CHALMERS
UNIVERSITY OF TECHNOLOGY

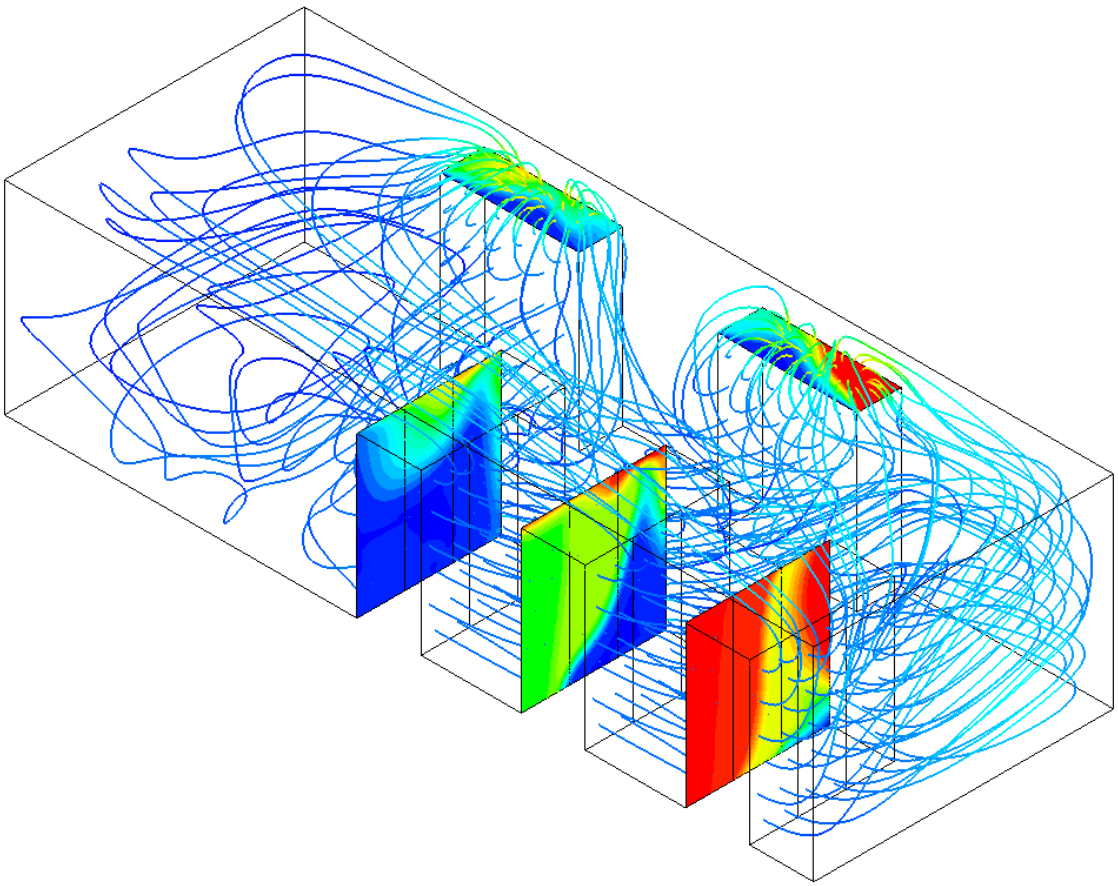

CFD Modeling of an Air-Cooled Data Center

Master's thesis in Applied Mechanics

EMELIE WIBRON

Department of Applied Mechanics
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2015

MASTER'S THESIS IN APPLIED MECHANICS

CFD Modeling of an Air-Cooled Data Center

EMELIE WIBRON

Department of Applied Mechanics
Division of Fluid Mechanics
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2015

CFD Modeling of an Air-Cooled Data Center
EMELIE WIBRON

Master's Thesis 2015:31
ISSN 1652-8557
Department of Applied Mechanics
Division of Fluid Mechanics
Chalmers University of Technology
SE-412 96 Gothenburg
Sweden
Telephone +46 (0)31-772 1000

Cover: Illustration of the temperature and flow field of the original setup in the server room at Fortlax data center

Chalmers Reproservice
Gothenburg, Sweden 2015

CFD Modeling of an Air-Cooled Data Center
Master's thesis in Applied Mechanics
EMELIE WIBRON
Department of Applied Mechanics
Division of Fluid Mechanics
Chalmers University of Technology

Abstract

Fortlax is an air-cooled data center in Norrbotten. The potential of future investments in data centers is prosperous, but sustainability is an increasingly important factor. It is important to make sure that the data centers are sufficiently cooled while too much forced cooling leads to economical losses and a waste of energy. The purpose of this thesis is to develop a CFD model of a server room at Fortlax data center that accurately predicts the temperature and flow field. The goal is to develop the CFD model in order to evaluate different cooling systems and also improve the placement of the server racks. Most of the server racks are only half full at the moment, but the heat load is supposed to increase in the near future. Therefore, it has been assumed that all the server racks generate full heat load. The commercial CFD software ANSYS CFX 16.0 has been used to perform the simulations. The current cooling system with the current placement of the server racks resulted in some insufficiently cooled server racks. The performance of the cooling system was significantly improved when the hard-floor configuration was replaced by a raised-floor configuration. The design strategy based on hot/cold-aisles has become the standard when raised-floor configurations are used. However, the raised-floor configuration with parallel rows of server racks turned out to be the setup that performed best based on the performance metrics that were used to evaluate the results in this thesis.

Keywords: CFD, ANSYS CFX, air-cooled, data center, hard-floor configuration, raised-floor configuration, perforated tiles, hot/cold-aisles

Contents

Preface

This thesis work has been carried out from January 2015 to June 2015 at the Division of Fluid Mechanics at Luleå Univeristy of Technology in collaboration with Fortlax. I am very greatful to Professor Staffan Lundström for giving me the opportunity to do this thesis work. Many thanks to my supervisor Ph.D. Anna-Lena Ljung for help, support and expertise. Thanks to everybody at the Division of Fluid Mechanics for making me feel welcome. I would also like to express my appreciation to Professor Sinisa Krajnovic at the Division of Fluid Mechanics at Chalmers University of Technology for being willing to take on the role as examiner.

Gothenburg June 2015

EMELIE WIBRON

Notations

Abbreviations

CFD	Computational Fluid Dynamics
CI	Capture Index
CRAC	Computer Room Air Condition
RANS	Reynolds-Averaged Navier-Stokes
RCI	Rack Cooling Index
RTI	Return Temperature Index

Roman upper case letters

A_{tile}	Area of a perforated tile
Ar	Archimedes number
E	Energy
H	Heat load
L	Characteristic length scale
Q	Volumetric flow rate
R_0	Universal gas constant
Re	Reynolds number
S	Source term
T	Temperature
U	Characteristic velocity scale
V	Volume

Roman lower case letters

c_p	Specific heat capacity
g	Gravitational acceleration
h	Enthalpy
k	Turbulent kinetic energy
$\dot{m}$	Mass flow rate
p	Pressure
$\overline{p}$	Steady pressure component
p'	Fluctuating pressure component
p_{abs}	Absolute pressure
t	Time
$\mathbf{u}$	Velocity vector
u, v, w	Velocity components
u_i	Velocity components in tensor notation
$\overline{u}$, $\overline{v}$, $\overline{w}$	Steady velocity components
u', v', w'	Fluctuating velocity components
u^+	Dimensionless velocity
u_τ	Friction velocity
w	Molecular weight
x, y, z	Cartesian coordinates
x_i	Cartesian coordinates in tensor notation
y^+	Dimensionless wall distance

Greek upper case letters

Γ	Diffusion coefficient

Greek lower case letters

β	Volumetric thermal expansion coefficient
δ_{ij}	Kronecker delta
ε	Dissipation of turbulent kinetic energy
λ	Thermal conductivity
μ	Dynamic viscosity
μ_t	Turbulent dynamic viscosity
ν	Kinematic viscosity
ϕ	Additional variable
ρ	Density
ρ_{ref}	Buoyancy reference density
σ	Percentage open area
τ_{ij}	Viscous stress components
τ_w	Wall shear stress

1 Introduction

1.1 Background

Fortlax is an air-cooled data center in Norrbotten that offers secure and effective solutions for storage and backup of digital information. The location offers advantages such as cold weather, reliable supply of power and a stable political climate[1]. The establishment of data centers such as Fortlax is one of the most important growth factors in the IT business. One of the reasons for the growth is that an increasing number of companies and organisations have started to outsource their data storage. Although the potential of future investments in data centers is prosperous, sustainability is an increasingly important factor. In 2010, the total electricity used by data centers was about 1.3% of all the electricity use in the world[2]. It is therefore of highest importance to consider the energy efficiency of data centers in general.

The main components of an air-cooled data center are Computer Room Air Conditioner (CRAC) units and server racks. The server racks dissipate heat and need to be cooled in order to make sure that the electronics operate in the temperature range recommended by the manufacturer. Otherwise there is a risk of overheating, resulting in malfunction or shut down to prevent hardware damages. This interruption is costly for business and needs to be prevented. It is therefore important to make sure that the data centers are sufficiently cooled while too much forced cooling leads to economical losses and a waste of energy. The CRAC units supply cold air into the data center. The cold air is supposed to enter through the front of the server racks and hot air will exit through the back. Depending on the distribution of CRAC units and server racks in the data center there is a risk that the cold air does not necessarily reach all the server racks to the desired extent. The cooling of data centers is therefore crucially dependent on the flow field of the air. Computational Fluid Dynamics (CFD) is an excellent tool to provide detailed information about the temperature and flow field in a data center. CFD modeling can be used to analyse both existing configurations and proposed configurations before they are built[3].

1.2 Purpose and goal

A server room at Fortlax data center will be considered in this thesis. Most of the server racks are only half full at the moment, but the heat load is supposed to increase in the near future. Therefore, it is assumed that all the server racks generate full heat load. The purpose of the thesis is to develop a CFD model of the server room at Fortlax data center that accurately predicts the temperature and flow field. To begin with, the original setup with the current cooling system and the current placement of the main components will be simulated. The CRAC units are fixed, but the server racks can be moved. The goal is to develop the CFD model in order to evaluate different cooling systems and also improve the placement of the server racks. Fortlax is currently planning to build another data center and the comparison of different cooling systems might provide interesting input. The commercial CFD software ANSYS CFX 16.0 will be used to perform the simulations.

1.3 Thesis outline

The thesis begins with an introduction in Chapter 1. It contains the background of the thesis and defines the purpose and the goal. Chapter 2 is a literature review of previous work within CFD modeling of air-cooled data centers. Chapter 3 contains basic theory about fluid mechanics in general and CFD modeling in particular. Chapter 4 describes the method. Geometry, mesh and boundary conditions are presented for the original setup as well as for the other setups that will be simulated. Which performance metrics that will be used to evaluate the results are also described. Chapter 5 presents the mesh convergence study and the results of the simulations. Finally, Chapter 6 contains discussion and conclusions.

2 Literature review

The cooling system in an air-cooled data center might be based on a raised-floor or a hard-floor configuration. When a raised-floor configuration is used, the CRAC units supply cold air into an under-floor space and the cold air enters the room through perforated tiles in the floor. The perforated tiles are removable and can be replaced by solid tiles which makes the configuration flexible. Strategies for the design in data centers are often based on hot/cold-aisles where the server racks are placed into a series of rows. Cold-aisles are formed between the front sides of two rows of server racks and hot-aisles are formed on the other sides. This design strategy has become the standard when raised-floor configurations are used. Perforated tiles are placed in the cold-aisles and solid tiles are placed in the hot-aisles. An example of a cold-aisle with perforated tiles is illustrated in Figure 2.1. The purpose of the design strategy is to prevent hot air exhausted by the back of a server rack to enter the front of another server rack[4].

A comparison of raised-floor and hard-floor configurations in a data center has been made. Room and ceiling return strategies were compared for both the cooling systems by CFD modeling. A ceiling return strategy means that a ceiling void space is placed directly above the hot-aisles. The purpose is to prevent the hot air from mixing with the cold air before returning to the CRAC units. The room return strategy does not isolate the hot air on its way back to the CRAC units, but it is still possible to use panels to cover the cold-aisles in order to minimise hot air recirculation. It was found that a ceiling return strategy for the return of hot air to the CRAC units was preferable for both raised-floor and hard-floor configurations. The raised-floor configuration performed better than the hard-floor configuration when the same return strategy was used for both the cooling systems[5].

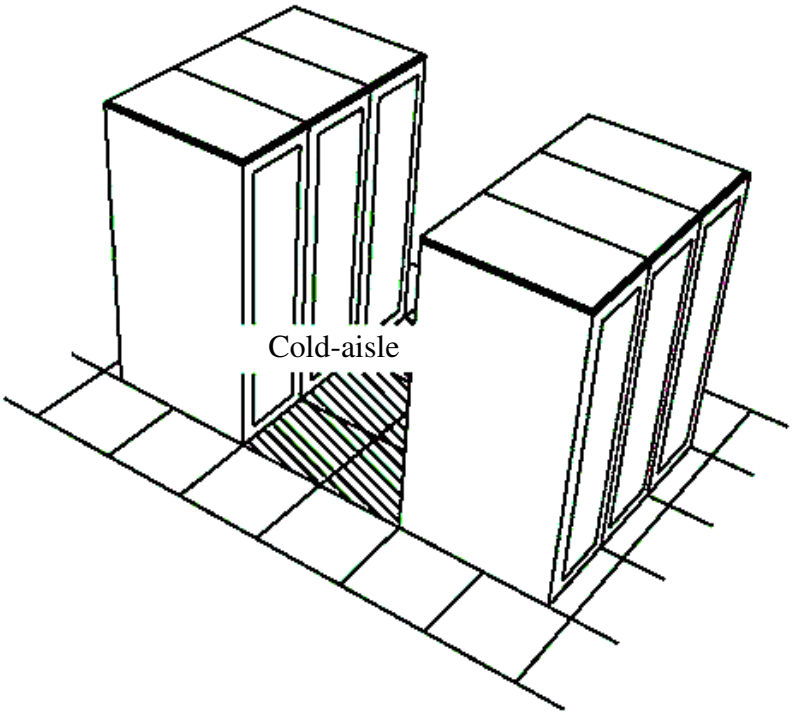

Figure 2.1: A cold-aisle with perforated tiles between two rows of server racks[6].

The flow through the server racks is driven by internal fans with a constant volumetric flow rate. The flow rate supplied by the CRAC units should be at least equal to the flow rate demanded by the server racks. If the cold air available to the server racks is insufficient, warm air will enter through the front of the server racks instead. When a raised-floor configuration is used and the cold air is supplied through perforated tiles, it is usually the top of the server racks that might not be sufficiently cooled. When a hot/cold-aisle configuration is used, hot air might be recirculated over the server racks or around them at the end of each row. This is what is called hot air recirculation and should be avoided if possible[3]. The heat load can vary significantly across the server racks in a data center. Therefore, just meeting the total flow rate requirement is not enough. Attention must also be paid to the distribution of the cold air in order to make sure that all server racks are sufficiently cooled[7]. Recommended and maximum allowable air temperatures at the front of the server racks are specified by ASHRAE's thermal guidelines. The recommended temperature range is $18 - 27°C$ and the allowable temperature range is $15 - 32°C$[8].

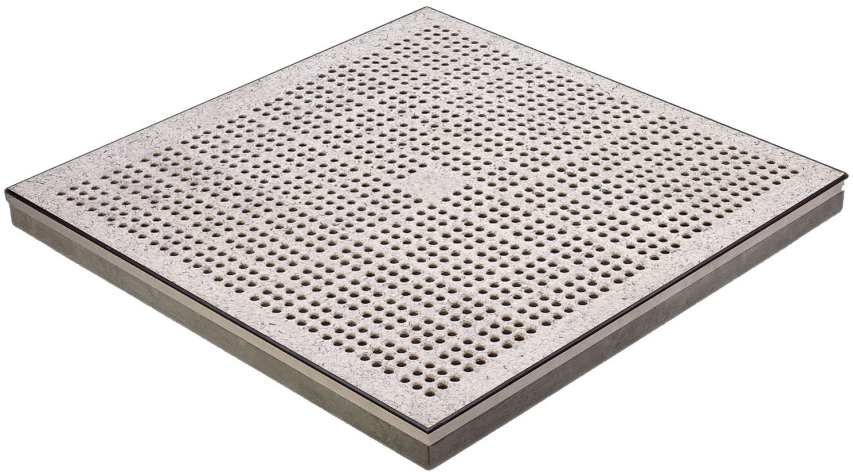

Figure 2.2: A perforated tile that can be used in a raised-floor configuration[9].

The perforated tiles in a raised-floor configuration have a large number of small openings which require a very fine mesh to resolve the flow when using CFD modeling. An example of a perforated tile is illustrated in Figure 2.2. Due to mesh size limitations, the perforated tiles are often modeled as fully open tiles with constant velocity. The reasoning behind this simplification is that since the openings in the perforated tiles are very small, the small jets will quickly merge into a single large jet. It has been concluded that the mixing level between the cold air entering a server room in a data center through the perforated tiles and the hot air exiting the server racks is low in this case. It results in hot and cold spots predicted by the CFD model that do not exist in reality. The deviations from measurements are largest in the area around the front and back of the server racks and can be as much as 5°C. Incorrect modeling of the flows that exits from perforated tiles and server racks, inadequacy of the turbulence model, boundary conditions on the floor, walls and ceiling that do not account for leakage, neglecting buoyancy and coarseness of the mesh have been identified as potential error sources[10].

These potential error sources have been further investigated to various extent. It has been further investigated if the simplified representation of the server racks in the CFD model is likely to cause the deviations from measurements. The most simple approach is to model a server rack as a "black box" in which the flow inside is not resolved. The flow enters the front of the server rack and then reappears at the back of the server rack with the appropriate increase of temperature included. This server rack model has been compared to more detailed models in order to investigate the extent to which the temperature and flow field in a data center is affected by the server rack model. It was found that there are no significant differences in the results obtained by the different server rack models. The pronounced hot and cold spots were still present. It indicates that the simplified representation of the server racks in the CFD model is not likely to cause the deviations from measurements[11].

The $k-\varepsilon$ model is well established as a suitable turbulence model to use in CFD modeling of data centers[11]. Inadequacy of the turbulence model has not been much furter investigated. A potential error source that has been further investigated and might have contributed to the hot and cold spots predicted by the CFD model is neglecting buoyancy. It has been showed that buoyancy often should be included in the CFD model and how the importance of including buoyancy can be evaluated[14]. Most of the CFD modeling of data centers has been made on relatively coarse meshes. The effect of mesh size on CFD modeling of data centers has been investigated with respect to solution time and relative accuracy. Hexahedral meshes were used and the distribution of the cells was made as uniform as possible for each cell size. It was found that accuracy improved significantly as the average cell size was reduced from 21 inches to about six inches. The results were not mesh independent even for a mesh with cell size as small as one inch, but only marginal improvements were found on accuracy when the cell size was reduced from six inches to smaller cell sizes. However, the solution time increased a lot more than the relative accuracy[12].

The low mixing level might be a result of the low momentum fluxes of the jets that enter the data center through the perforated tiles when they are modeled as fully open. The velocity is set to $\mathrm{Q}/\mathrm{A}_{\mathrm{tile}}$ in order to obtain the correct mass flow rate and the resulting momentum flux is equal to $\rho\mathrm{Q}\,(\mathrm{Q}/\mathrm{A}_{\mathrm{tile}})$. The velocity is lower than the correct jet velocity through the small openings in the perforated tiles. An alternative approach of modeling a perforated tile is to use a momentum source to correct for the momentum deficit. The flow area is still considered to be the full area of the perforated tile. The correct jet velocity through the small openings in the perforated tile is equal to $\mathrm{Q}/\sigma\mathrm{A}_{\mathrm{tile}}$ where σ is the percentage open area and the resulting momentum flux is equal to $\rho Q(\mathrm{Q}/\sigma\mathrm{A}_{\mathrm{tile}})$. Therefore, the correct momentum flux is not obtained if the perforated tile has a percentage open area less than 100%. The percentage open area is what characterises a perforated tile and 25% is the most common value. When a momentum source is used, a body force term is added in the momentum equation in the vertical direction to correct for the momentum deficit. It has been concluded that a fully open tile model is adequate for perforated tiles with a percentage open area of 50% or larger. Adding a momentum source to correct for the momentum deficit is sufficiently accurate to match the velocity profile for perforated tiles with lower perforation. However, the temperature field closest to the perforated tiles might not be accurately predicted when the flow through the small openings is not resolved[13].

CFD modeling can provide detailed information about the temperature and flow field in data centers and it has a capacity to offer a huge amount of data. There are many different performance metrics that can be used when analyzing this data. The Rack Cooling Index (RCI) focuses on the temperature at the front of the server racks. It evaluates how many of the server racks that have temperatures higher than the maximum recommended temperature at the front and how much the limit is exceeded for each server rack. The RCI can be improved by supplying air at an increased flow rate or at a decreased temperature. Both these alternatives result in an increased energy consumption. The Return Temperature Index (RTI) measures the energy performance of the cooling system. The optimal situation is that the difference between the temperature of the air returned to and supplied from a CRAC unit is equal to the temperature rise across a server rack. Deviations from this indicates a declining energy performance. A combination of RCI and RTI is therefore useful as performance metrics[15]. The Capture Index (CI) is a performance metric that focuses on the flow field. It evaluates how much of the air that enters the front of a server rack that originates from the CRAC units or the perforated tiles and not from the back of another server rack[16].

3 Theory

3.1 Governing equations

The governing equations of fluid mechanics are the conservation laws for mass, momentum and energy. These equations can be stated in differential or integral form by applying the conservation laws to a point or an extended region of the fluid, respectively. The differential form of the governing equations will be described here. The point of the fluid is considered to be an infinitely small volume element, but not so small that individual molecules influence macroscopic properties. The fluid is considered to be a continuum and all fluid properties are functions of space and time.

The fact that the mass of a fluid is conserved gives the continuity equation for a compressible fluid

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0 \tag{3.1}$$

The first term represents the rate of change of mass and the second term represents the net rate of flow of mass out of the volume.

Both surface and body forces are acting on the fluid. The fact that the rate of change of momentum is equal to the sum of the forces gives the momentum equations

$$\frac{\partial (\rho u)}{\partial t} + \nabla \cdot (\rho u \mathbf{u}) = \frac{\partial (-p + \tau_{xx})}{\partial x} + \frac{\partial \tau_{yx}}{\partial y} + \frac{\partial \tau_{zx}}{\partial z} + S_{Mx} \tag{3.2}$$

$$\frac{\partial (\rho v)}{\partial t} + \nabla \cdot (\rho v \mathbf{u}) = \frac{\partial \tau_{xy}}{\partial x} + \frac{\partial (-p + \tau_{yy})}{\partial y} + \frac{\partial \tau_{zy}}{\partial z} + S_{My} \tag{3.3}$$

$$\frac{\partial (\rho w)}{\partial t} + \nabla \cdot (\rho w \mathbf{u}) = \frac{\partial \tau_{xz}}{\partial x} + \frac{\partial \tau_{yz}}{\partial y} + \frac{\partial (-p + \tau_{zz})}{\partial z} + S_{Mz} \tag{3.4}$$

The terms on the left hand side represent the rate of change of momentum and the net rate of flow of momentum out of the volume, respectively. The terms on the right hand side represent surface stresses and the rate of change of momentum due to sources. Both pressure, which is a normal stress, and viscous stresses are included in the terms that represent surface stresses.

The fact that the rate of change of energy is equal to the sum of the rate of heat addition to and the work done on the fluid gives the energy equation

$$\begin{aligned} \frac{\partial (\rho E)}{\partial t} + \nabla \cdot (\rho E \mathbf{u}) = &- \nabla \cdot (p \mathbf{u}) + \left[\frac{\partial (u \tau_{xx})}{\partial x} + \frac{\partial (u \tau_{yx})}{\partial y} + \frac{\partial (u \tau_{zx})}{\partial z} + \frac{\partial (v \tau_{xy})}{\partial x} \right. \\ &\left. + \frac{\partial (v \tau_{yy})}{\partial y} + \frac{\partial (v \tau_{yz})}{\partial z} + \frac{\partial (w \tau_{xz})}{\partial x} + \frac{\partial (w \tau_{yz})}{\partial y} + \frac{\partial (w \tau_{zz})}{\partial z} \right] \\ &+ \nabla \cdot (\lambda \nabla T) + S_E \end{aligned} \tag{3.5}$$

The terms on the left hand side represent the rate of change of energy and the net rate of flow of energy out of the volume, respectively. The terms on the first two rows on the right hand side represent the total rate of work done by surface stresses. The last two terms represent the rate of heat addition due to heat conduction across the boundaries of the volume and the rate of change of energy due to sources[18].

The in total five governing equations include seven unknown variables. To close the system of equations, constitutive equations of state for density and enthalpy are needed. The general form of these equations read

$$\rho = \rho(p,T) \quad dh = \left.\frac{\partial h}{\partial T}\right|_p dT + \left.\frac{\partial h}{\partial p}\right|_T dp = c_p dT + \left.\frac{\partial h}{\partial p}\right|_T dp \quad c_p = c_p(p,T) \tag{3.6}$$

For an ideal gas, the ideal gas law is used to calculate density and the specific heat capacity can be at most a function of temperature. The constitutive equations for an ideal gas read

$$\rho = \frac{wp_{abs}}{R_0 T} \quad dh = c_p dT \quad c_p = c_p(T) \tag{3.7}$$

This is one of a range of various special cases of the constitutive equations for particular material types that can be chosen in ANSYS CFX.

Transport equations for additional variables can also be defined. Additional variables are scalar components that are transported through the flow. An additional variable is typically interpreted as a concentration in ANSYS CFX. A volumetric additional variable is defined as

$$\phi = \frac{\text{conserved quantity}}{\text{volume of fluid}} \tag{3.8}$$

and if the conserved quantity has unit of mass, the additional variable will have the same unit as concentration. The transport equation for an additional variable reads

$$\frac{\partial(\rho\phi)}{\partial t} + \nabla\cdot(\rho\phi\mathbf{u}) = \nabla\cdot(\Gamma\nabla\phi) + S_\phi \tag{3.9}$$

The terms on the left hand side represent the rate of change of the additional variable and the net rate of flow of the additional variable out of the volume, respectively. The terms on the right hand side represent the rate of change of the additional variable due to diffusion and the rate of change of the additional variable due to sources. For most applications, the transport of additional variables is a both convective and diffusive process which requires that a value of the diffusion coefficient is specified. But there are applications where the transport of the additional variable is a purely convective process and diffusion effects can then be neglected.

Various source terms can be included in ANSYS CFX. Momentum sources are implemented as forces per unit volume acting on the fluid. They can only be specified on subdomains, not on boundaries and points. Subdomains are regions contained in the computational domain. General momentum sources can be defined if the predefined momentum source models are insufficient. Different sources can be specified in different directions according to

$$S_{Mx} = S_{\text{general},x}\mathbf{i} \quad S_{My} = S_{\text{general},y}\mathbf{j} \quad S_{Mz} = S_{\text{general},z}\mathbf{k} \tag{3.10}$$

Buoyancy source terms can be included for flows in which gravity is important. Buoyancy is driven by variations in density. When buoyancy is included, the following source term is added to the momentum equation in the vertical direction

$$S_{\text{buoyancy},y} = (\rho - \rho_{\text{ref}})g \tag{3.11}$$

The full buoyancy model evaluates the density difference directly and is the only option when a fluid with density as a function of temperature is considered in ANSYS CFX[17].

3.2 Turbulence modeling

The Reynolds number is a measure of the relative importance of inertia forces and viscous forces. It is expressed as

$$Re = \frac{\rho U L}{\mu} \tag{3.12}$$

At values of Reynolds number below some critical value, the flow is smooth and steady. This regime is called laminar flow. At values of Reynolds number above the same critical value, the flow behaviour changes drastically. This regime is called turbulent flow. Turbulent flows are characterised by random fluctuations and chaotic motion. Most industrially relevant flows are turbulent. However, most often the time-averaged properties are of greatest interest. The governing equations of the steady mean flow are called the Reynolds-Averaged Navier-Stokes (RANS) equations. They are obtained by introducing the Reynolds decomposition where the flow variables are decomposed into a steady and a fluctuating component according to

$$u(t) = \overline{u} + u'(t) \quad v(t) = \overline{v} + v'(t) \quad w(t) = \overline{w} + w'(t) \quad p(t) = \overline{p} + p'(t) \tag{3.13}$$

The RANS equations are obtained by substituting the above relations into the momentum equations and taking the time average. They read

$$\begin{aligned} \frac{\partial \rho \overline{u}}{\partial t} + \nabla \cdot (\rho \overline{u} \overline{\mathbf{u}}) &= \frac{\partial}{\partial x}(-\overline{p} + \tau_{xx} - \rho \overline{u'^2}) \\ &+ \frac{\partial}{\partial y}(\tau_{xy} - \rho \overline{u'v'}) + \frac{\partial}{\partial z}(\tau_{xz} - \rho \overline{u'w'}) + S_{Mx} \end{aligned} \tag{3.14}$$

$$\begin{aligned} \frac{\partial \rho \overline{v}}{\partial t} + \nabla \cdot (\rho \overline{v} \overline{\mathbf{u}}) &= \frac{\partial}{\partial x}(\tau_{yx} - \rho \overline{u'v'}) \\ &+ \frac{\partial}{\partial y}(-\overline{p} + \tau_{yy} - \rho \overline{v'^2}) + \frac{\partial}{\partial z}(\tau_{yz} - \rho \overline{v'w'}) + S_{My} \end{aligned} \tag{3.15}$$

$$\begin{aligned} \frac{\partial \rho \overline{w}}{\partial t} + \nabla \cdot (\rho \overline{w} \overline{\mathbf{u}}) &= \frac{\partial}{\partial x}(\tau_{zx} - \rho \overline{u'w'}) \\ &+ \frac{\partial}{\partial y}(\tau_{zy} - \rho \overline{v'w'}) + \frac{\partial}{\partial z}(-\overline{p} + \tau_{zz} - \rho \overline{w'^2}) + S_{My} \end{aligned} \tag{3.16}$$

The RANS equations include additional stress terms called Reynolds stresses. A turbulence model is needed to account for the Reynolds stresses and close the system of equations.

Turbulence models based on the RANS equations focus on the mean flow and what effect turbulence has on the mean flow properties. The $k-\varepsilon$ model is a turbulence model that solves additional transport equations for the two quantities turbulent kinetic energy and rate of dissipation of turbulent kinetic energy. It is one of the most widely used turbulence models because of its excellent performance for many industrially relevent flows and it is also the most widely validated turbulence model. However, there are a number of important cases when the $k-\varepsilon$ model performes very bad. The $k-\varepsilon$ model assumes isotropic turbulence and is based on the Boussinesq approximation which assumes that the Reynolds stresses are proportional to the mean velocity gradients according to

$$-\rho \overline{u_i' u_j'} = \mu_t \left(\frac{\partial \overline{u}_i}{\partial x_j} + \frac{\partial \overline{u}_j}{\partial x_i} \right) - \frac{2}{3} \rho k \delta_{ij} \tag{3.17}$$

The additional transport equations read

$$\frac{\partial(\rho k)}{\partial t} + \nabla \cdot (\rho k \mathbf{u}) = \nabla \cdot \left[\frac{\mu_t}{\sigma_k} \nabla k \right] + \mu_t \left(\frac{\partial \overline{u}_i}{\partial x_j} + \frac{\partial \overline{u}_j}{\partial x_i} \right)^2 - \rho \varepsilon \tag{3.18}$$

$$\frac{\partial(\rho \varepsilon)}{\partial t} + \nabla \cdot (\rho \varepsilon \mathbf{u}) = \nabla \cdot \left[\frac{\mu_t}{\sigma_k} \nabla \varepsilon \right] + C_{1\varepsilon} \frac{\varepsilon}{k} \mu_t \left(\frac{\partial \overline{u}_i}{\partial x_j} + \frac{\partial \overline{u}_j}{\partial x_i} \right)^2 - C_{2\varepsilon} \rho \frac{\varepsilon^2}{k} \tag{3.19}$$

The terms on the left hand side represent rate of change of turbulent kinetic energy or dissipation of turbulent kinetic energy and transport by convection. The terms on the right hand side represent transport by diffusion, rate of production and rate of destruction. The constants read

$$C_\mu = 0.09 \quad \sigma_k = 1.00 \quad \sigma_\varepsilon = 1.30 \quad C_{1\varepsilon} = 1.44 \quad C_{2\varepsilon} = 1.92 \tag{3.20}$$

and have been determined by experimental data fitting for a wide range of flows[18].

3.3 Near wall modeling

The velocity approaches zero close to a solid wall. Consider the Reynolds number based on the wall distance according to

$$Re_y = \frac{\rho U y}{\mu} \tag{3.21}$$

Far away from the wall, inertia forces dominate over viscous forces. But when the wall distance is decreased, the Reynolds number will decrease as well. Before the wall distance reaches zero there will be a range of values where the Reynolds number is of order one. From here and closer to the wall, the magnitude of viscous forces will be equal to or larger than the magnitude of inertia forces. This region is called the linear sub-layer because of the linear relationship between the dimensionless expressions for velocity and wall distance

$$u^+ = \frac{\overline{u}}{u_\tau} = \frac{u_\tau y}{\nu} = y^+ \tag{3.22}$$

where the friction velocity is defined as $u_\tau = (\tau_w/\rho)^{1/2}$. This viscous sublayer is extremely thin, it is only valid for $y^+ < 5$. Outside the viscous sublayer, there is a region where both inertia forces and viscous forces are important. It is valid for $30 < y^+ < 500$. This region is called the log-law layer because of the logarithmic relationship between the dimensionless expressions for velocity and wall distance

$$u^+ = \frac{1}{\kappa} \ln(y^+) + B = \frac{1}{\kappa} \ln(E y^+) \tag{3.23}$$

where von Karman's constant is equal to $\kappa \approx 0.4$ and the additive constant is equal to $B \approx 5.5$ (or $E \approx 9.8$) for all turbulent flows past smooth walls at high Reynolds number[18].

The $k-\varepsilon$ model includes a wall function approach that can be used when there is no interest to resolve the entire boundary layer. When using the wall function approach, the log-law is used to represent the boundary layer. The near-wall grid points should be located at a dimensionless wall distance within the range where the log-law layer is valid. ANSYS CFX allows one to choose between standard and scalable wall functions. The scalable wall functions ignore grid points below the lower limit and are therefore preferable[17].

3.4 Discretisation method

The governing equations of fluid mechanics do seldom have an analytic solution. Therefore, CFD modeling is used to solve them. The differential equations are approximated as a system of algebraic equations by using a discretisation method. The most common discretisation methods are the finite difference method, the finite element method and the finite volume method. ANSYS CFX uses an element-based finite volume method. The computational domain is divided into discrete control volumes used to conserve quantities such as mass, momentum and energy. The control volumes can have any kind of shape in two or three dimensions and are represented by the mesh. Hexahedrons and tetrahedrons are common shapes of the control volumes when a three dimensional computational domain is considered. The governing equations are integrated over the control volumes in order to give discretised equations at the nodes which are located at the center of the control volumes. The resulting system of algebraic equations is then solved. Numerical values of all fluid properties and solution variables are stored at the nodes[17].

3.5 Errors and uncertainties

It should be noted that CFD modeling always includes errors and uncertainties to some extent. Errors are defined as recognisable deficiencies in the CFD model that are not caused by lack of knowledge. There are three recognised sources of numerical errors. Roundoff errors are a result of the fact that the computer represents real numbers with a finite number of significant digits. In general, roundoff errors can be controlled. Since CFD modeling is an iterative process, iterative convergence errors are the difference between the numerical and the exact solution. Limited resources in form of time and computing power imply that the solution is considered to be converged when it is sufficiently close to the exact solution. Root-mean-square values of the residuals are often used as convergence criteria. Discretisation errors occur as a result of representing the governing equations as algebraic expressions on a discrete domain even though the fluid properties in reality varies continously. Refining the mesh until the solution does not change anymore is called a mesh convergence study.

Uncertainties are defined as potential deficiencies in the CFD model that are caused by lack of knowledge. There are two different kinds of uncertainties, input uncertainty and physical model uncertainty. Input uncertainty result from that it may be difficult to represent the geometry of the domain, boundary conditions and fluid properties of the real flow in the CFD model. When defining the geometry, the shape and size of regions of interest must be set. The geometry in the CFD model will probably not be exactly the same as in reality. When specifying the boundary conditions, these values might not be known with a satisfying accuracy. The location of open boundaries should be sufficiently far away from the area of interest in order to make sure that it does not affect the flow field in this area. All fluid properties depend to some extent on the local value of flow parameters. Errors are introduced if assumptions about the fluid properties are inaccurate. Physical model uncertainty results from that CFD modeling involves semi-empirical submodels such as turbulence models. Complex flows may not be accurately described by these submodels[18].

4 Method

4.1 Geometry and mesh

The main components of the original setup in the server room at Fortlax data center are two CRAC units and nine server racks which are arranged in three parallel rows. All the server racks are facing the same direction and the flow through them is in the positive x-direction. The geometry is generated in ANSYS DesignModeler. The dimensions of the server room are $10.10 \times 3.74 \times 2.54$ m. The geometry is illustrated in Figure 4.1. It is simplified by assuming that the CRAC units and all the server racks are rectangular boxes. All other components in the server room such as cables are neglected. It is also assumed that there is no space between the components and the adjacent walls. The flow inside the CRAC units and the server racks is not resolved. All components are modeled as "black boxes". Boundary conditions are applied to the sides where the flow goes out from and into the CRAC units and the server racks. These sides represent the inlets and outlets of the computational domain.

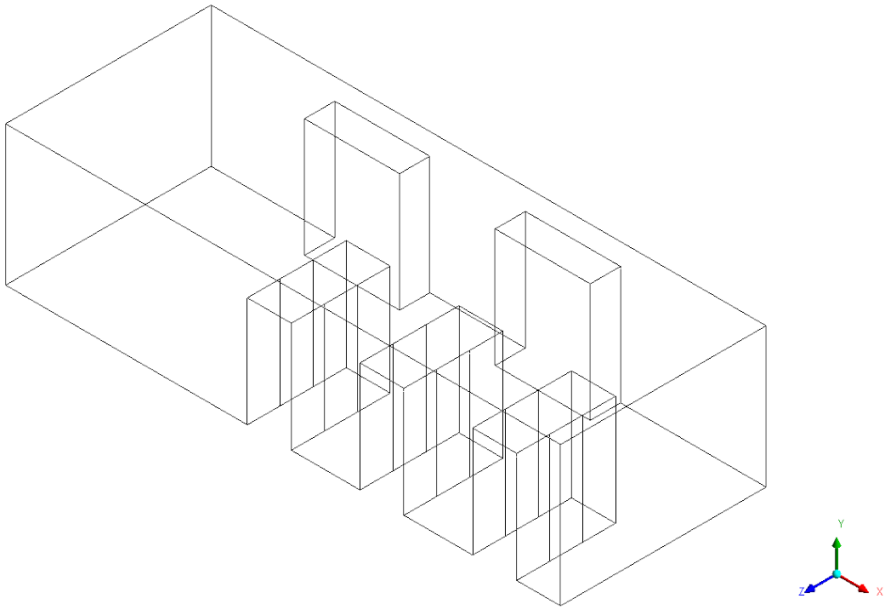

Figure 4.1: The geometry of the original setup in the server room at Fortlax data center.

The mesh is generated in ANSYS ICEM CFD. A hexahedral mesh is used because of the geometrical simplicity. The computational domain is divided into rectangular blocks which are connected to each other by edges and vertices. The edges and vertices of the blocks are associated with the curves and points of the geometry. The blocks are then divided into small hexahedral cells. The maximum cell size setting is used to generate meshes with different sizes in order to do a mesh convergence study on the original setup. The mesh with the maximum cell size used during the rest of the simulations is illustrated in Figure 4.2. The value of the maximum cell size is equal to 5 cm. The aspect ratio is kept as close as one as possible but there are deviations at some locations as a result of the dimensions of the geometry.

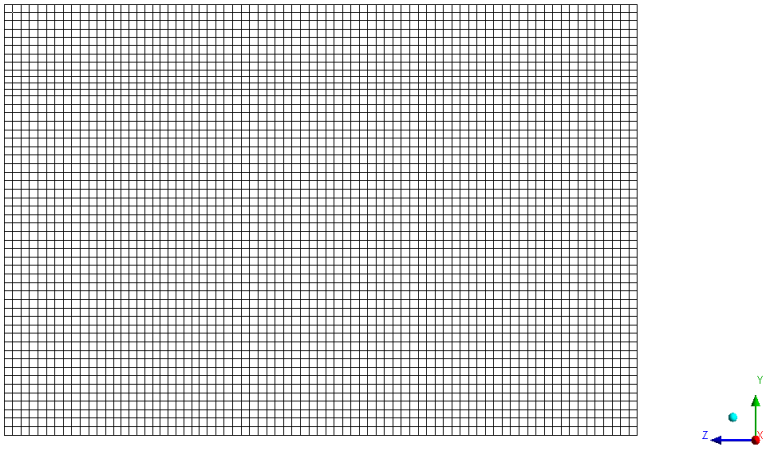

Figure 4.2: The hexahedral mesh with 754 054 *nodes viewed from the short side.*

The $k-\varepsilon$ model with scalable wall functions is used. The dimensionless wall distance from the first interior node is illustrated in Figure 4.3. The values of y^+ are well below 500 everywhere except for above the CRAC units. It results in a contribution to the physical model uncertainty. In order to keep the values of y^+ below 500 everywhere, a mesh where the cell size was decreased from the global value 5 cm to some smaller value near the walls was tested. But this mesh results in convergence problems. The solution as well as the residuals are oscillating which indicate that a steady solution cannot be found. Unsteady effects seem to be present and it is captured when the mesh is refined. Therefore, refinements near the walls are not used.

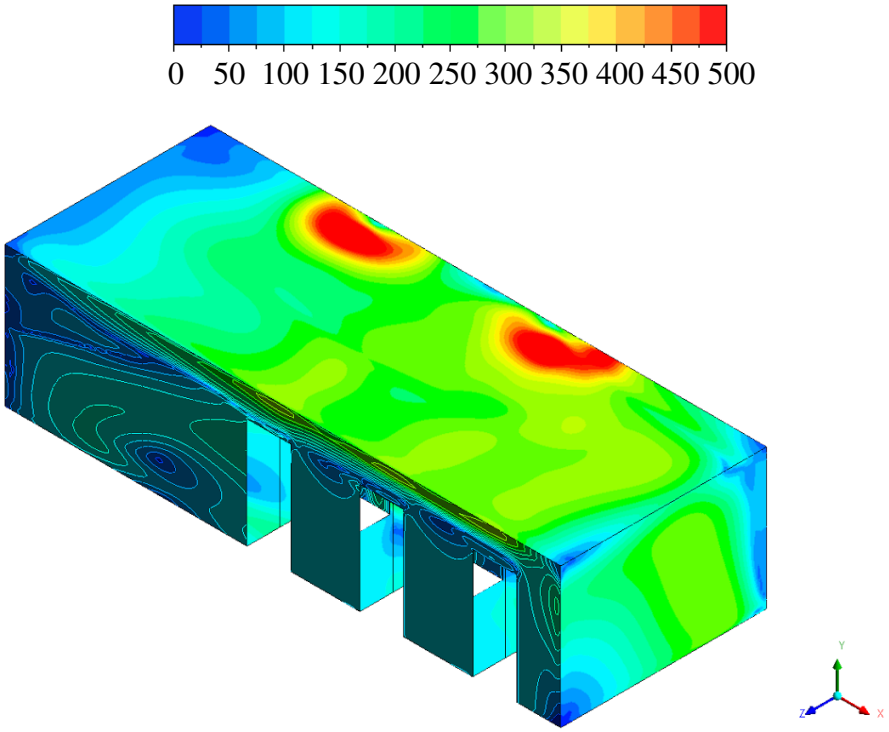

Figure 4.3: The values of y^+ from the first interior node.

4.2 Boundary conditions

The walls, the floor and the ceiling are assumed to be adiabatic. It is also assumed that there is no leakage. Thermal energy is used to model heat transfer. The CRAC units supply air at a velocity of 0.7 m/s and a temperature of 17°C. The mass flow rate should be kept constant between the inlet and the outlet located at both CRAC units, respectively. The location of the inlet and outlet boundary conditions on a CRAC unit is illustrated in Figure 4.4. Observe that the names inlet and outlet refer to the computational domain and should not be confused with the flow direction in the CRAC unit.

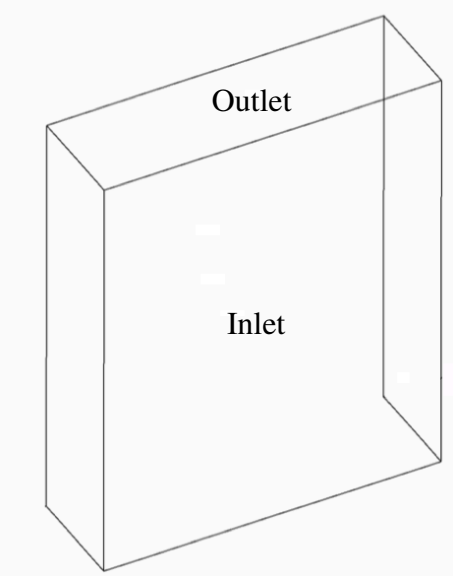

Figure 4.4: Location of the boundary conditions on a CRAC unit.

The internal fans in the server racks draw air through the front at a velocity of 0.58 m/s. Each server rack generates a heat load of 5.7 kW. Observe that it has been assumed that all the server racks generate full heat load. The heat load results in a temperature difference between the back and the front of the server rack. The temperature difference is calculated as

$$\Delta T = \frac{H}{\dot{m} c_p} \tag{4.1}$$

It is implemented by setting the temperature at the back of the server rack equal to the average of the temperature at the front plus the temperature difference. Since the back of a server rack is one single surface, the temperature of the air exhausted by the back of a server rack is uniform. In the original setup in the server room, the hot air exhausted by the back of a server rack enters the front of another server rack. Therefore, the temperature of the exhausted air might affect the temperature at the front of the next server rack so much that it is questionable if the server racks can be modeled as one single surface. It will be further investigated and an alternative approach is described in Section 4.4. The ideal gas law is used to calculate density. Since the density of the air varies with temperature, the mass flow rate will not be the same for all the server racks. But the mass flow rate should be kept constant between the outlets and the inlets located at all the server racks, respectively.

4.3 Buoyancy

The Archimedes number is a dimensionless parameter that can be used to estimate the relative importance of buoyancy and inertia forces. It is expressed as

$$Ar = \frac{\beta g L \Delta T}{U^2} \tag{4.2}$$

where the volumetric thermal expansion coefficient is equal to $\beta = 1/T$ for an ideal gas. The characteristic length scale should be in the vertical direction and the temperature difference should be evaluated between hot and cold features. If the Archimedes number is of order one, buoyancy and inertia forces are of the same order of magnitude and both must be taken into account in the CFD model to give accurate results. Buoyancy should only be neglected when the Archimedes number is very small.

The characteristic length scale is taken as the height of the server racks. The temperature difference could be taken as the global temperature difference between the outlet and inlet on a CRAC unit. But since there is no available specification by the manufacturer of what heat removal capacity the CRAC units have, that temperature difference is unknown. The local temperature difference between the back and the front of a server rack is therefore used instead. It has a smaller value than the global temperature difference and by using Equation 4.1 it is found that the local temperature difference is approximately equal to 7°C. The characteristic velocity is taken as the velocity of the air that enters the front of the server racks. Based on these characteristic values and by using Equation 4.2 it is found that the Archimedes number is approximately equal to 1.36 which indicates that buoyancy should be included in the CFD model. The full buoyancy model is used. It evaluates the difference between the density and the buoyancy reference density. The buoyancy reference density is set to 1.185 kg/m^3.

4.4 Setups

To begin with, the current cooling system with the current placement of the main components in the server room will be simulated. Observe that it has been assumed that all the server racks generate full heat load. A mesh convergence study will also be made on the original setup. A range of modifications to the original setup will then be made in order to develop the CFD model. The server racks will be modeled both as having one surface and four surfaces in the vertical direction on the back and front of each server rack. The two different ways to model a row of server racks are illustrated in Figure 4.5. The temperature condition is applied separately to each surface which results in a non-uniform temperature of the flow exhausted by the back of a server rack when four surfaces are used. When this is the only modification, the setup will still be refered to as the original setup. The other setups will be more different. Three more server racks will be included so that there is in total four rows instead of three. The boundary conditions on the server racks are the same as in the original setup. It means that the heat load and the velocity of the air that enters through the fronts are unchanged. Two different design strategies will be used. The server racks will be placed both in parallel rows similar to in the original setup and based on a hot/cold-aisle design strategy. The two different ways to model the server racks will be used for both design strategies.

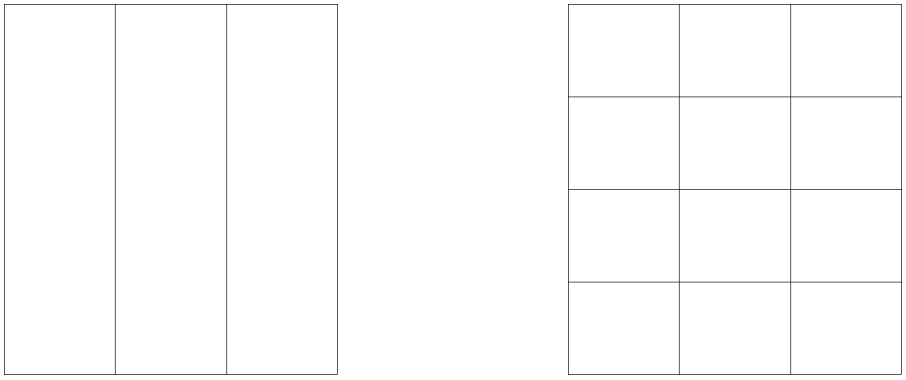

(a) One surface. *(b) Four surfaces.*

Figure 4.5: Two ways to model the server racks in a row.

The location of the CRAC units is the same as in the original setup. The total flow rate supplied by the CRAC units is increased so that it is equal to the flow rate demanded by the server racks. It results in that the CRAC units supply air at a velocity of 1.12 m/s. The purpose of this modification is further explained in Section 4.5. The temperature of the air supplied by the CRAC units is unchanged. The hard-floor configuration will also be exchanged for a raised-floor configuration. The boundary conditions of the CRAC units are different when a raised-floor configuration is considered. The inlets located at the CRAC units do not longer exist. Inlets are instead located at the perforated tiles which are placed on the floor directly in front of the server racks. The dimensions of the perforated tiles are 60×60 cm and the percentage open area is assumed to be equal to 25%. The flow in the under-floor space is not resolved. The total mass flow rate should be kept constant between the perforated tiles and the outlets located at the CRAC units. The geometry of the raised-floor configuration with hot/cold-aisles and server racks modeled as four surfaces is illustrated in Figure 4.6.

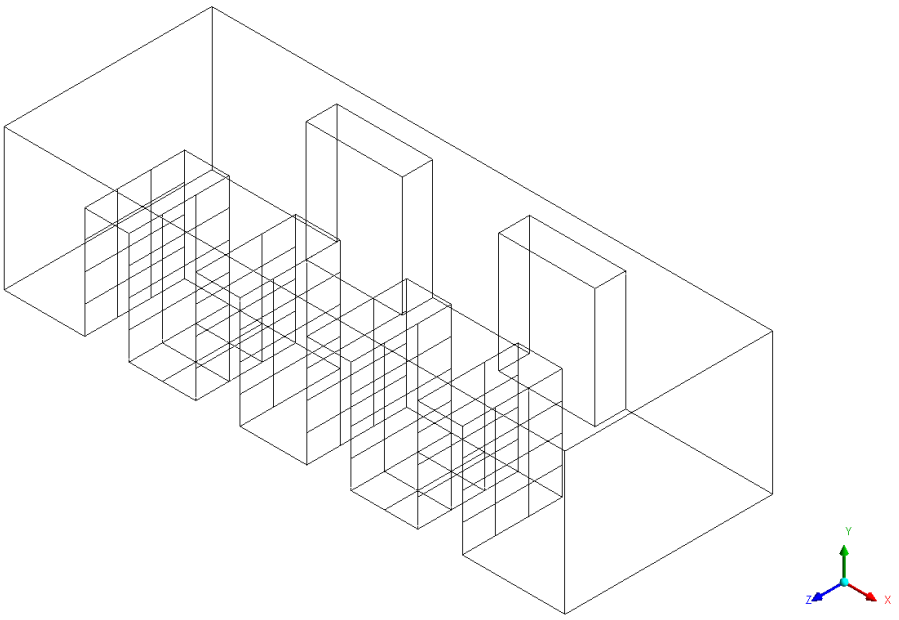

Figure 4.6: The raised-floor configuration with hot/cold-aisles and four surfaces.

The perforated tiles will be modeled by using a momentum source to correct for the momentum deficit that occurs when the perforated tiles are modeled as fully open. The reasoning behind this model was further explained in Chapter 2. The body force per unit volume in the vertical direction is expressed as

$$S_{\text{general},y} = \frac{1}{\text{V}} \rho Q \left(\frac{Q}{\sigma A_{\text{tile}}} - \frac{Q}{A_{\text{tile}}} \right) \tag{4.3}$$

where the volume refers to the region above each of the perforated tiles where the body force is distributed. The region above a perforated tile is a subdomain with the same area as the perforated tile and a height equal to 10 cm. It is illustrated in Figure 4.7.

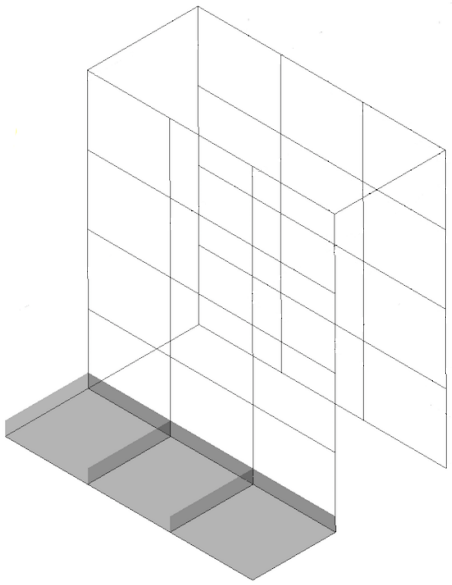

Figure 4.7: The subdomains above a row of perforated tiles.

In total eight different setups will be evaluated in addition to the original setup. It is all the resulting combinations of the two different ways to model the server racks, the two different design strategies and the two different cooling systems.

4.5 Performance metrics

A combination of Rack Cooling Index, Return Temperature Index and Capture Index will be used to evaluate and compare the results of the setups. The RCI has two parts, one that focuses on the high end of the temperature range and one that focuses on the low end. Only the high end of the temperature range will be considered here. Over-temperature is what it is called when the maximum value of the temperature at the front of a server rack exceeds the maximum recommended temperature. A summation of the over-temperature across all the server racks is called the total over-temperature. The RCI is defined as

$$\text{RCI} = \left(1 - \frac{\text{Total over-temperature}}{\text{Max allowable over-temperature}} \right) 100\,[\%] \tag{4.4}$$

where the max allowable over-temperature is a summation of the difference between the maximum allowable and maximum recommended temperature across all the server racks.

If the temperature at the front of all server racks is below the maximum recommended temperature, the RCI is equal to 100% which is the ideal value. If the temperature at the front of at least one server rack exceeds the maximum recommended temperature, the RCI is less than 100%. The RCI will be evaluated for all setups. If the temperature at the front of at least one server rack exceeds the maximum allowable temperature, the value is marked with a warning flag "*".

The RTI will not be used as a performance metrics in the true sense of the word. It will not be used as a parameter to compare the setups. Instead it is assigned a constant value. The RTI is defined as

$$\text{RTI} = \left(\frac{\Delta T_{\text{CRAC units}}}{\Delta T_{\text{Server racks}}} \right) 100\,[\%] \tag{4.5}$$

where $\Delta T_{\text{CRAC units}}$ is the temperature difference between the air that returns to and the air that is supplied by the CRAC units and $\Delta T_{\text{Server racks}}$ is the temperature increase from the front to the back of the server racks. Both values are weighted averages. The RTI will be equal to 100% for all setups. It is achieved by setting the flow rate supplied by the CRAC units equal to the flow rate demanded by the server racks. The purpose is to use a realistic constant flow rate in order to make the comparison of the other performance metrics meaningful since for example the RCI can be improved by supplying air at an increased flow rate.

The CI is typically used as a performance metric for setups with hot/cold-aisles. However, the use of CI is not restricted to certain setups. It has two parts, one that focuses on the cold-aisles and one that focuses on the hot-aisles. Only the CI for cold-aisles will be considered here. The CI is then defined as the fraction of air that enters the front of a server rack that orginates from the CRAC units or the perforated tiles. All the server racks are assigned individual values. If all the air that enters the front of a server rack orginates from the CRAC units or the perforated tiles, the CI is equal to 100% which is the ideal value. However, a low value does not necessarily imply that the server rack is insufficiently cooled. Most of the air that enters the front of a server rack might orginate from the surrounding room environment instead of from the CRAC units or the perforated tiles. The CI will be low but if the temperature of the surrounding room environment is low enough, the server rack will be sufficiently cooled. This situation is more complex and unpredictable. Therefore, high values of CI are preferable in data centers even though it might not be necessary. The CI will be evaluated for all the server racks in all the setups. To be able to evaluate the CI, an additional variable is introduced in the CFD model. A volumetric additional variable is defined and a transport equation is solved for it. The transport of this additional variable is a purely convective process. The concentration is set equal to one at the inlets located on the CRAC units or the perforated tiles and equal to zero at the inlets located on the server racks. The concentration is then evaluated at the front of the server racks. It represents the fraction of air that orginates from the CRAC units of the perforated tiles.

w

5 Results

5.1 Mesh convergence study

The geometry and boundary conditions used during the mesh convergence study are the same as stated earlier for the original setup. A convergence criteria was set as a target of 10^{-6} for the root-mean-square residuals. The values of imbalances should approach zero and the monitor points of the flow field should be steady. Meshes with cell sizes equal to 20 cm, 15 cm, 10 cm 5 cm and 2.5 cm have been used. However, convergence could not be achieved for the finest mesh. The solution as well as the residuals and the values of imbalances are oscillating which indicate that a steady solution cannot be found. Similar to when the mesh was refined close to the wall, unsteady effects seem to be present and the finest mesh captures it in contrast to the meshes with larger cell sizes.

The temperature along vertical lines at the front and back of the server racks is used to compare meshes with different sizes and to study mesh convergence. The vertical lines are located at the center of each row. The geometry with the vertical lines included is illustrated in Figure 5.1. The temperature along these lines for the meshes with cell sizes equal to 20 cm, 15 cm, 10 cm and 5 cm is illustrated in Figure 5.2. The rows of server racks are numbered in the positive x-direction, the same direction as the flow direction. It can be seen that the differences between the meshes get larger for each row. It might be an indication of that as already mentioned, the implementation of the boundary condition of the temperature should be investigated. The results are not mesh independent for a cell size of 5 cm, but since it is in line with previous studies and since a steady solution could not be obtained on a finer mesh, a cell size of 5 cm is used for the rest of the simulations.

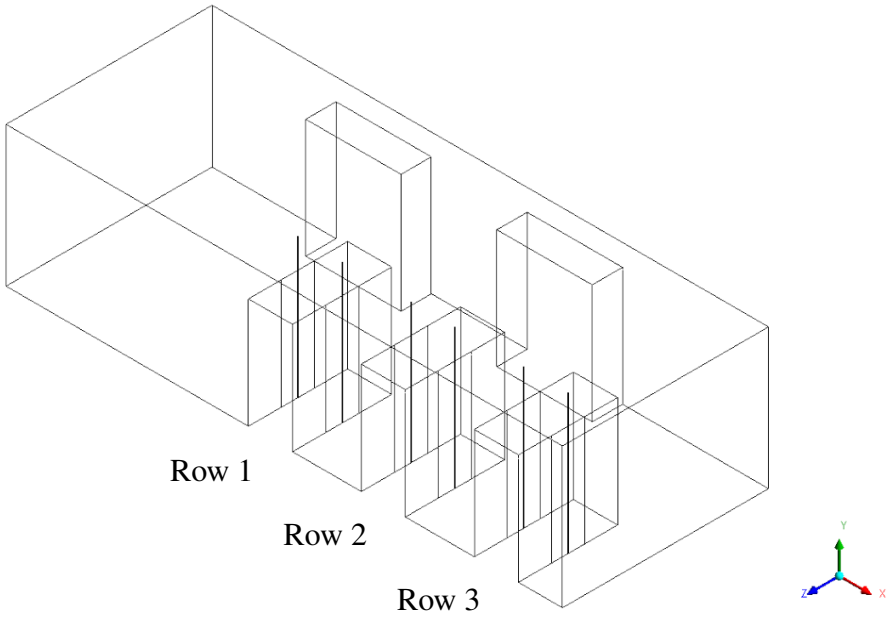

Figure 5.1: The geometry of the original setup with the vertical lines included.

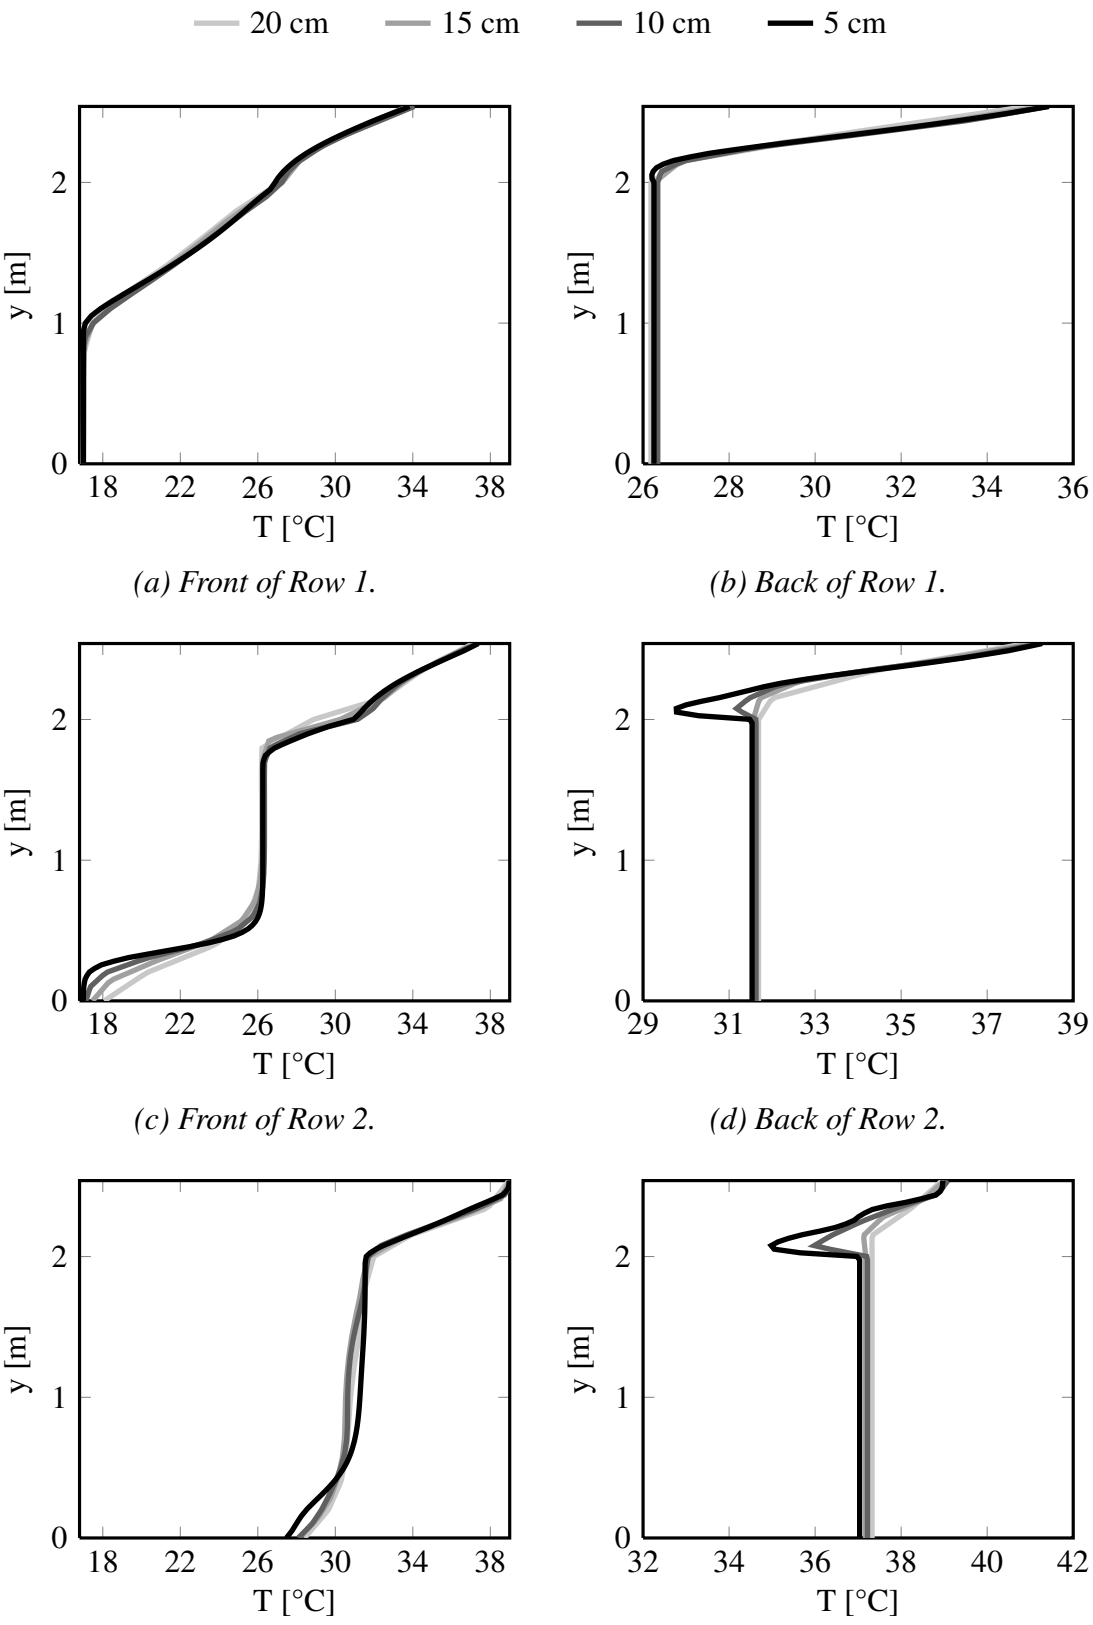

(a) Front of Row 1.

(b) Back of Row 1.

(c) Front of Row 2.

(d) Back of Row 2.

(e) Front of Row 3.

(f) Back of Row 3.

Figure 5.2: Temperature along vertical lines at the center of the rows of server racks.

5.2 Original setup

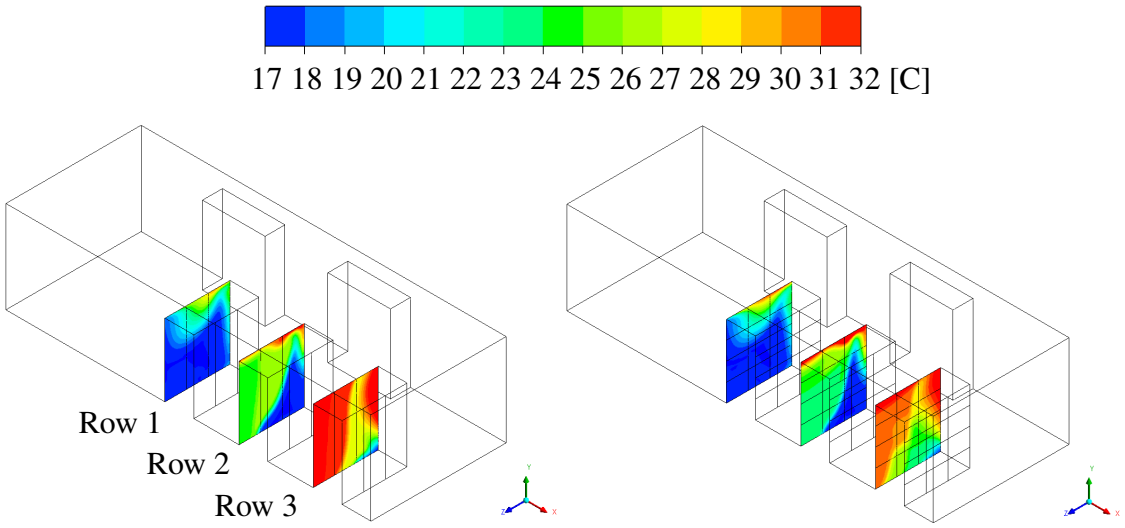

(a) One surface. *(b) Four surfaces.*

Figure 5.3: Temperature at the front of the server racks in the original setup.

The temperature distribution at the front of the server racks in the original setup is illustrated in Figure 5.3. The temperature scale ranges from 17°C which is the temperature of the air supplied by the CRAC units to 32°C which is the maximum allowed temperature at the front of the server racks. It can be seen that the temperature exceeds the maximum allowed temperature at a small part of the server racks in Row 2 and at a large part of the server racks in Row 3 both when the server racks are modeled as one surface and four surfaces. When the server racks are modeled as four surfaces, the maximum allowed temperature is also exceeded at a small part of the server racks in Row 1. It means that if the heat load in the server room at Fortlax data center would be increased to three full rows of server racks, the effect of the current cooling system needs to be increased. The RCI for the original setup has been calculated by using Equation 4.4 and is illustrated in Figure 5.4. Setup (a) and Setup (b) refer to the setups with the server racks modeled as one surface and four surfaces which are illustrated in Figure 5.3 (a) and Figure 5.3 (b), respectively. The temperature is far above the maximum recommended temperature of 27°C in large parts of the original setup. Therefore, the values of RCI are very poor and even negative when the server racks are modeled as four surfaces.

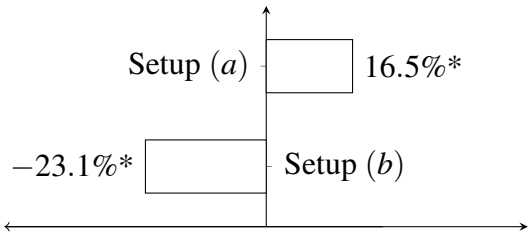

Figure 5.4: RCI for the original setup.

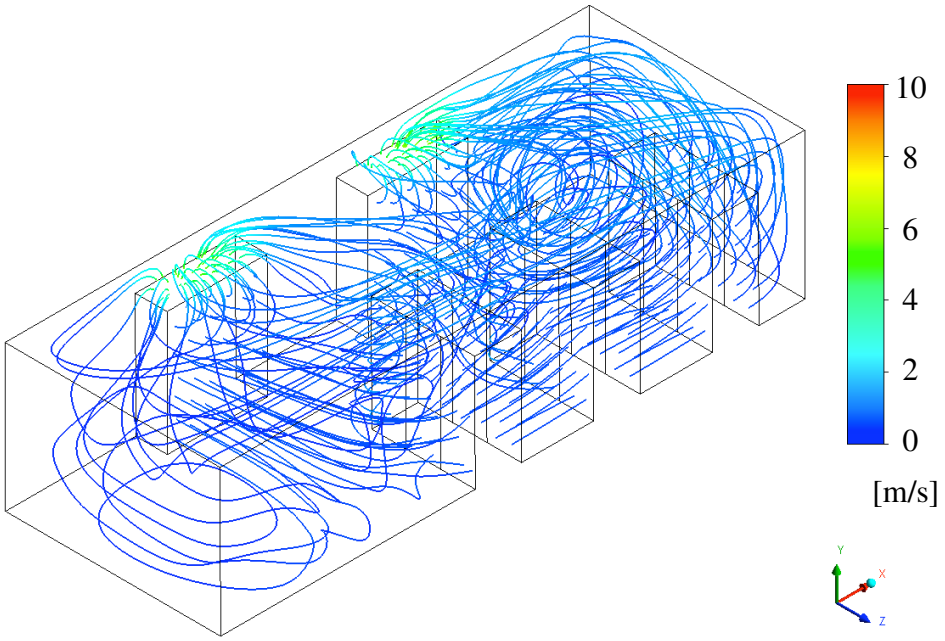

Figure 5.5: Streamlines of the flow in the original setup.

Why especially the server racks in Row 3 in the original setup are insufficiently cooled can be further explained by looking at the flow field. The flow field of the original setup with the server racks modeled as one surface is illustrated in Figure 5.5. There is no significant difference in the flow field when the server racks are modeled as four surfaces. The colors of the streamlines represent velocity and the same velocity scale will be used in the following sections. The flow direction through the server racks is in the positive x-direction. It can be seen that most of the hot air exhausted by the back of a server rack enters the front of another server. The CI for the original setup is illustrated in Figure 5.6. It shows that most of the cold air supplied by the CRAC units does not reach the server racks in Row 2 and Row 3.

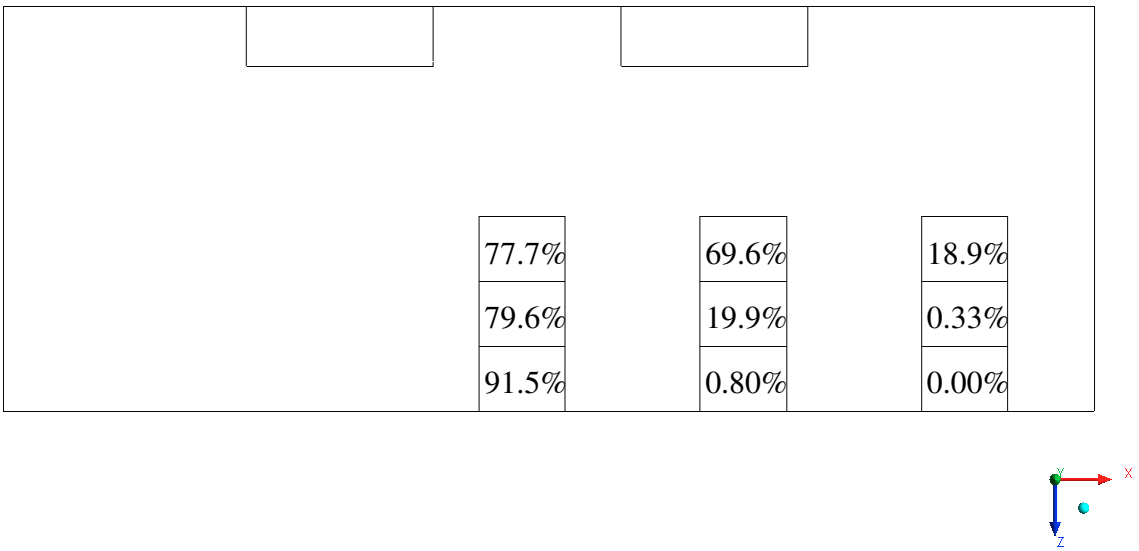

Figure 5.6: CI for the original setup.

5.3 Hard-floor configurations

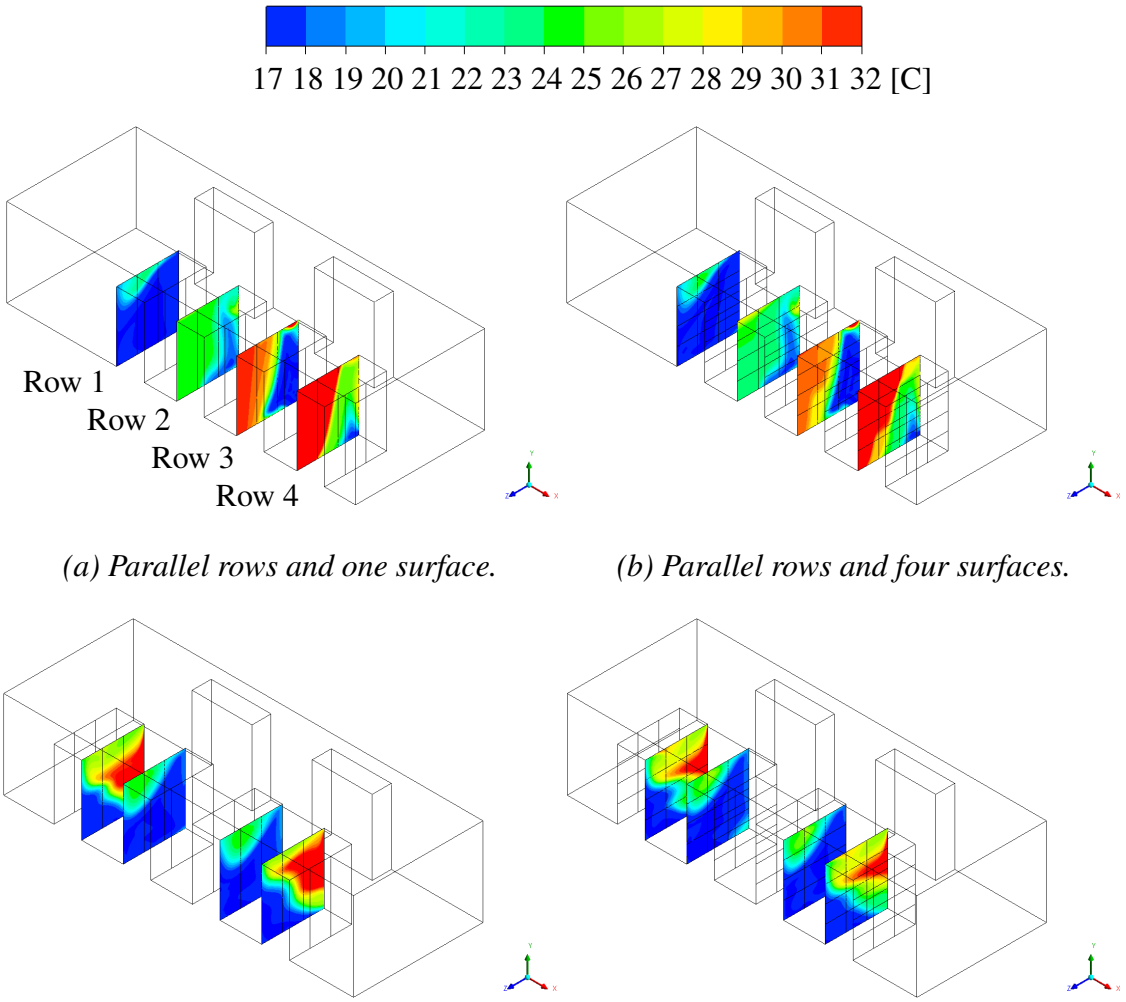

(a) Parallel rows and one surface.

(b) Parallel rows and four surfaces.

(c) Hot/cold-aisles and one surface.

(c) Hot/cold aisles and four surfaces.

Figure 5.7: Temperature at the front of the server racks in the hard-floor configurations.

The temperature distribution at the front of the server racks in the setups with hard-floor configurations is illustrated in Figure 5.7. The temperature scale is the same as in the previous section. In the setups with parallel rows, the temperature exceeds the maximum allowed temperature at large parts of the server racks in Row 3 and Row 4. In the setups when a hot/cold-aisle design strategy is used, the temperature exceeds the maximum allowed temperature at large parts of the server racks in Row 1 and Row 4. It means that the if the heat load in the server room at Fortlax data center would be increased to four full rows of server racks, the cooling system needs to be revised. There are some differences in the temperature distribution at the front of the server racks when they are modeled as one surface compared to four surfaces. When parallel rows are used, the difference increases for each row in the flow direction. When hot/cold-aisles are used, larger parts at the end of the rows exceed the maximum allowed temperature when the server racks are modeled as one surface compared to four surfaces.

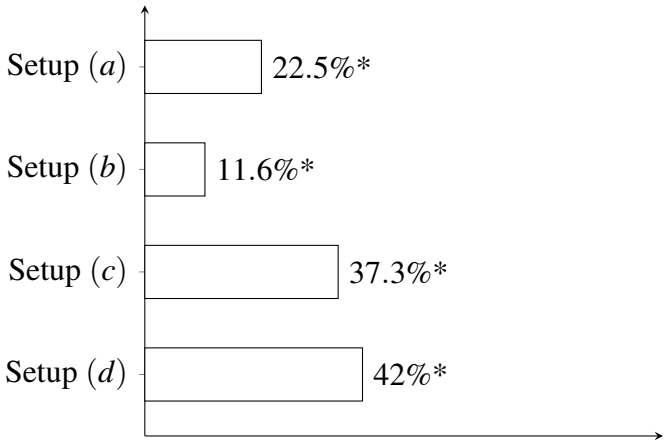

Figure 5.8: RCI for the setups with hard-floor configurations.

The RCI for the setups with hard-floor configurations has been calculated by using Equation 4.4 and is illustrated in Figure 5.8. Setup (a) refers to the setup illustrated in Figure 5.7 (a) and the same applies to the other setups. Large parts in all the setups exceed the maximum allowed temperature at the front of the server racks and even larger parts exceed the maximum recommended temperature. It results in very poor values of the RCI. In the setups with parallel rows, the over-temperature in Row 3 and Row 4 is lower when the server racks are modeled as one surface compared to four surfaces which results in higher RCI. In the setups with hot/cold-aisles, the difference is not as large even though the over-temperature is higher in Row 1 and Row 4 when the server racks are modeled as one surface compared to four surfaces which results in lower RCI.

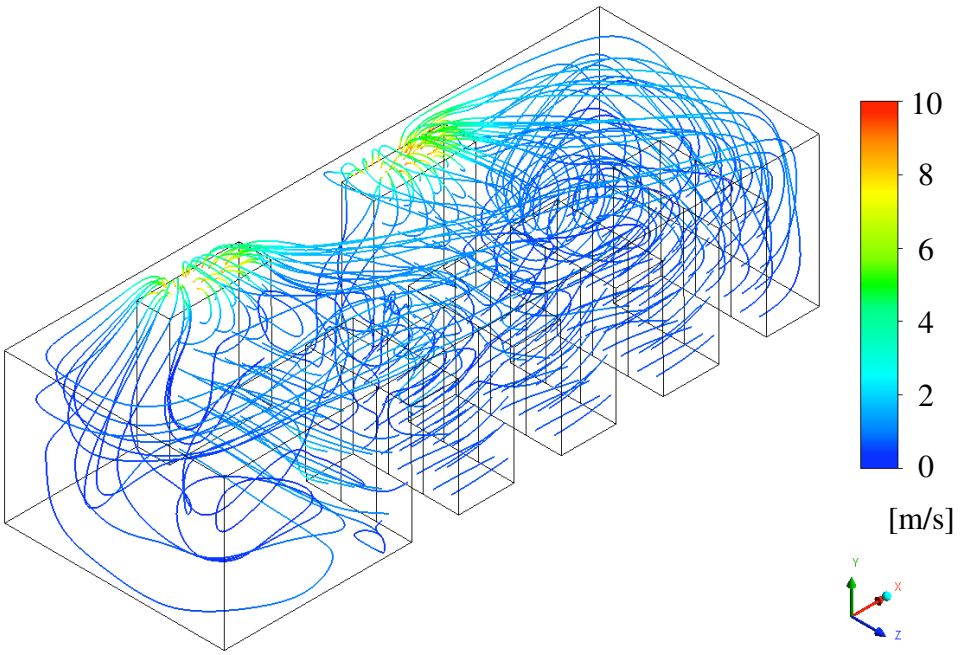

(a) Parallel rows and one surface.

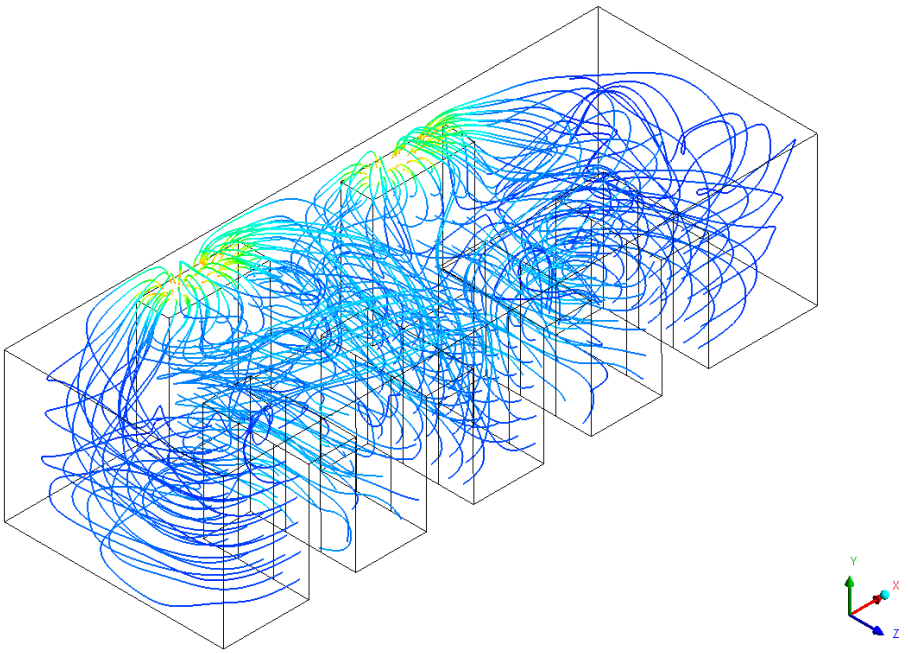

(b) Hot/cold-aisles and one surface.

Figure 5.9: Streamlines of the flow in the hard-floor configurations.

Why especially Row 3 and Row 4 are insufficiently cooled in the setups with parallel rows and Row 1 and Row 4 in the setups with hot/cold-aisles can be further explained by looking at the flow field. The flow field of the hard-floor configurations is illustrated in Figure 5.9. There are no significant differences in the flow field when the server racks are modeled as one surface or four surfaces. Therefore, only the flow field for the setups with the server racks modeled as one surface is illustrated here. The colors of the streamlines represent velocity.

In the setups with parallel rows, the flow direction through the server racks is in the positive x-direction. It can be seen that most of the hot air exhausted by the back of a server rack enters the front of another server rack. If the temperature of the air exhausted by the back of a server rack would be sufficiently low, the server rack where the air enters the front would be sufficiently cooled. This is what happens from Row 1 to Row 2 and the temperature that enters the server racks in Row 2 is well below the maximum allowable temperature. But since this procedure is repeated again and again, the server racks in Row 3 and Row 4 are insufficiently cooled. The same reasoning applies to Row 3 in the original setup in the previous section.

In the setups with hot/cold-aisles, it can be seen that the flow field in the server room is divided into two almost equal parts where each CRAC unit supplies cold air to two of the rows with server racks. The air that is exhausted by Row 2 and Row 3 meet and is pushed up toward the ceiling before it returns to the CRAC units. Some fraction of the air supplied by the CRAC units return directly to the CRAC units without cooling any server racks. Since the flow rate supplied by the CRAC units is equal to the flow rate demanded by the server racks, it results in hot air recirculation. This is what happens at the end of Row 1 and Row 4. The air exhausted by the back of the server racks is recirculated and enters the front of the server racks in the same row again.

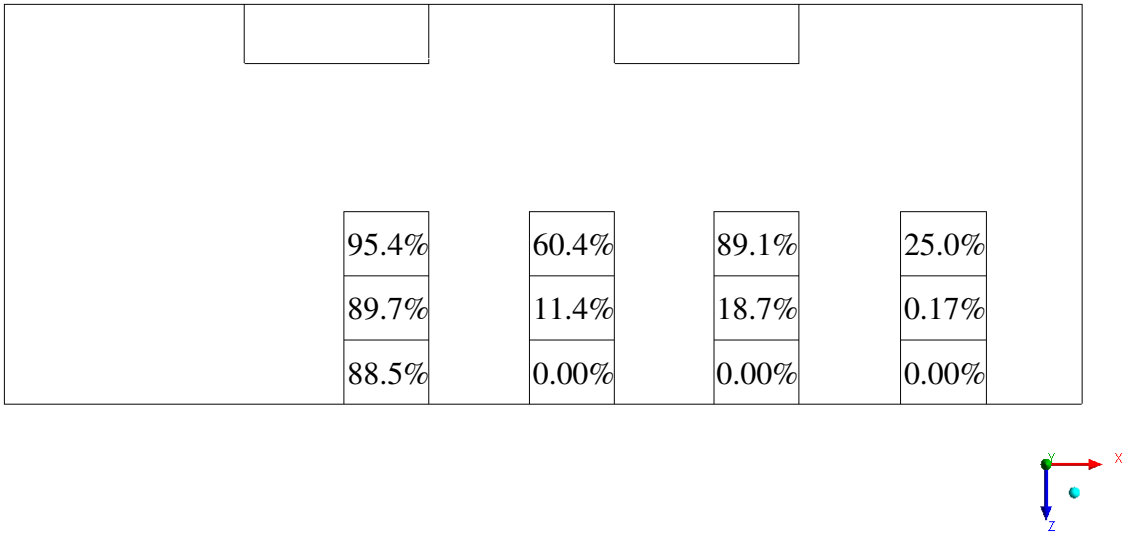

(a) Parallel rows and one surface.

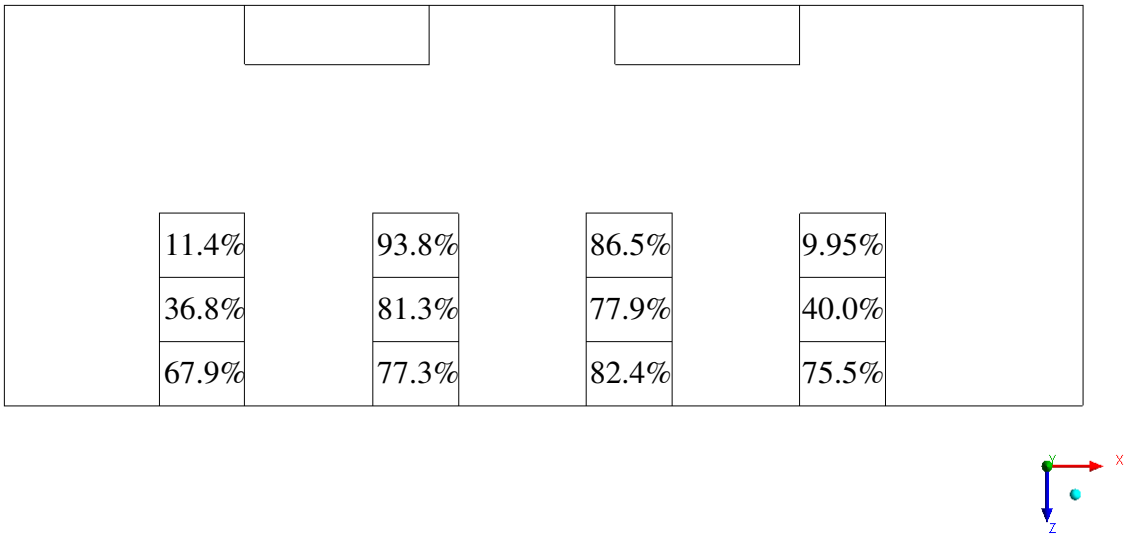

(b) Hot/cold-aisles and one surface.

Figure 5.10: CI for the setups with hard-floor configurations.

The CI for the setups with hard-floor configurations is illustrated in Figure 5.10. The values of the CI support the reasoning based on the flow field. In the setups with parallel rows, it can be seen that all the air that enters the front of the server racks closest to the wall in Row 2, Row 3 and Row 4 already has been exhausted by the back of another server rack. Most of the cold air supplied by the CRAC units does not reach the server racks in these rows. No air that enters the front of the server racks closest to the wall in Row 2, Row 3 and Row 4 orginates from the CRAC units and it results in CI equal to zero. However, the server racks in Row 2 are sufficiently cooled as a result of the relatively low temperature of the air exhausted by the server racks in Row 1. In the setups with hot/cold-aisles, it can be seen that the values of CI are higher for the server racks closest to the wall. The server racks at the end of Row 1 and Row 4 have the most poor values of CI as a result of the hot air recirculation that occurs there.

5.4 Raised-floor configurations

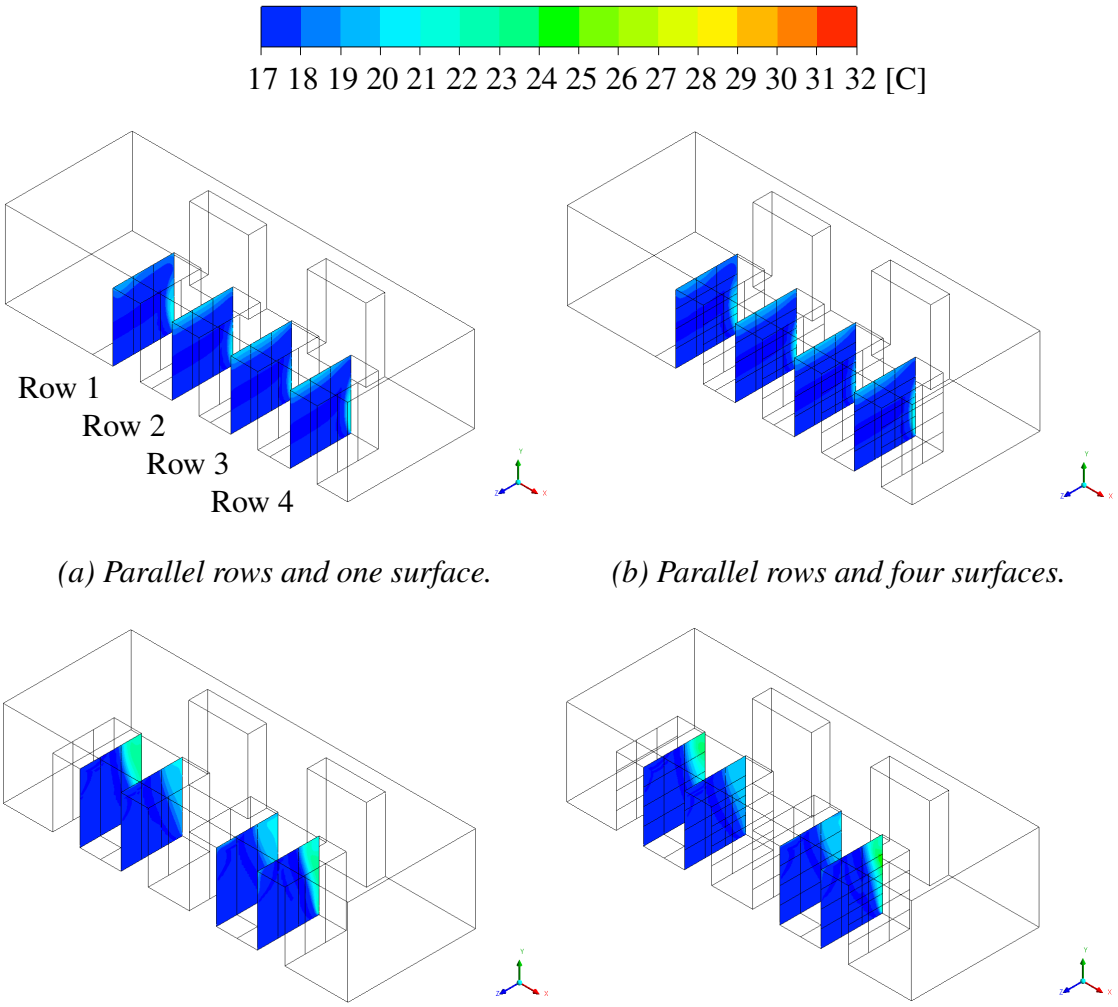

(a) Parallel rows and one surface.

(b) Parallel rows and four surfaces.

(c) Hot/cold-aisles and one surface.

(c) Hot/cold aisles and four surfaces.

Figure 5.11: Temperature at the front of the server racks in the raised-floor configurations.

The temperature distribution at the front of the server racks in the different setups with raised-floor configurations is illustrated in Figure 5.11. The temperature scale is the same as in the previous sections in order to be able to compare the results. However, a much smaller temperature scale could have been used for these setups. The temperature does not exceed either the maximum allowed temperature or the maximum recommended temperature at any part of the server racks in any of the setups. It means that if the heat load in the server room at Fortlax data center would be increased to four full rows of server racks, the cooling system would be sufficient for all the setups based on a raised-floor configuration. There are no significant differences in the temperature distributions at the front of the server racks when they are modeled as one surface compared to four surfaces. This is true both when parallel rows and hot/cold-aisles are used.

Figure 5.12: RCI for the setups with raised-floor configurations.

The RCI for the setups with raised-floor configurations has been calculated by using Equation 4.4 and is illustrated in Figure 5.12. No part of the server racks in any setup exceeds the maximum recommended temperature at the front of the server racks. It results in ideal RCI values for all the setups. Why the temperature distribution at the front of the server racks is improved compared to the hard-floor configurations can be further explained by looking at the flow field. The flow field of the raised-floor configurations is illustrated in Figure 5.13. There are no significant differences in the flow field when the server racks are modeled as one surface or four surfaces. Therefore, only the flow field for the setups with the server racks modeled as one surface is illustrated here. The colors of the streamlines represent velocity.

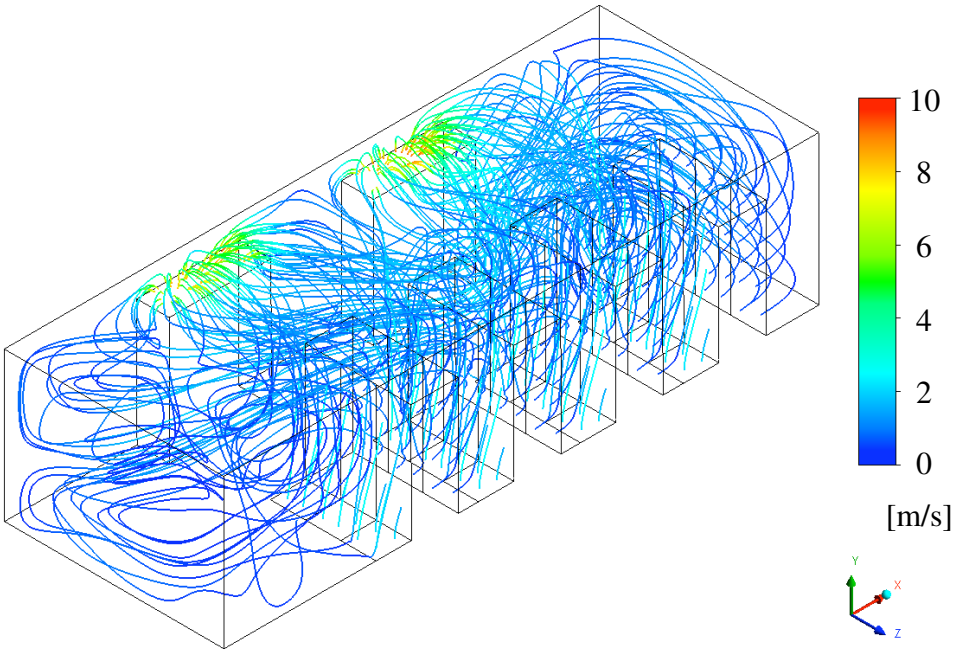

(a) Parallel rows and one surface.

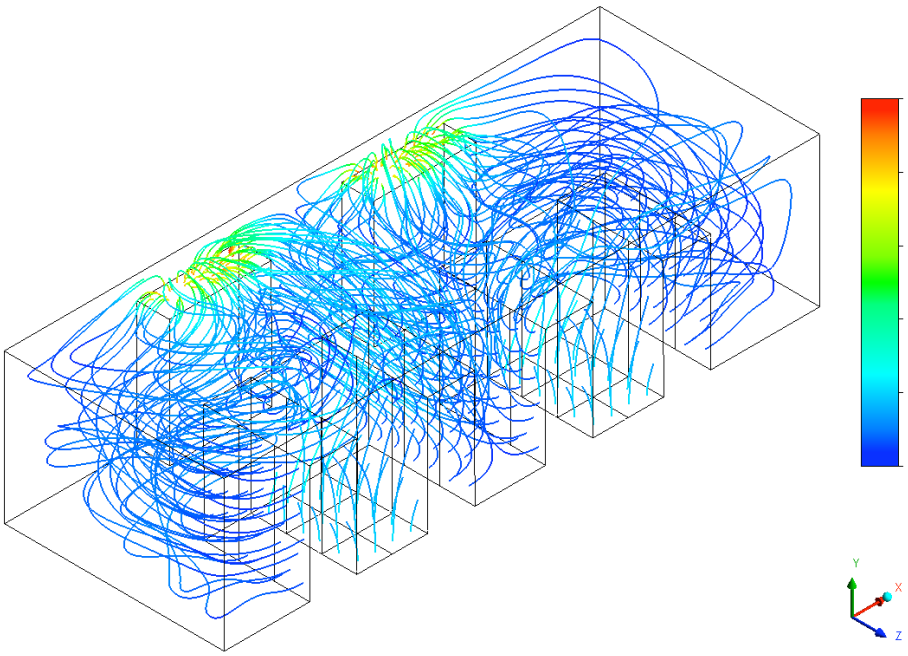

(b) Hot/cold-aisles and one surface.

Figure 5.13: Streamlines of the flow in the raised-floor configurations.

In the setups with parallel rows, the flow direction through the server racks is in the positive x-direction. It can be seen that almost all the air supplied by the perforated tiles enters the front of the server racks. Most of the hot air exhausted by the back of the server racks is pushed up towards the ceiling and returns to the CRAC units. The problem that some fraction of the air supplied by the CRAC units return directly to the CRAC units without cooling any server racks is basically eliminated. There is some hot air recirculation at the top of the server racks and at the end of the rows. There is also a small fraction of the air exhausted by the back of the server racks that enters the front of the server racks in the next row. However, the presence of this combination is to a very low extent and it does not cause any problems with too high temperatures at the front of the server racks.

In the setups with hot/cold-aisles, it can be seen that the flow field in the server room is divided into two almost equal parts again. The air that is exhausted by Row 2 and Row 3 meet and is pushed up toward the ceiling before it returns to the CRAC units. The CRAC units do not supply the cold air when raised-floor configurations are used, but each CRAC unit recieves the hot air exhausted from two of the rows with server racks. There is no problem with air that returns directly to the CRAC units in this setup either. There is some hot air recirculation at the end of the rows, but this is to a low extent in the setup with this design strategy as well and does not cause any problems.

The CI for the setups with raised-floor configurations is illustrated in Figure 5.14. It supports the reasoning about the flow field. The recirculation in the setup with parallel rows does not affect the CI as much as it does in the setup with hot/cold-aisles. However, the recirculation does not affect the temperature distribution at the front of the server racks so much that it exceeds the maximum recommended temperature and it is kept well below the maximum allowable temperature.

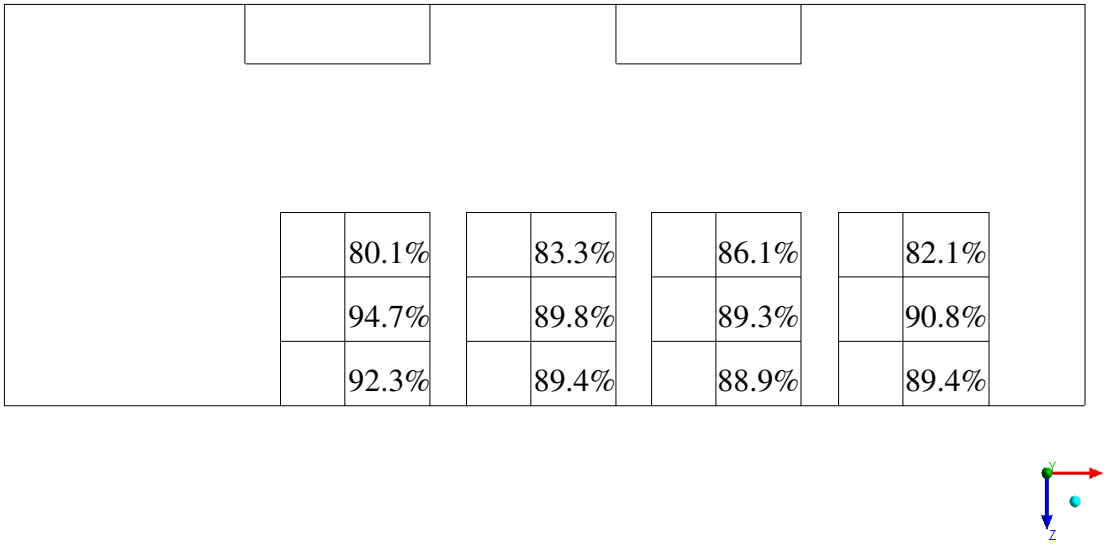

(a) Parallel rows and one surface.

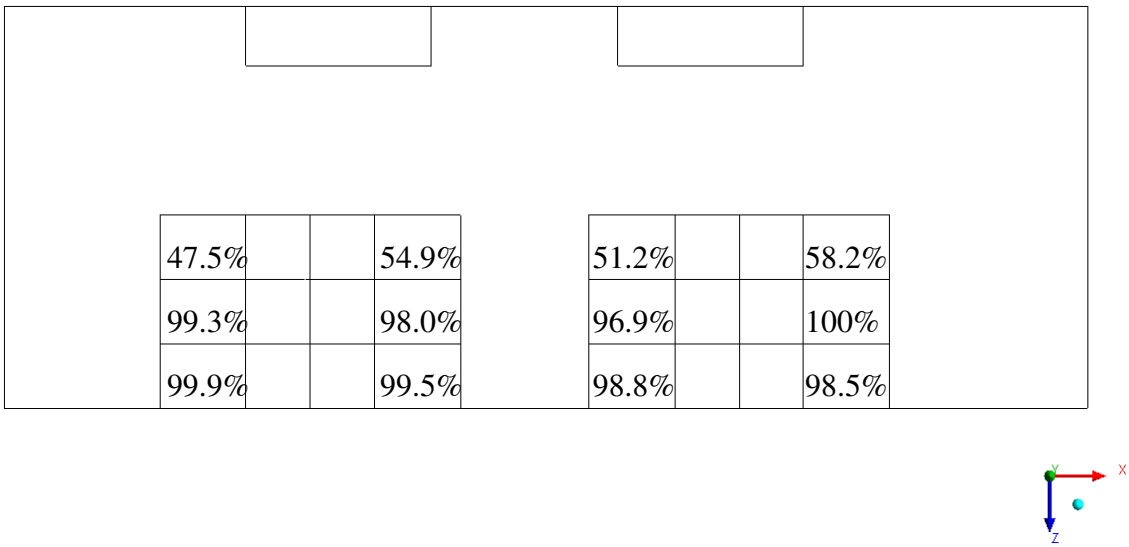

(b) Hot/cold-aisles and one surface.

Figure 5.14: CI for the setups with raised-floor configurations.

6 Discussion and conclusions

The results of the mesh convergence study were not as good as desired. The results were not mesh independent for a cell size of 5 cm even though there were not large differences in the results between the different mesh sizes. There were convergence problems when the cell size was decreased to some value smaller than 5 cm. It would have been desirable to achieve convergence on finer meshes in order to keep the values of the dimensionless wall distance from the first interior node well below 500 everywhere. It should be possible to achieve a steady solution on finer meshes as well. However, as a result of the limited time schedule and the available computer resources it could not be achieved in this thesis. It is possible that the implementation of the boundary conditions is part of the problem. CFD modeling of data centers on finer meshes have not been much further investigated in general. It is likely that it is because of mesh size limitations when significatly larger server rooms than the one considered in this thesis are simulated. The fact that it has been shown that the relative accuracy is only marginally improved when the cell size is further decreased is probably also a contributing factor. Based on this reasoning, the results obtained in this thesis are not likely to be much different on a mesh with a different cell size even though it would have been desirable to achieve convergence on finer meshes.

Important to mention is that there were some input uncertainties in the CFD model. The geometry of the domain was simplified, but it is not likely that the components that have been neglected such as cables are so important that it is worth the increase in mesh size that would be required to represent the geometry of the server room in more detail. The fluid properties of the air do probably not introduce significant errors into the CFD model either. The boundary conditions have been the most difficult to represent in the CFD model. The velocity and the temperature of the air supplied by the CRAC units as well as the velocity of the air that enters the front of the server racks have been considered as given values. But in practice, they are actually measured and these values are not known with a satisfying accuracy. More accurate measurements of these values should be performed in order to give more reliable results. Measured values of the temperature and flow field with a satisfying accuracy should also be produced in order to validate the simulations.

No general conclusions have been made about how many surfaces in the vertical direction that should be used to model the server racks. Measured values of the temperature and flow field could be used to determine which approach that best matches experimental values. How many surfaces the server racks should be modeled as might be related to the appearance of the equipment in the server racks. How many internal fans that are present and the location of them probably affect the flow field of the air exhausted by the back of the server racks. It might also be possible to use a different approach than using the average of the temperature at the front of a server rack when implementing the boundary condition that determines the temperature at the back of the server rack. It is not certain that the temperature of the air exhausted by the back of a server rack is uniform over one or even four surfaces. No indications have been found that modeling a server rack as one surface or four surfaces might be more or less appropriate for the different design strategies, respectively.

It has been shown that the heat load in the original setup in the server room at Fortlax data center cannot be increased to three full rows of server racks. The current cooling system with the current placement of the main components results in some insufficiently cooled server racks since the cold air supplied by the CRAC units does not reach all the server racks. The temperature at large parts of the front of the server racks exceeds the maximum allowable temperature. It is likely that the cooling system would be sufficient for two full rows of server racks with a heat load of 5.7 kW each, but exactly what total heat load that can be sufficiently cooled has not been concluded. It could be investigated by varying the heat load across the server racks. As already mentioned, in order to evaluate how accurately the CFD model predicts the temperature and flow field in the server room at Fortlax data center, experimental values with a satisfying accuracy should be produced. The perforated tiles in the raised-floor configurations have been modeled by using a momentum source to correct for the momentum deficit. The inlets located at the back of the server racks as well as at the CRAC units are also perforated surfaces. It is therefore likely that these also should be modeled by using momentum sources to correct for the momentum deficits in order to improve the accuracy of the CFD model.

It is not possible to use a design strategy based on hot/cold-aisles when only three rows of server racks are used. It was one of the main reasons for including a fourth row with three more server racks. Four rows of server racks could be used even if some of the server racks are only half full in order to reduce the total heat load. The temperature and flow field of the setups with four parallel rows of server racks based on a hard-floor configuration are very similar to the original setup. All the setups based on hard-floor configurations were insufficiently cooled even though the flow rate supplied by the CRAC units was increased to equal the flow rate demanded by the server racks. The comparison of the hard-floor and the raised-floor configurations was done in order to provide interesting input when Fortlax is planning to build another data center. However, it could be investigated how a ceiling return strategy would affect the performance of the cooling system and if it would be possible to implement it in the server room at Fortlax data center. Based on the current room return strategy, panels might then be used to cover the hot-aisles in order to minimize hot air recircuation when a hot/cold-aisle design strategy is used.

It could also be investigated how varying the flow rate or the temperature of the air supplied by the CRAC units would affect the temperature distribution at the front of the server racks. The temperature difference between the back and the front of a server rack is calculated by using Equation 4.1 and is approximately constant. It means that a lower temperature of the air supplied by the CRAC units would result in an equal reduction in temperature of the air that enters the server racks. This is true for all the rows since the temperature of the air exhausted by the back of the server racks in Row 1 affects the temperature of the air that enters the front of the server racks in Row 2 and so on. However, it should be noted that ASHRAE's thermal guidelines not only specifies maximum allowable and maximum recommended temperatures but also minimum allowable and minimum recommended temperatures at the front of the server racks. The minimum allowable temperature at the front of the server racks is equal to 15°C. There is therefore a limit to how much the supply temperature can be used to improve the performance of the cooling system.

The benefits of using a raised-floor configuration compared to a hard-floor configuration are clear. The performance of the cooling system was significantly improved when the hard-floor configuration was replaced by a raised-floor configuration regardless of what design strategy that was used for the server racks. The main difference is that the perforated tiles supply cold air to the server racks right in front of them. Most of the air that enters the front of the server racks orginates from the perforated tiles which results in temperatures close to the supply temperature and therefore almost ideal CI values for most of the server racks. The design strategy based on hot/cold-aisles has become the standard when raised-floor configurations are used. However, the raised-floor configuration with parallel rows of server racks turned out to be the setup that performed best based on the performance metrics that were used to evaluate the results. It is not unlikely that another choice of performance metrics could have given another result. It is therefore recommended to consider raised-floor configurations in general when building the new data center. This recommendation is only based on the performance of the cooling system and not on any other factors that also might be important for Fortlax.

References

[1] Fortlax [Internet]. Piteå, Sweden: Fortlax; 2014 [cited 2015 Feb 12]. Available from: http://www.fortlax.se/om-fortlax/.

[2] Koomey J. Growth in Data center electricity use 2005 to 2010 [Internet]. Oakland, USA: Analytics Press; 2011 [cited 2015 Feb 10] Available from: http://www.analyticspress.com/datacenters.html/.

[3] Patankar SV. Airflow and Cooling in a Data Center. Journal of Heat Transfer. 2010;132(7):073001-01-073001-17.

[4] Sullivan RF. Alternating Cold and Hot Aisles Provides More Reliable Cooling for Server Farms. 2010. White Paper from the Uptime Institute.

[5] Srinarayana N, Fakhim B, Behnia M, Armfield SW. A Comparative Study of Raised-Floor and Hard-Floor Configurations in an Air-Cooled Data Centre. Proceedings of the 13th IEEE ITHERM Conference; 2012 May 30- Jun 1; San Diego, USA.

[6] VanGlider JW, Shrivastava SK, Zhang X; American Power Conversion Corporation. System and method for evaluating equipment rack cooling performance. United States patent 7991592. 2012 Aug 11.

[7] Karki KC, Radmehr A, Patankar SV. Use of Computational Fluid Dynamics for Calculating Flow Rates Through Perforated Tiles in Raised-Floor Data Centers. HVAC&R Research. 2003;9(2);153-166.

[8] ASHRAE Inc. 2011 Thermal Guidelines for Data Processing Environments. 2011. White paper from ASHRAE.

[9] Data Center Products [Internet]. Jessup, USA: Tate; 2015 [cited 2015 May 15]. Available from: http://tateinc.com/products/airflow_perfs.aspx/.

[10] Iyengar M, Schmidth RR, Hamann H, VanGlider J. Comparison Between Numerical and Experimental Temperature Distributions in a Small Data Center Test Cell. Proceedings of the ASME 2007 InterPACK Conference; 2007 Jul 8-12; Vancouver, Canada.

[11] Zhang X, VanGlider JW, Iyengar M, Schmidth RR. Effect of Rack Modeling Detail on the Numerical Results of a Data Center Test Cell. Proceedings of the 11th IEEE ITHERM Conference; 2008 May 28-31; Orlando, USA.

[12] VanGlider JW, Zhang X. Coarse-Grid CFD: The Effect of Grid Size on Data Center Modeling. ASHRAE Transactions. 2008;114(2):166-181.

[13] Abdelmaksoud WA, Khalifa HE, Dang TQ, Elhadidi B. Experimental and Computational Study of Perforated Floor Tile in Data Centers. Proceedings of the 12th IEEE ITHERM Conference; 2010 Jun 2-5; Las Vegas, USA.

[14] Adbdelmaksoud WA, Khalifa HE, Dang TQ, Schmidt RR, Iyengar M. Improved CFD Modeling of a Small Data Center Test Cell. Proceedings of the 12th IEEE ITHERM Conference; 2010 June 2-5; Las Vegas, USA.

[15] Herrlin MK. Airflow and Cooling Performance of Data Centers: Two Performance Metrics. ASHRAE Transactions. 2008;114(2):182-187.

[16] VanGlider JW, Shrivastava SK. Capture Index: An Airflow-Based Rack Cooling Performance Metric. ASHRAE Transactions. 2007;113(1):126-136.

[17] ANSYS Inc. ANSYS CFX Solver Theory Guide. 16th release. Canonsburg, USA.

[18] Versteeg HK, Malalasekera W. An Introduction to Computational Fluid Dynamics: The Finite Volume Method. 2nd edition. Harlow: Pearson Prentice Hall; 2007.