

ENERGY ANALYSIS OF REAR DOOR HEAT EXCHANGERS IN DATA
CENTERS WITH DYNAMIC WORKLOAD DISTRIBUTION

by

SAI ABHIDEEP PUNDLA

Presented to the Faculty of the Graduate School of
The University of Texas at Arlington in Partial Fulfillment
of the Requirements
for the Degree of

MASTER OF SCIENCE IN AEROSPACE ENGINEERING

THE UNIVERSITY OF TEXAS AT ARLINGTON

DECEMBER 2022

Copyright © by SAI ABHIDEEP PUNDLA 2022

All Rights Reserved

ACKNOWLEDGEMENTS

First of all, I would like to thank my supervising professor Dr. Dereje Agonafer for constantly advising and encouraging me. His invaluable advice proved to be extremely helpful during the entire course of my graduate studies. I wish to thank my committee members Dr. Haji Sheikh and Dr. Miguel Amaya for their interest in my research and for taking time to be in my committee.

I would like to express my deep gratitude to my parents, Radha Rani Sreeram and Rama Krishna Pundla, who have encouraged me through every step i have taken and also sponsored my graduate studies. I would also like to extend my appreciation to my brother Sandeep and Sister in law Pratyusha

I am especially grateful to Vibin Simon Shalom for his interest in my research and for the helpful discussions and invaluable comments. I wish also to thank my friends from EMNSPC for taking the time to critically evaluate this manuscript and my research work.

I am grateful to all the professors who taught me during the years I spent as a graduate student.

DEC 08, 2022

ABSTRACT

ENERGY ANALYSIS OF REAR DOOR HEAT EXCHANGERS IN DATA CENTERS WITH DYNAMIC WORKLOAD DISTRIBUTION

SAI ABHIDEEP PUNDLA, M.S. Aerospace Engineering

The University of Texas at Arlington, 2022

Supervising Professor: Dr. Dereje Agonafer

In recent years, applications including IoT (Internet of Things), content delivery, and 5G have created a large demand for low-latency access to data processing and data storage. Traditional centralized data centers weren't designed with those use cases in mind. Small data centers such as edge computing data centers and colocation data centers house more than half of all servers across the United States. Rear Door Heat exchangers (RDHx) provide an energy-efficient cooling solution to traditional CRAC/CRAH - based cooling methods by localizing the heat removal from the rack. With more control over air distribution through a shorter path between the hot air and the heat removal, this cooling method is efficient and predictable and can easily be implemented into existing data centers. In this study, RDHx implemented in a Edge computing center and Colocation Data center models using commercially available CFD software (6SigmaRoom). TCO (Total Cost of Ownership) and Cooling costs of the data center models are calculated and compared to traditional CRA based cooling methods

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	iii
ABSTRACT	iv
LIST OF ILLUSTRATIONS	vii
LIST OF TABLES	viii
Chapter	Page
1. INTRODUCTION	1
1.1 Data Centers Energy Consumption	1
1.2 Heat Generation and Cooling of Data Centers	2
1.3 Rear Door Heat Exchangers	3
1.4 Total Cost of Ownership	4
2. Literature Survey	5
3. Methodogy	7
3.1 Hot Aisle Containment	7
3.2 Data center layout	8
3.2.1 case A	10
3.2.2 case B	10
3.3 CAD modelling	11
3.4 Simulation termination strategy	12
4. Results and Discussion	14
4.1 Total Cost of Ownership	14
4.2 Assumptions	17
5. Conclusion	19

REFERENCES 20

LIST OF ILLUSTRATIONS

Figure	Page
1.1 Trends in heat loads of the Data Center	2
1.2 Underfloor and overhead air supply	3
1.3 IBM Rear Door Heat Exchanger	4
3.1 Modelling in 6SigmaRoom	8
3.2 Case A layout	10
3.3 Case B layout	11
3.4 CFD Residuals in 6SigmaRoom	13
4.1 CRAC units for case A	15
4.2 CRAC units for case B	16
4.3 Cost of infrastructure comparison	18
4.4 Cooling costs of different configurations	18

LIST OF TABLES

Table

Page

CHAPTER 1

INTRODUCTION

Data centers are centralized facilities that house networking and computing equipment for remote data storage, distribution, and processing. This equipment plays a vital role in running the daily tasks of any IT or government organization. Apart from this Data Centers are particularly important for cloud service providers for backing up huge amounts of user data and keeping critical applications up and running during nominal and peak demand periods.

1.1 Data Centers Energy Consumption

The average power consumption for cooling (electricity and related infrastructures) and IT equipment is 50 and 45 percent, respectively, according to ASHRAE TC 9.9 (2011) [1]. According to estimates [2] made in 2018, the total amount of energy used by data centers throughout the world increased to 205 terawatt hours, which is roughly equivalent to 1% of total global electricity consumption. Some of the largest data centers in the world can each contain many tens of thousands of IT devices, and they require more than 100 megawatts (MW) of power capacity. This amount of power is sufficient to provide electricity to approximately 80,000 homes in the United States (U.S. DOE 2020) [3]. The Power Usage Effectiveness (PUE) metric, which is one of the metrics that can be measured to assess how efficient a data center is [4], is the ratio of the amount of energy that is consumed by IT equipment to the overall amount of energy that is consumed in a data center.

1.2 Heat Generation and Cooling of Data Centers

The typical and expected power consumption per rack can be seen from Fig 1.1

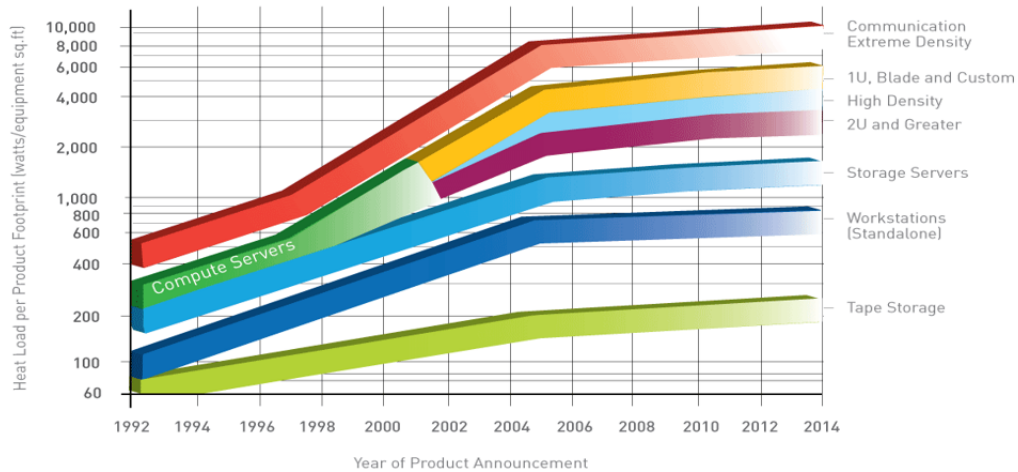


Figure 1.1: Trends in heat loads of the Data Center

Underfloor and overhead air supply designs, which are both often employed, are shown in Figs. 1.2 and 1.3. Both of these arrangements use a hot aisle–cold aisle layout to separate the supply of chilled air from the hot air. The front face of the rack, which serves as the equipment’s air intake, is positioned in front of perforated tiles. A hot aisle is formed by a rack facing the backside of another rack from which hot air exhausts. Fig. 1.2 depicts an example of an underfloor air supply configuration.

The chilled air is delivered into the area beneath the elevated floor by the computer room air conditioning unit (CRAC). Perforated tiles allow the chilly air to enter the space where it passes through the racks and becomes heated. The heated air then goes back to the CRAC inlet. In the overhead supply design (Fig. 1.3), diffusers above the room allow cooled air to enter. The hot air is then expelled from




Figure 1.2: Underfloor and overhead air supply

the room through wall vents after passing through racks. After passing through a heat exchanger, this hot air is finally delivered back as cold air through diffusers.

1.3 Rear Door Heat Exchangers

The significant temperature gradients at the intake suggest that air cooling for highly powered clusters is useless. In these circumstances, a hybrid cooling system with air cooling supported by a liquid to air heat exchanger can be effective. Using a water-cooled heat exchanger attached to the rack's back door is one method of reducing the impact of hot air recirculation in data centers. The heat exchanger greatly lowers the air temperature that is exhausted from the back of the rack while also removing a significant amount of heat from the rack. A plate fin and flat tube "radiator" design or a traditional fin and tube heat exchanger can be used. When receiving hot exhaust air for one coolant stream, the heat exchanger's inlet is aligned with the rear of the frame and has a planar geometry. The heat exchanger received cooled water for the cold coolant stream. Heat exchangers are connected to water plumbing headers at their inlet and outlet by flexible hose lines.



Figure 1.3: IBM Rear Door Heat Exchanger

1.4 Total Cost of Ownership

The total cost of ownership is a comprehensive tool that calculates the cost of each component, from the chip to the chiller, including all capital and operational expenses. The cost estimates include costs for planning, preparation, deployment, power, infrastructure, IT equipment, software, cooling, and maintenance. Numerous books have been devoted to simplifying cost calculations and developing methods for cost computation using the best practices for operating a data center.

The rack power density is increasing over time so that high-powered compute nodes can be utilized to meet the ever-increasing computational demand. Adoption of novel cooling methods to tackle cooling difficulties in order to preserve standard IT equipment reliability. Periodically, the internal and external infrastructure of the data center must be updated to fulfill cooling requirements. The proper selection of equipment, racks, rails, cable management, aisle containment, heat exchangers, Power Distribution Units (PDUs), software acquisition and security, network equipment, emergency power supply, and building management systems, among others, incurs additional expenses. Therefore, the costs of the goods are interdependent, and a change in one parameter impacts the others. Consequently, some assumptions and simplifications are made in order to apply a single TCO model to a data center of choice.

CHAPTER 2

Literature Survey

Conventional methods of air cooling have a number of important drawbacks when it comes to heat dissipation. These drawbacks include high energy consumption at the CRAC/CRAH units, humidity excursions, recirculation into the cold aisle, substantial site survey work, and so on [6]. In spite of the fact that there are optimization solutions for air conditioning provisioning, there are constraints due to the increased demand in high power demanding devices [7]. According to the data from the industry, the growth rate of servers around the world is 2.5 times 10^6 every year, and it is anticipated that this trend will continue [8]. In addition, the ASHRAE figure (Fig.) for the amount of heat load per 42U rack demonstrates that rack density is steadily growing. Because of this, it is challenging for traditional methods of air-cooling to provide increased capacity [9]. As a result, there is room for improvement in the techniques used to cool the air; one of these techniques is the retrofitting of existing datacenters with rear door heat exchangers (RDHx).

In most cases, the individual components that make up hybrid cooling systems are self-contained units. One example of this is the side car, which is an air-to-liquid heat exchanger that is attached to the side of a server cabinet. IBM has come out with an enclosed heat exchanger, which has the capability of removing up to 35 kW of power [10]. Another type of hybrid cooling system is exemplified by a fully enclosed server cabinet [11], which typically has a heat exchanger in the shape of a V installed on the bottom of the cabinet. A heat exchanger known as the RDHx (Rear Door Heat eXchanger) may be found at the back of the cabinet, directly in

front of where the hot exhaust air from the servers exits. It was demonstrated in [12] that utilizing an RDHx in various configurations of a data center can have a number of advantages. These advantages include 1) reducing the number of CRAH units that are required, 2) eliminating hot spots by removing heat closer to the source, and 3) allowing higher chilled water temperatures while still maintaining IT inlet temperatures within recommended ranges. [13] In a hybrid fashion, RDHXs were utilized in conjunction with air conditioning at the room level [14].

Mulay et.al [15] has established that utilizing a hybrid cooling technique can increase the cooling of a data center, which in turn may allow for the full population of racks as well as improved system infrastructure management.

Schmidt et al [16] studied water based heat exchangers and identified the effects of various failure modes of traditional air based cooling

CHAPTER 3

Methodogy

3.1 Hot Aisle Containment

A hot aisle containment system is one in which the hot aisle is contained and the remaining white space in the data center functions as a large cold aisle. Through roof panels and air ducts, the hot air is returned to the CRAC return.

Depending on the scale of the data center, a HAC system can be implemented in many ways. Hot air is returned to the CRAC unit, conditioned, and provided back as cold airflow for smaller data centers, as depicted in Fig. For bigger data centers, heated airflow can be channeled to a CRAH (Computer Room Air Handler) machine, which employs chilled water, cooling coils, and fans to condition hot air. This configuration can be used in conjunction with ASE to reduce cooling expenses in data centers. A further practical advantage of a HAC system is that, by confining the heat flow, the temperature of the working area is kept within OSHA (Occupational Safety and Health Standards) guidelines.

John Neimann et al. [20] analyzed the energy savings of both types of containment systems in a white paper. Their investigation revealed that the HAC system generated annual savings of \$40,000 in cooling expenses and a 13% decrease in annual PUE (Power Usage Effectiveness). Considering that the remainder of the data center is hot due to the CAC system, with temperatures reaching 38°C in large data centers, it can be hazardous to the safe functioning of unracked ITE. In light of these considerations, it is determined that a HAC is the better alternative for energy savings.

3.2 Data center layout




Figure 3.1: Modelling in 6SigmaRoom

There are numerous commercial CFD software that can accurately estimate the airflow patterns caused by mechanical devices and thermal sources. The majority of them, such as 6SigmaRoom , are executable on personal PCs. Using CFD methodologies, Future Facilities' 6SigmaRoom can estimate air flow and heat transport explicitly for Data Centers. In lieu of or in addition to the physical model, it provides a rapid method for estimating the indoor data center environment. It functions as an easy-to-use tool that enables the user to build a virtual facility model that may be used to evaluate and compare new designs, to model and troubleshoot existing facilities, and to provide a foundation for ongoing change management; this is known as predictive simulation. Applications comprising design, ventilation, and air conditioning system for all types and sizes of data centers, from low to high-density. The method closely

resembles a field model for estimating the turbulent convective air flow route within a data center.

As depicted in the figure, this program enables the user to create a virtual data center in full detail, granting access to a variety of architectural and cooling options, just as in a physical data center. Consequently, the airflow and thermal analysis conclude.

There are other commercial data center design products available on the market, but 6SigmaRoom's broad vendor libraries, which include ACU's, PDU's, ITE, fans, etc., simplify and expedite the data center design process.

3.2.1 case A

In this study, we consider two scenarios of common scale data centers. The first case the data center racks are of 16kW capacity but about 20% of the racks are running at 100% capacity and the rest of them are running at idle power (30% of total power). For the CRAC, there are 8 cooling units placed as shown in the layout. The cooling properties of the CU's are shown in the Fig, these can be adjusted in the 6SigmaRoom software

The dimensions of the data center are given in the table.

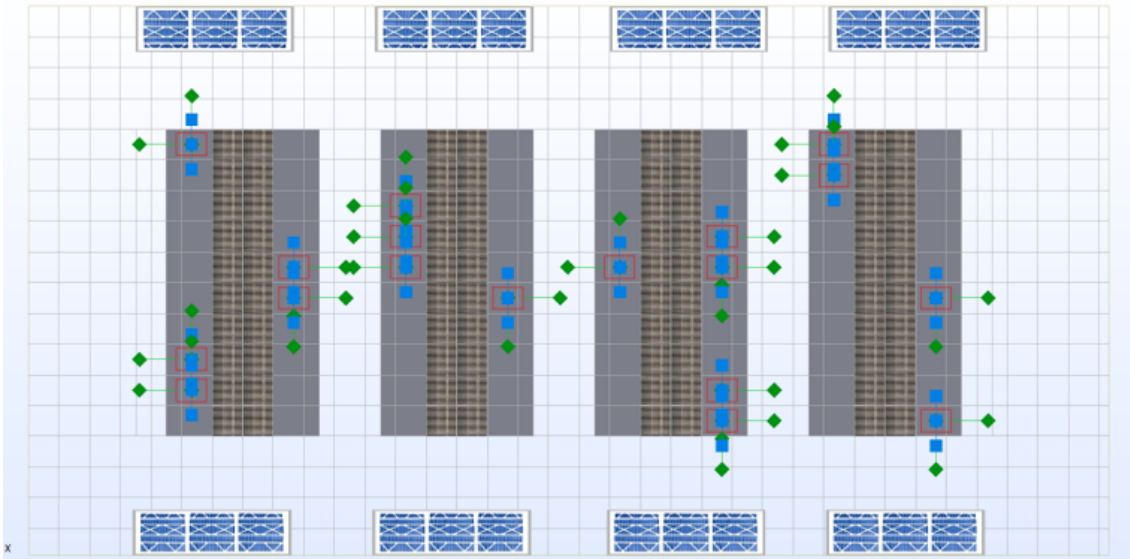


Figure 3.2: Case A layout

3.2.2 case B

In case B, instead of Random racks running at full power capacity, the layout includes rows of higher capacity rack . these are modelled to reflect a real life colocation or a third party data center.




Figure 3.3: Case B layout

3.3 CAD modelling

In this part, we will quickly outline the processes that were taken to create the CAD model depicted in Figure 1. The measurements of the space, which serve as the room's boundaries or walls, are sketched as part of a room layout. This compartment contains all of the technical space, including power supply, cooling equipment, and ITE.

2. Once the technical space has been defined, the raised floor option is selected from the object panel as depicted in Fig, and the needed raised floor height is given. Depending on the design requirements, supporting structures such as columns and beams can also be included during this phase. In the current model, a column was installed to support the artificial ceiling.

After specifying the flooring and room size, cabinets are arranged according to the design specifications. In this phase, the cabinet power limit and the servers to be stacked are also selected from the 6SigmaRoom vendor library, which contains actual equipment from specific manufacturers.

After arranging the cabinets and stacking them with servers, the room's cooling requirements can be assessed and the ACU units are arranged.

The final stage was separating the hot aisles. As shown in Fig., a false ceiling is selected from the architectural node in the object panel, and the height from the elevated floor is set. Following this, an aisle enclosure is constructed, and holes are cut in the roof panel and roof ducts. The next stage is to link the ACUs in the aisle to the false ceiling so that all hot airflow is returned to the ACU without mixing with cold airflow.

3.4 Simulation termination strategy

The CFD solver completes the specified number of iterations and terminates the solution depending on the default value (which is 1) or the value of the termination factor that has been assigned. The termination factor determines how much of allowable numerical error in the calculation required to terminate the simulation.

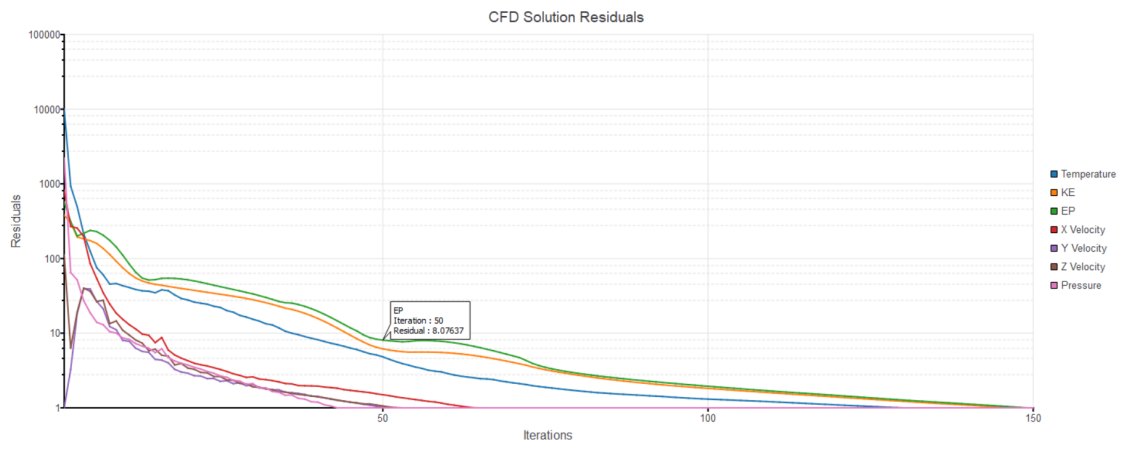


Figure 3.4: CFD Residuals in 6SigmaRoom

CHAPTER 4

Results and Discussion

The simulations were run until termination strategy was achieved. The temperature plots of the data center at a distance of 1m from the floor are plotted as shown in Fig

4.1 Total Cost of Ownership

The specified TCO model classifies the total cost into the categories listed below. The cost breakdown includes infrastructure, server, network, electricity, maintenance, and other expenses. In the publication, the individual cost is given in detail.

The infrastructure cost C_i is determined by the server and network power consumption, energy efficiency, rack occupancy, and land, building, and cooling infrastructure capital expenses.

$$C_i = \frac{C_{\text{sqm}} \times A_{\text{rack}} \times N_{\text{rack}} \times K_{\text{occupancy}}}{12 \times T_{\text{building}}} + \frac{(\text{SPUE} \times P_{\text{server}} + P_{\text{net}}) \times C_{cp,w}}{12 \times T_{\text{cooling}}} \quad (4.1)$$

The server cost consists of the server acquisition cost, which can either be a lumped server model or include all of the server's components. The model also includes the server's online rate to account for heat dissipation based on usage. In this research, the components also include the cost of heatsinks and cold plates for air and liquid cooling.




Figure 4.1: CRAC units for case A

$$C_p = \frac{C_{e-KWh} \times 30 \times 24}{1000} \times PUE \times (SPUE \times P_{\text{server}} + P_{\text{net}})$$

$$P_{\text{server}} = \begin{cases} P_{\text{server-peak}} = P_{ss-\text{peak}} \times U_{ss} \times P_{ss-\text{idle}} \times (1 - U_{ss}) \\ P_{\text{server-avg}} = P_{ss-\text{avg}} \times U_{ss} \times P_{ss-\text{idle}} \times (1 - U_{ss}) \end{cases} \quad (4.2)$$

The cost of acquiring a network comprises the core/cable and rack network equipment costs, as well as their respective power consumption. The cost is held constant for study and comparative purposes.

$$C_n = \frac{C_{\text{core-node}} \times N_{\text{server}}}{12 \times T_{\text{core}}} + \frac{C_{\text{netperrack}} \times N_{\text{rack}}}{12 \times T_{\text{net-rack}}} \quad (4.3)$$

The overall cost of electrical power consists of the cost per kilowatt-hour consumed by IT equipment, PDUs, power distribution losses, and cooling procedures. The server's power consumption and cooling can be tuned for air and liquid cool-




Figure 4.2: CRAC units for case B

ing, and efficiency metrics such as PUE and SPUE are factored into the total power consumption calculations.

$$\begin{aligned}
 C_m = & \left[\frac{C_{\text{server}} \times N_{\text{server}} \times ARR_{\text{server}}}{12} \right. \\
 & + (C_{\text{rep-cost}} \times ARR_{\text{server}} \\
 & \left. + \sum_i AFR_{\text{comp},i} \times N_{\text{comp},i} \times C_{\text{comp},i}) \right. \\
 & \left. \times (T_{\text{server}} - T_{\text{warranty}}) + C_{\text{labor}} \right] \times \frac{1}{12 \times T_{\text{server}}}
 \end{aligned} \tag{4.4}$$

As the maintenance cost is proportional to the failure rate of components, we hold it constant for the purpose of comparison in our analysis.

Priority should be given to the failure of IT equipment rather than the server as it is the heat that is affecting data center performance. The replacement cost, cold spare cost, and labor hourly rate are added to the failure rate in the maintenance cost calculation.

4.2 Assumptions

For purposes of comparison, the expenses of land acquisition, electricity distribution, and cooling infrastructure are held constant. The model accounts for an additional space (1.25 times) for cooling and power distribution in addition to the usual 42U unit rack dimensions. Infrastructure, cooling, and servers are assumed to have depreciation years of 15, 10, and 5 years, respectively. The cost of electricity per kilowatt-hour is based on an average price of \$0.1167 per kilowatt-hour. The peak power for the network equipment is provided as an assumption for all scenarios, and maximum utilization is specified (100). For the purpose of simplicity, the cost of networking equipment for all racks is the same in all described situations. For air cooling, the server acquisition cost is calculated using a simplified model, but for liquid cooling, it is computed using a model that accounts for the cost of each component in the liquid cooling loop. The warranty period for IT equipment is three years from the date of purchase. The number of cold spares and cost of replacement are computed based on the depreciation and SLA contract years, the failure rate of each component, and the labor cost per hour. For the analysis, the per-IT-equipment and maintenance expenses for CRAC, In-row Cooler (IRC), Overhead Cooling Unit (OCU), Bottom Cooling Unit (BCU), and Indirect Evaporative Cooling are held constant (IDEC). However, the figures are modified due to the installation of liquid-cooled IT equipment components. For air and liquid cooling, the server power usage and power consumption (peak and idle) are presented differently. Case studies feature a variety of PUE values.

The total cost of ownership is determined by summing the expenses for infrastructure, server, network, electricity, and maintenance.

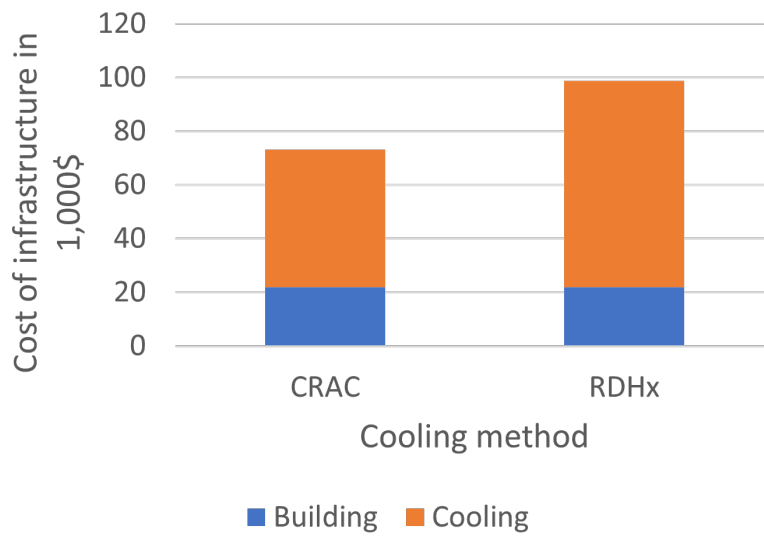


Figure 4.3: Cost of infrastructure comparison

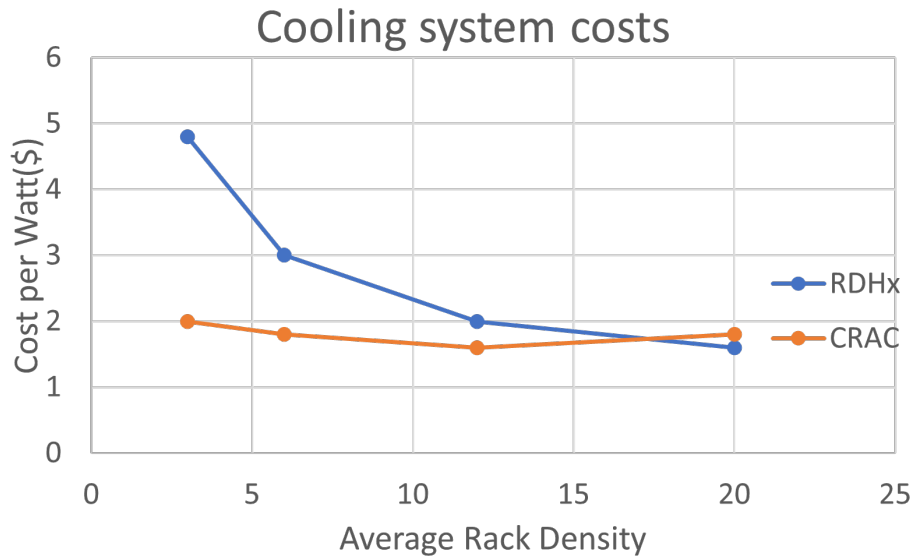


Figure 4.4: Cooling costs of different configurations

CHAPTER 5

Conclusion

Based on the thermal analysis carried out in the air-cooled data center under consideration, Rear Door Heat exchangers are more energy efficient in cooling data centers with 16kW racks.

Two real-life case scenarios were modeled and the simulation was carried out in 6SigmaRoom. In both cases, RDHxs cooled the data center without the need for additional containment. Temperature hotspots are eliminated due to localized cooling of IT equipment with a much less energy usage.

TCO analysis was carried out for both the cases and for both configurations. Although the initial capital expenditure for RDhx is high, it has been shown that the operational expenditure covers the high initial costs within a year. This is true for both cases, thus making RDH a cost and energy-efficient alternative to traditional CRAC cooling methods

REFERENCES

- [1] M. K. Patterson, “Liquid cooling guidelines,” in *Proceedings of the 2011 workshop on Energy Efficiency: HPC System and Datacenters - EE-HPC-WG '11*. ACM Press, 2011. [Online]. Available: <https://doi.org/10.1145%2F2159344.2159349>
- [2] E. Masanet, A. Shehabi, N. Lei, S. Smith, and J. Koomey, “Recalibrating global data center energy-use estimates,” *Science*, vol. 367, no. 6481, pp. 984–986, feb 2020. [Online]. Available: <https://doi.org/10.1126%2Fscience.aba3758>
- [3] A. Levinson, “How much energy do building energy codes really save? evidence from california,” Tech. Rep., dec 2014. [Online]. Available: <https://doi.org/10.3386%2Fw20797>
- [4] C. L. Belady and C. G. Malone, “Metrics and an infrastructure model to evaluate data center efficiency,” in *ASME 2007 InterPACK Conference, Volume 1*. ASMEDC, jan 2007. [Online]. Available: <https://doi.org/10.1115%2Fipack2007-33338>
- [5] A. Shehabi, S. Smith, D. Sartor, R. Brown, M. Herrlin, J. Koomey, E. Masanet, N. Horner, I. Azevedo, and W. Lintner, “United states data center energy usage report,” Tech. Rep., jun 2016. [Online]. Available: <https://doi.org/10.2172%2F1372902>
- [6] R. Schmidt and M. Iyengar, “Server rack rear door heat exchanger and the new ASHRAE recommended environmental guidelines,” in *ASME 2009 InterPACK Conference, Volume 2*. ASMEDC, jan 2009. [Online]. Available: <https://doi.org/10.1115%2Finterpack2009-89212>

- [7] V. Mulay, S. Karajgikar, D. Agonafer, R. Schmidt, M. Iyengar, and J. Nigen, “Computational study of hybrid cooling solution for thermal management of data centers,” in *ASME 2007 InterPACK Conference, Volume 1*. ASMEDC, jan 2007. [Online]. Available: <https://doi.org/10.1115%2Fipack2007-33000>
- [8] R. Udakeri, V. Mulay, and D. Agonafer, “Comparison of overhead supply and underfloor supply with rear heat exchanger in high density data center clusters,” in *2008 Twenty-fourth Annual IEEE Semiconductor Thermal Measurement and Management Symposium*. IEEE, mar 2008. [Online]. Available: <https://doi.org/10.1109%2Fstherm.2008.4509385>
- [9] K. Nemati, T. Gao, B. T. Murray, and B. Sammakia, “Experimental characterization of the rear door fans and heat exchanger of a fully-enclosed, hybrid-cooled server cabinet,” in *2015 31st Thermal Measurement, Modeling & Management Symposium (SEMI-THERM)*. IEEE, mar 2015. [Online]. Available: <https://doi.org/10.1109%2Fsemi-therm.2015.7100154>
- [10] K. Nemati, H. A. Alissa, B. T. Murray, K. Schneebeli, and B. Sammakia, “Experimental failure analysis of a rear door heat exchanger with localized containment,” *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 7, no. 6, pp. 882–892, jun 2017. [Online]. Available: <https://doi.org/10.1109%2Ftcpmt.2017.2682863>
- [11] V. S. Simon, H. Modi, K. B. Sivaraju, P. Bansode, S. Saini, P. Shahi, S. Karajgikar, V. Mulay, and D. Agonafer, “Feasibility study of rear door heat exchanger for a high capacity data center,” in *ASME 2022 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*. American Society of Mechanical Engineers, oct 2022. [Online]. Available: <https://doi.org/10.1115%2Fipack2022-97494>