1), pp. 1–27.
- Habibi Khalaj, A., and Halgamuge, S. K. 2017. "A Review on Efficient Thermal Management of Air- and Liquid-Cooled Data Centers: From Chip to the Cooling System," *Applied Energy* (205:March), Elsevier, pp. 1165–1188.
- Ham, S. W., Park, J. S., and Jeong, J. W. 2015. "Optimum Supply Air Temperature Ranges of Various Air-Side Economizers in a Modular Data Center," *Applied Thermal Engineering* (77), Elsevier Ltd, pp. 163–179.
- Heller, B., Seetharaman, S., Mahadevan, P., Yiakoumis, Y., Sharma, P., Banerjee, S., and McKeown, N. 2010. "ElasticTree: Saving Energy in Data Center Networks," *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation*, pp. 17–17.
- Iyengar, M., David, M., Parida, P., Kamath, V., Kochuparambil, B., Graybill, D., Schultz, M., Gaynes, M., Simons, R., Schmidt, R., and Chainer, T. 2012. "Server Liquid Cooling with Chiller-Less Data Center Design to Enable Significant Energy Savings," *Annual IEEE Semiconductor*



- Thermal Measurement and Management Symposium* (1), pp. 212–223.
- Liu, L., Wang, H., Liu, X., Jin, X., He, W. B., Wang, Q. B., and Chen, Y. 2009. “GreenCloud: A New Architecture for Green Data Center,” *Proceedings of the 6th International Conference Industry Session on Autonomic Computing and Communications Industry Session*, pp. 29–38.
- Liu, Z., Chen, Y., Bash, C., Wierman, A., Gmach, D., Wang, Z., Marwah, M., and Hyser, C. 2012. “Renewable and Cooling Aware Workload Management for Sustainable Data Centers,” in *Proceedings of the 12th ACM SIGMETRICS*, New York, New York, USA: ACM Press, pp. 175–186.
- Mone, G. 2012. “Redesigning the Data Center,” *Communications of the ACM* (55:10), p. 14.
- Mousavi, A., Vyatkin, V., Berezovskaya, Y., and Zhang, X. 2015. “Cyber-Physical Design of Data Centers Cooling Systems Automation,” *Proceedings - 14th IEEE International Conference on Trust, Security and Privacy in Computing and Communications, TrustCom 2015* (3), pp. 254–260.
- Ni, J., and Bai, X. 2017. “A Review of Air Conditioning Energy Performance in Data Centers,” *Renewable and Sustainable Energy Reviews* (67), Elsevier, pp. 625–640.
- Ni, J., Jin, B., Zhang, B., and Wang, X. 2017. “Simulation of Thermal Distribution and Airflow for Efficient Energy Consumption in a Small Data Centers,” *Sustainability* (9:4), p. 664.
- Ohadi, M. M., Dessiatoun, S. V., Choo, K., Pecht, M., and Lawler, J. V. 2012. “A Comparison Analysis of Air, Liquid, and Two-Phase Cooling of Data Centers,” *Annual IEEE Semiconductor Thermal Measurement and Management Symposium*, pp. 58–63.
- Pakbaznia, E., and Pedram, M. 2009. “Minimizing Data Center Cooling and Server Power Costs,” *Proceedings of the 14th ACM/IEEE international symposium on Low power electronics and design - ISLPED '09*, p. 145.
- Parolini, L., Sinopoli, B., and Krogh, B. H. 2008. “Reducing Data Center Energy Consumption via Coordinated Cooling and Load Management,” *Proceedings of the 2008 conference on Power aware computing and systems*, p. 14.
- Patterson, M. K. 2008. “The Effect of Data Center Temperature on Energy Efficiency,” *2008 11th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, I-THERM*, pp. 1167–1174.
- Rambo, J., and 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.
- Rong, H., Zhang, H., Xiao, S., Li, C., and Hu, C. 2016. “Optimizing Energy Consumption for Data Centers,” *Renewable and Sustainable Energy Reviews* (58), Elsevier, pp. 674–691.
- Sankar, S., Kansal, A., and Liu, J. 2013. “Towards a Holistic Data Center Simulator,” *Microsoft Research*, p. 3.
- Song, Z., Zhang, X., and Eriksson, C. 2015. “Data Center Energy and Cost Saving Evaluation,” *Energy Procedia* (75), Elsevier B.V., pp. 1255–1260.
- Zhang, Y., Wei, Z., and Zhang, M. 2017. “Free Cooling Technologies for Data Centers: Energy Saving Mechanism and Applications,” *Energy Procedia* (143), Elsevier B.V., pp. 410–415.
- Zimmermann, S., Meijer, I., Tiwari, M. K., Paredes, S., Michel, B., and Poulikakos, D. 2012. “Aquasar: A Hot Water Cooled Data Center with Direct Energy Reuse,” *Energy* (43:1), Elsevier Ltd, pp. 237–245.