



UNIVERSITY OF LEEDS

This is a repository copy of *Three computational methods for analysing thermal airflow distributions in the cooling of data centers*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/126518/>

Version: Accepted Version

Article:

de Boer, GN orcid.org/0000-0002-5647-1771, Johns, A, Delbosc, N et al. (4 more authors) (2018) Three computational methods for analysing thermal airflow distributions in the cooling of data centers. *International Journal of Numerical Methods for Heat and Fluid Flow*, 28 (2). pp. 271-288. ISSN 0961-5539

<https://doi.org/10.1108/HFF-10-2016-0431>

(c) 2018, Emerald Publishing Limited. Published by Emerald Publishing Limited. This is an author produced version of a paper published in the *International Journal of Numerical Methods for Heat and Fluid Flow*. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

Unless indicated otherwise, fulltext items are protected by copyright with all rights reserved. The copyright exception in section 29 of the Copyright, Designs and Patents Act 1988 allows the making of a single copy solely for the purpose of non-commercial research or private study within the limits of fair dealing. The publisher or other rights-holder may allow further reproduction and re-use of this version - refer to the White Rose Research Online record for this item. Where records identify the publisher as the copyright holder, users can verify any specific terms of use on the publisher's website.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Three computational methods for analysing thermal airflow distributions in the cooling of data centers

de Boer GN¹, Johns A¹, Delbosc N^{1,3}, Burdett D¹, Tatchell-Evans M², Summers JL^{1,4}, Baudot, R^{5,6}

¹ School of Mechanical Engineering, University of Leeds, Leeds, LS2 9JT, UK.

² School of Chemical and Process Engineering, University of Leeds, Leeds, LS2 9JT, UK.

³ DS SIMULIA, S.L.U. Calle Orense 58, planta 12, 28020 Madrid, SPAIN.

⁴ RISE SICS North, Björkskataleden 112 (fd Traktorvägen 1) 973 47 Luleå. SWEDEN.

⁵ 4Energy, Nottingham Science Park, Jesse Boot Avenue, Nottingham, NG7 2RU, UK.

⁶ RBCD Solutions, 105 Wilford Crescent East, Nottingham, NG2 2ED, UK.

Abstract

This paper uses three different numerical simulation technologies to analyse the thermal airflow distribution in a simple example data center layout. The numerical approaches used are based on finite element, finite volume and lattice Boltzmann methods and are respectively implemented via commercial Multiphysics software, opensource CFD code and GPU based code developed by the authors. Each method includes an appropriate turbulence model and the simulation results focus on comparison of the three methods when applied to 2 rows of datacom racks with cool air supplied by a computer room air conditioner and distributed via an underfloor plenum. Good quantitative agreement between the three methods is seen in terms of the inlet temperatures to the Datacom equipment.

Key Words: Data center cooling, Lattice Boltzmann, Finite Element, Finite Volume.

1. Introduction

Data centers are facilities hosting information and communications technology (ICT) infrastructure, usually laid out in rows of 2m tall racks. Their operation provides digital services across a range of sectors. The quantity and density of these ICT systems results in a distributed and complex dynamic generation of heat throughout the facility that needs to be transported away from the ICT, referred to as datacom, equipment. This is usually achieved by air cooling, where the heat is transferred and ultimately rejected to the outside environment. The growth of these facilities has been considerable, with data centers consuming only 0.12% of US energy consumption for the year 2000 [1] and that figure having risen to over 2% only ten years later, in 2010 [2]. Globally this figure is currently closer to 1.5%, at a growth of roughly 11% per year over the last decade [3], and approximately 45% of this energy is used in the removal of heat [4]. Improvements to the energy efficiency of these facilities are rapidly becoming paramount, due to the need to reduce both running costs and environmental impact [5].

Correct management of air distribution throughout these facilities is one way of reducing inefficiencies in data centers. Computational fluid dynamics (CFD) is a useful tool in achieving this, with detailed modelling of the internal environment allowing the prediction of hot spots, bypass air and recirculation, and other inefficiencies [6]. This enables air distribution to be optimised to minimise energy consumption whilst ensuring a suitable thermal environment is provided, both in existing and new data centers. CFD has been used successfully to investigate the impact of floor grilles [7], optimise placement of datacom equipment [8] and study the effect of aisle containment [9]. The major challenges in producing accurate models are the multiple length-scales [6], from the chip to the room level, and accurately capturing the various modes of thermal transport and flow regimes present therein [10].

The use of CFD is increasingly being applied in the analysis of air distributions in data centers as it offers a much greater resolution of thermal flow information compared to most experimental approaches, as demonstrated in [11], and affords data center operators the ability to predict and monitor unfavourable air flow patterns. There exists a range of modelling methods and software packages available for this task, each with individual advantages and disadvantages. This paper analyses the air flow through a simple data center layout for three distinct modelling strategies to identify the trade-offs between each methodology. The three applied methods are; i) a developed lattice Boltzmann method (LBM) based simulation code for rapid execution on a graphical processing unit (GPU), ii) an opensource CFD finite volume method (FVM) based package OpenFOAM, and the commercial CFD finite element method (FEM) based software COMSOL Multiphysics. Previous studies have demonstrated through experimental validation that CFD models can accurately predict data center air flows and thermal environments [7-11]. However, the work presented here aims to demonstrate the range of distinct numerical methods available, which produce similar results for established data center configurations and identifies the advantages and disadvantages of each approach.

2. Theory

The relevant theory relating to the thermodynamics and fluid mechanics of data center cooling is outlined. This covers the general description of a data center and the associated cooling operation, the physical phenomena which describe the fluid flow and heat transfer mechanisms, and a description of the computational methods used to model this behaviour.

2.1 Continuum Mechanics

The physics of fluid flow and heat transfer in the data center are governed by the conservation of mass, momentum and energy. A steady state condition is assumed such that the operating conditions of the data center do not change with time, which in the transient sense accounts to averaging over a number of instances in a given period with stationary operating conditions. Buoyancy forces are modelled due to the convection of heat produced by the Datacom units in which energy and momentum transports are coupled using the Boussinesq approximation [12]. Material properties of air (density, viscosity, specific heat capacity, thermal conductivity, and thermal expansion coefficient) are assumed constant over the range of temperatures and pressures observed in the data center.

The conservation of mass (1), momentum (2), and energy (3) are governed in the data center using continuum mechanics by the Reynolds-Averaged Navier-Stokes (RANS) equations for steady, incompressible fluid flow,

$$\vec{\nabla} \cdot \vec{\mathbf{u}} = 0 \quad (1)$$

$$\rho \vec{\mathbf{u}} \cdot \vec{\nabla} \vec{\mathbf{u}} = -\vec{\nabla} p + \mu_{eff} \nabla^2 \vec{\mathbf{u}} - \rho \vec{\mathbf{g}} \beta (T - T_{ref}) \quad (2)$$

$$\rho c_p \vec{\mathbf{u}} \cdot \vec{\nabla} T = \kappa_{eff} \nabla^2 T \quad (3)$$

where $\vec{\mathbf{u}}$ is the fluid velocity vector in m/s, p is the fluid pressure in Pa, T is the fluid temperature in K, $\mu_{eff} = \mu + \mu_t$ and $\mu = 1.85 \times 10^{-5}$ Pa.s is the dynamic viscosity, μ_t is the turbulent dynamic viscosity, $\kappa_{eff} = \kappa + \mu_t c_p / Pr_t$, $\kappa = 0.287$ W/(m.K) is the fluid thermal conductivity, Pr_t is the turbulent Prandtl number taken to be 0.9 for all simulations, $\beta = 3.43 \times 10^{-3}$ 1/K is the thermal expansion coefficient, $\vec{\mathbf{g}} = (0,0,-9.81)$ m/s² is the acceleration due to gravity, and $T_{ref} = 293$ K is a reference temperature for neutral buoyancy. For laminar flow μ_t is zero everywhere, which is not the case for airflows in data centers. A turbulence model must also be considered in the model since inertial forces dominate over viscous forces in the fluid flow. This is undertaken by the inclusion of the k- ϵ model that yields a varying (non-zero) turbulent viscosity, μ_t , in equations (1-3). The k- ϵ model was chosen to represent the turbulent behaviour of the data center airflow because research has been published previously showing that this provides an accurate representation of the experimental observations made in data centers under typical operating conditions [13].

In order to provide a solution to the RANS equations a discrete numerical procedure is required due to the nonlinearity of the fluid mechanics. In this paper two methods are used for this purpose: the Finite Element Method (FEM); and the Finite Volume Method (FVM). The Finite Difference Method (FDM) may also be used to solve the RANS equations [14] however this type of analysis is not explored in this paper. A third method of solving for the conservation of mass, momentum and energy is based on a different formulation than that of continuum mechanics, namely the Lattice Boltzmann Method (LBM), the details of which are given in Section 2.2.

2.1.1 Finite Element Method

In the FEM the domain of interest is discretised into a number of elements n and upon which piecewise continuous basis functions are used to express the field variables (sometimes referred to as the primitives), which for the RANS would be $\vec{\mathbf{u}}$, p , T , k and ϵ . This is shown in Figure 1 using one-dimension as an example. The discrete form of a variable $\varphi(x)$ is given by the following equation,

$$\varphi(x) = \sum_{i=1}^n \varphi_i N_i(x) \quad (4)$$

where φ_i are the nodal values of the variable, and $N_i(x)$ are the shape functions of the response. The shape functions describe the influence of a node on the total response, where for linear elements $N_i(x) = 1$ at one location and are zero elsewhere varying linearly to zero at the two adjacent nodes to the location where $N_i(x) = 1$. For the three-dimensional solution to the RANS second order tetrahedral based finite elements are applied in this work.

The nodal equations are subsequently formulated by a weighted integral (Galerkin) approach, letting $\mathcal{L}(\varphi) = s$ be the strong (differential) form of the equation to be solved. The weak (integral) form is written as follows,

$$\int_{x_i}^{x_{i+1}} \mathcal{L}(\varphi) N_i(x) dx = \int_{x_i}^{x_{i+1}} s N_i(x) dx \quad (5)$$

from which the discrete nodal equation for φ can be derived. The numerical solution to φ is subsequently calculated subject to the imposed boundary conditions, which are either Neumann (gradient-based) or Dirichlet (value-based) in description. The FEM provides a means by which the solution varies over the domain and the equations are solved to fit this response. The method grew from the development of discrete structural mechanics but can also be applied to any differential equations such as those governing fluid flow and heat transfer. In this work the computer software COMSOL Multiphysics¹ has been used to solve the data center flow problem using the FEM.




Figure 1 – Sketch of linear finite elements along the x-dimension

2.1.2 Finite Volume Method

In the FVM method the domain is discretised into a number of control volumes as shown in Figure 2 for the example one-dimensional case. The nodes are located at the centres of the cells. The faces of the cell form the control volume boundary and are halfway between the centres of the adjacent cells.




Figure 3 – Sketch of a control volume for the finite volume discretisation of the x-dimension

The nodal equation is obtained by integrating over the control volumes, for the case $\mathcal{L}(\varphi) = s$ this is described by the following equation,

$$\int_{x_{i-1/2}}^{x_{i+1/2}} \mathcal{L}(\varphi) dx = \int_{x_{i-1/2}}^{x_{i+1/2}} s dx \quad (6)$$

and for which the same boundary conditions as the FEM can be imposed to obtain a numerical solution. The approach imposed by the FVM leads to a conservative scheme where the flux on a boundary is equal to the opposing flux in the adjacent cell, making it a useful method in solving partial differential equations based on conservation laws. Unlike the FEM, the FVM grew from the development of CFD and is therefore applicable to the solution of the fluid flow and heat transfer in the data center. In this work the open source computer software OpenFOAM [15] has been used to solve the thermal airflow in the data center problem using the FVM.

2.2 Lattice Boltzmann Method

An alternative to the macroscopic continuum mechanics formulation is the Lattice Boltzmann Method (LBM), in which the governing equations of fluid flow and heat transfer are instead derived based on the collision and propagation of particles at a mesoscopic scale. Figure 3 depicts how the LBM can be regarded as a mesoscopic

¹ www.comsol.com/comsol-multiphysics

method, where the right hand side of the figure represents fluid flow that would be described by RANS and the length scales are such that the so-called continuum hypothesis holds. On the left hand side of Figure 4 there is a representation of the microscopic nature of fluid flow where methods would deterministically describe the motion of individual fluid molecules. Moving towards the middle diagram however, the scale is such that the number of particles becomes larger and it is more convenient to average all particle properties over a certain volume. A density distribution function, $f(\vec{x}, \vec{u}, t)$, is adopted that represents the number of particles per unit volume having velocity \vec{u} within the volume centred at \vec{x} and at time t . Discretising the particle velocities on a lattice as shown in Figure 4, a discrete distribution function is defined that propagates and collides as defined by the lattice Boltzmann equation (7), which is in fact a special discretisation of the Boltzmann equation; the Navier-Stokes equations can be derived via a low Mach number expansion of the lattice Boltzmann equation [16].




Figure 3 – Description of fluid motion at different scales.




Figure 4 – Discretisation of the microscopic velocities in 3D (D3Q19), which is a nineteen velocity (three speed) lattice.

$$f_i(\vec{x} + \vec{e}_i \Delta t, t + 1) - f_i(\vec{x}, t) = F_i \Delta t - \frac{f_i(\vec{x}, t) - f_i^{eq}(\vec{x}, t)}{\tau} \quad (7)$$

where the subscript i refers to one of the 19 particle velocities, \vec{e}_i , in the D3Q19 lattice (Figure 4). The left hand side of equation (7) represents the propagation step that is associated with the advection term in the fluid mechanics equations. The far right hand term of equation (8) is the collision term, which is based on the Bhatnager-Gross-Krook (BGK) approximation [17], where the distribution function is close to a local equilibrium and relaxes to this value with some characteristic time, τ . The term involving $F_i = \vec{e}_i \cdot \vec{F} / 2$ is the body force term in the momentum equation due to buoyancy based on the low Mach number Boussinesq approximation, $\vec{F} = -\vec{g} \beta (T - T_{ref})$, where T is the fluid temperature [18]. The local equilibrium, f_i^{eq} , for each velocity direction is based on a local Mach number expansion of the Maxwell-Boltzmann distribution up to second order in the fluid velocity, \vec{u} , and is given by the following equation,

$$f_i^{eq}(\vec{x}, t) = \rho w_i \left(1 + 3(\vec{e}_i \cdot \vec{u}(\vec{x}, t)) + \frac{9}{2}(\vec{e}_i \cdot \vec{u}(\vec{x}, t))^2 - \frac{3}{2}(\vec{u}(\vec{x}, t) \cdot \vec{u}(\vec{x}, t)) \right) \quad (8)$$

where w_i are weights associated with the lattice and the macroscopic values of density, ρ , and velocity, \vec{u} , are obtained from the following equations

$$\rho(\vec{x}, t) = \sum_0^{18} f_i(\vec{x}, t) , \quad \rho\vec{u}(\vec{x}, t) = \sum_0^{18} \vec{e}_i f_i(\vec{x}, t) + \frac{\Delta t}{2} \vec{F} \quad (9)$$

The fluid kinematic viscosity, ν , is fixed based on the collision relaxation time, τ , the lattice space, Δ , and the time step, Δt , and is given as

$$\nu = \frac{\Delta^2}{\Delta t} \left(\frac{\tau}{2} - \frac{1}{2} \right) \quad (10)$$

Equations (9) are moments of the distribution function and as is common in kinetic theories, it is possible to obtain higher order moments to include the thermal aspects of the fluid flow. However, thermal fluctuations in reality require many more particle velocities and the D3Q19 would therefore be insufficient. But, in data center airflows, the viscous heat dissipation and compressive work done by the pressure are negligible so that the temperature field is passively advected by the fluid and obeys a simple advection-diffusion equation. There are two distinct approaches for thermal LBM models, higher order (multispeed) lattice stencils [19] or double distribution function models [20]. In this work the latter is used and the temperature equation can be solved on a D3Q6 lattice using an independent set of temperature distribution functions, T_j .

$$T_j(\vec{x} + \vec{c}_j \Delta t, t + 1) - T_j(\vec{x}, t) = - \frac{T_j(\vec{x}, t) - T_j^{eq}(\vec{x}, t)}{\tau_T} \quad (11)$$

where the subscript j here is over the 6 lattice directions with particle velocities, \vec{c}_j , of the D3Q6 lattice. As before, the left hand side of equation (11) captures the advection of temperature and the right hand side its diffusion based on the relaxation time, τ_T , which links to the diffusion coefficient for the energy equation.

$$\frac{k}{\rho c_p} = \frac{(2\tau_T - 1) \Delta^2}{4 \Delta t} \quad (12)$$

and following [21], the local temperature and the equilibrium temperature distribution are respectively given by

$$T(\vec{x}, t) = \sum_1^6 T_j(\vec{x}, t) ; \quad T_j^{eq}(\vec{x}, t) = \frac{T(\vec{x}, t)}{6} (1 + 3\vec{c}_j \cdot \vec{u}(\vec{x}, t)) \quad (13)$$

The modelling of thermal airflows in data centers requires a model to capture the effects of turbulence. For the application of the LBM in this work a large eddy simulation (LES) turbulence model is used that makes use of the fact that the local momentum stress tensor, $S_{\alpha\beta}$, is directly available at each time step and does not require the solution of additional governing equations as in the RANS application for FEM and FVM, which helps to speed up the computation by virtue of the fact that [22]

$$S_{\alpha\beta} = \frac{1}{2} \left(\frac{\partial u_\alpha}{\partial x_\beta} + \frac{\partial u_\beta}{\partial x_\alpha} \right) = \sum_0^{18} e_{i\alpha} e_{i\beta} (f_i - f_i^{eq}) \quad (14)$$

The Smagorinsky turbulence model can then be used to calculate the local kinematic turbulent viscosity, ν_t , using

$$\nu_t = \frac{1}{6} \left(\sqrt{v^2 + 18C_s^2 \Delta^2 \sqrt{S_{\alpha\beta} S_{\alpha\beta}}} \right) \quad (15)$$

where the constant $C_s > 0$, and essentially dictates the impact of the sub-grid eddies on the fluid mechanics. The incorporation of the Smagorinsky LES model requires varying relaxation times, τ and τ_T , of the double distribution functions to be adjusted to give $\tau = \frac{1}{2} + 3(\nu + \nu_t)$ and $\tau_T = \frac{1}{2} + 2 \left(\frac{k}{\rho c_p} + \frac{\nu_t}{Pr_t} \right)$, where Pr_t is the turbulent Prandtl number taken as 0.9 in the simulations presented in this work.

3. Materials and Methods

3.1 Data Centre Geometry

A benchmark case study is explored in this work, in which the geometry of a data center with the necessary components is designed to allow for a simple implementation for each of the numerical methods considered while keeping most of the characteristics of a typical data center in operation. The data center floor space has an area of 28.8m^2 , and houses ten racks ($0.6\text{m} \times 1\text{m} \times 2\text{m}$ in size) organised in two rows, see Figure 5. The datacom units in the racks have the fronts facing each other across the cold aisle of the center. Each rack is separated into two datacom units along the height, so that the digital workload between the top half and bottom half of the rack can be allocated different values. This is a simplification as a rack would normally be composed of many more datacom systems. A computer room air conditioner (CRAC) unit is placed in line with the cold aisle at one end of the data center, and feeds cool air into the plenum under the floor. An illustration of this is given in Figure 5 and a more detailed schematic is provided in Figure 6.




Figure 5 – Schematic of the benchmark case study data center in cooling operation. Temperatures and airflows indicated by the colour (warm – red; cool - blue) and direction of the arrows.




Figure 6 – Plan and front view with dimensions of the benchmark case study data center

Cold air is supplied to the room at a constant temperature T_{CRAC} and a constant flow rate Q_{CRAC} through the bottom of the CRAC unit, which coincides with the top of the plenum as indicated in the schematic of Figure 6. Cold air travels in the plenum and then enters the room through the floor vents, modelled here as open holes and filling the floor space between the two rows of racks. This is a simplified model, as in reality flow through the plenum is affected by legs supporting the floor and by cables and pipes running through the plenum. Note also that floor grilles used in real data centres contain detailed features that straighten the flow and affect its momentum, and they can sometimes be oriented to change the angle of the flow to provide the cool air where it is required. The air enters each datacom unit through their front-faces and exits from their back-faces after an increase in momentum and a rise in temperature as heat is removed from the datacom microprocessors by the flow of air. The volume flow rate through each datacom system Q_{datacom} and the temperature rise $\Delta T_{\text{datacom}}$ are dependent on the datacom unit's power draw P_{datacom} as described by Eq. (16),

$$P_{\text{datacom}} = \rho c_p Q_{\text{datacom}} \Delta T_{\text{datacom}} \quad (16)$$

where $\rho = 1.23 \text{ kg/m}^3$ and $c_p = 1.01 \text{ kJ/(kg.K)}$ are the density and specific heat capacity at constant pressure of air respectively. Eq. (16) results from the conservation of energy transferred from the datacom microprocessors to the air flow by convection. Therefore knowing the datacom unit's power demand and air flow rate yields a value of temperature difference between the front and the back of the unit. The inlet to each datacom system is treated as an outlet of the data center domain and the outlet from each datacom system an inlet to the data center domain.

3.2 Boundary Conditions

The CRAC unit is responsible for removing hot air from the room and providing cool air to the room. The CRAC exhaust is treated as a flow domain inlet providing air at a uniform temperature of 289 K and a flow rate of 1.134 m³/s. The velocity profile across the exhaust is uniform, and is directed vertically downwards into the plenum. The CRAC intake is treated as a domain outlet and uses a Dirichlet boundary condition for pressure, $p = 0$, and von Neumann (zero gradient) boundary conditions for the velocity, temperature and turbulence fields. The turbulence variables (k - ϵ) for the FVM and FEM approaches are specified using a turbulent intensity of 5% and a mixing length scale of 0.14 m (7% of the CRAC characteristic length scale).

In this simple study all datacom casings and the CRAC unit panels which are not intake or exhaust components are treated in the same manner as the boundary of the domain, namely the data center walls. It is assumed that no heat is transferred into or out of the domain through these boundaries, therefore adiabatic boundary conditions are prescribed for the thermal boundary conditions. Boundary conditions for velocity at the walls are treated as no-slip, and for pressure, zero gradient boundary conditions are specified. For the turbulence fields in both FEM and FVM standard k - ϵ wall functions have been used as detailed in [23] and [24] respectively.

When a datacom unit is in operation air enters it through its front-face and exits from the back-face after an increase in momentum and a rise in temperature. The inlet to each datacom system is treated as an outlet of the fluid domain and the outlet from each datacom system is treated as an inlet to the fluid domain. However, the temperature rise in each datacom system, $\Delta T_{datacom}$, is dependent on the datacom unit's power draw, $P_{datacom}$ and air flow rate, $Q_{datacom}$, via equation (16). In this model, flow rate and temperature are linked using the following equation, from Summers et al. [25]:

$$T(\vec{x}_{rear}, t) = \frac{1}{A_{front}} \int_{A_{front}} T(\vec{x}_{front}, t) dA + \frac{P_{datacom}}{\rho c_p Q_{datacom}} \quad (17)$$

The power distribution across the datacom units for this particular study is specified in Table 1 below and the positioning of the datacom equipment in the data center is labelled according to Figure 7. For units which draw zero power, no flow passes through them and the boundaries are treated as impermeable walls.




Figure 7 – Layout of datacom units as half rack units.

Left Row				Right Row			
Bottom		Top		Bottom		Top	
1	2.0 kW	6	0.0 kW	11	0.0 kW	16	4.0 kW
2	1.0 kW	7	1.0 kW	12	2.0 kW	17	0.0 kW
3	2.0 kW	8	2.0 kW	13	1.0 kW	18	1.0 kW
4	0.0 kW	9	0.0 kW	14	0.0 kW	19	1.0 kW
5	1.0 kW	10	1.0 kW	15	0.0 kW	20	1.0 kW

Table 1 – Power draw of each half rack datacom equipment by distribution.

The power draw of each datacom unit (two per rack) is used to define the linked thermal boundary condition between the front and back of each datacom unit based on equation (17). For each datacom unit, velocity is specified by assuming a uniform distribution of $Q_{datacom}$ over the outward facing area, given by points \vec{x}_{front} . Flow is positive into the domain and negative leaving the domain. For power draws of 4 kW, 2 kW and 1 kW the datacom volumetric flow rates are specified as 0.221 m³/s, 0.133 m³/s and 0.083 m³/s respectively. On both sides of the datacom units zero gradient boundary conditions are specified for pressure. The k- ϵ turbulence variables are given in both the FEM and FVM for each datacom by using a turbulent intensity of 5% and mixing length scale of 0.07 m (7% of the datacom characteristic length scale).

3.3 Solution Procedures

3.3.1 Finite Element Method

Comsol Multiphysics has its own built-in meshing tool to construct the Finite Element mesh for the data center geometry. Second order tetrahedral type elements were specified and for the final mesh minimum and maximum length scales of 0.015 m and 0.15 m were chosen respectively. The meshing algorithm places more elements in near wall regions, subsequently the size of elements grows towards regions of empty space. The problem is setup using the incompressible solver with the Boussinesq approximation to account for the buoyancy forces. The effect of mesh refinement on the solution is determined by examining the CRAC return temperature and two datacom inlet temperatures as functions of the number of finite elements. The result of this mesh refinement study is shown in Figure 8 along with the time taken to obtain the solution.




Figure 8 – Mesh refinement on the calculated boundary temperatures at certain points in the data center based on the FEM.

3.3.2 Finite Volume Method

The computational mesh used for the FVM was created using the tool `snappyHexMesh` which is a standard




Figure 9 – FVM mesh refinement on the calculated boundary temperatures at certain points in the data center.

meshing tool provided with OpenFOAM. The initial computational mesh was comprised of $80 \times 46 \times 30$ uniformly spaced control volumes created using `blockMesh`. To capture wall effects, `snappyHexMesh` was then used to include two levels of surface refinement around each component in the domain. The resulting mesh was composed of cells with a minimum volume of $1.95499 \times 10^{-6} \text{ m}^3$ and a maximum volume of $9.16675 \times 10^{-4} \text{ m}^3$. Further details of the meshing procedure for the data center case study can be found in Summers et al. [25]. In order to ensure an adequate mesh density, the CRAC return temperature and two datacom inlet temperatures are examined as a function of the number of control volumes. The result of the mesh refinement is shown in Figure 9 along with the computational time, which is seen to increase with the number of control volumes. The simulations are based on the OpenFOAM `buoyantBoussinesqSimpleFoam` solver, with the addition of a new boundary condition for the temperature field that incorporates equation (17).

3.3.3 Lattice Boltzmann Method

The implementation of the LBM approach to the simulation of thermal airflow in the data center case study under investigation here is very different from the FEM and FVM. Firstly the LBM is an inherently time-dependent solver of the Navier-Stokes equations, so for comparison in this work it is important to first integrate to steady-state, but with the LES turbulence model applied to the LBM there is still fluctuation in the macroscopic field variables. A detailed discussion of the LBM implementation is found in Delbosc et al. [26] which uses graphical processing units (GPUs) rather than central processing units (CPUs) used by the FEM and FVM. The LBM is a local based method and is amenable to the vector-like processing on GPUs. Since the computation is bound to the GPU, it is possible to visualise the fluid flow directly through the graphics interface, enabling a real-time CFD approach as depicted in Figure 10.




Figure 10 – Contrasting the solution procedures of traditional CFD on the left (FEM and FVM) with real-time CFD on the right (LBM).




Figure 11 – Constraints on the real-time rate as a function of the number of nodes in the lattice.

The LBM simulations of the case study data center are run for 180 seconds of real time and then the macroscopic temperature is averaged over a further 120 seconds of real time. The number of nodes for the LBM simulations is

chosen so that the real-time rate is greater than or equal to 1, see Figure 11. This rate depends on some physical parameters of the problem, namely a characteristic length scale, L , and a ratio of the characteristic velocity to the lattice Boltzmann velocity, C_v , the latter of which must be kept small to reduce compressibility errors that are inherent in the LBM as applied here [27]. To contrast the computational time of the LBM running on GPUs, the real-time rate was kept at around unity, so that for the 300 seconds of real time equates to approximately 300 seconds of computational time and a lattice with up to a million nodes.

Each computational approach has its own performance constraints, levels of accuracy, stability and convergence conditions, where for FEM and FVM a tolerance is imposed on the discretised equation residuals. The FEM solution requires 403,825 elements to meet a convergence criterion of 10^{-5} and took 52,814 seconds to compute on 4 cores of an Intel i7-860S, 2.53GHz CPU. The FVM solution requires 1,303,093 cells to meet a convergence criterion of 10^{-5} and took 19,996 seconds to compute on a single core of an Intel Xeon E5-2620, 2.0GHz CPU. The LBM solution on the other hand computes on a regular grid composed of 753,984 lattice points with a time-step of 0.005s and takes 300 seconds to compute (i.e. real-time) on an nVIDIA Tesla K40. By considering the thermal design power of each compute engine (CPU or GPU) and the simulation time for each method, the energy costs for each simulation method were calculated to be 1.18kWh, 0.53kWh and 0.05kWh for FEM, FVM and LBM respectively. For the inherently transient LBM approach the time to solution is based on 300 seconds to remove simulation transience and a further 300 seconds to obtain a time averaged steady state solution.

4. Results & Discussion

4.1 Thermal Airflow Distributions

The results of the simulations from the three computational approaches, namely FEM, FVM and LBM, are presented graphically by taking cross sections through the data center in 3 perpendicular planes. Each plot shows the temperature field in Kelvin and the airflow is represented qualitatively by velocity vectors on the same plot.




Figure 12 – Cross section plane through hot and cold aisles at 3.4m from the end of the data center, $y = 3.4$. Results show temperature and air velocity for (a) FVM, (b) LBM and (c) FEM. Temperature scales are in Kelvin

Figure 12 shows a cross section cutting through the two hot aisles and the single cold aisle, at a plane located at $y=3.4\text{m}$ (see Figure 6 for schematic). The numerical solution from all three methods demonstrates one of the most common flows found in data centers, that is where the hot air exiting from datacom units into the hot aisles is drawn into the cold aisle as a result of insufficient cold air being supplied to the cold aisle. The direct effect of this is that the air entering the datacom equipment mixes with hot recirculating air.

Figure 13 shows a cross section that runs down the middle of the data center through the cold aisle. The results show how the air is fed into the cold aisle from under the floor and also how the hot air returns back to the CRAC. The results clearly demonstrate the necessity for floor tiles that can straighten the flow and angle it directly upwards for delivery in front of the datacom units. All three approaches depict the momentum of the airflow directing cool air away from the racks closest to the CRAC. Figure 13 also clearly shows the path of the hot air returning back to the CRAC.




Figure 13 – Cross section plane down the middle of the cold aisle at 2.4m from the side of the data center, $x=2.4$. Results show temperature and air velocity for (a) FVM, (b) LBM and (c) FEM. Temperature scales are in Kelvin.

Figure 14 shows a horizontal cross section through the data center at a height of 1.5m above the floor, which is above the plenum. Here it is possible to see how the airflow patterns carry the hot air from the end of the hot aisles back to the CRAC unit. It is also possible to see bypass airflow at the far end of the two rows, at the opposite end of the data center from the CRAC unit. Bypass air is cool air supplied by the CRAC that bypasses the inlets of the datacom units, traversing directly into the hot aisles. It is also possible to see the recirculating air, where the hot exhaust from the datacom units of the left row is drawn around the front of the left row directly into the inlet of the datacom equipment at the front of the right row.

The results laid out in Figures 12, 13 and 14 from the different simulation technologies are in qualitative agreement as seen graphically by the temperature profiles and velocity vectors clearly depicting recirculating and bypass airflows in the same regions. There is clearly a compelling case to enclose either the cold aisle or the hot aisles, both of which are becoming a common practice [28].




Figure 14 – Cross section horizontal plane at 1.5m from the data center floor, $z = 1.5$. Results show temperature and air velocity for (a) FVM, (b) LBM and (c) FEM. Temperature scales are in Kelvin.

4.2 Datacom Inlet Temperatures

The three methods can be quantitatively assessed by comparing the airflow temperatures at different points within the simplified data center layout. A numerical comparison of the average datacom inlet temperatures is a strong test of the agreement between the predictions of the 3 modelling methods.

Figure 15 shows the average temperature at each of the datacom inlets across the three computational approaches for all 20 half rack units. For some datacom inlets all three methods give the same temperature to within 1K, but for inlets of datacom systems towards the center of the aisle, there is up to a 4K difference between the hottest and coldest prediction. Overall, the predicted temperature at the datacom inlets agree well, but the LBM appears to give consistently a slightly higher temperature (by about 1K) than the other methods. This may be a side effect of using an LES turbulence model as opposed to the RANS model used in both the FEM (Comsol) and FVM (OpenFOAM), but is also likely to be affected by the choice of the initial temperature field, which has not been explored in this work.

The average difference in inlet temperatures between the three computational methods is found to be 0.53K, 1.27K and 0.76K for FVM-FEM, FVM-LBM and FEM-LBM respectively. While the LBM is the fastest method for the given resolution and time-step, it can display spurious fluctuations that appear as a checkerboard pattern and are due to numerical instabilities inherent to the method. However, with a resolution of 4 Kelvin for inlet temperature predictions, the numerical approach can provide very fast response to data center design and operation issues.




Figure 15 – Inlet temperatures of datacom equipment based on their position in the data center (for the ID refer to Figure 7) as calculated from the three computational methods.

5. Conclusions

The simulation results of thermal air flows in a simplified data center obtained with three computational methods using different meshing technologies and turbulence models demonstrate good agreement in terms of the overall flow structures and the average temperatures at each datacom inlet. The level of simulation accuracy of FVM, FEM and LBM are similar and there are clear advantages of using the LBM in respect to computational performance and applicability for transient flows, which also offers the potential to inform on the dynamic nature of real thermal air management of data centers.

This work has also demonstrated the computational performance of both the FEM and FVM as well as identifying further aspects of the modelling approach and assumptions that position these computational techniques as valuable tools for analysing thermal air management of data centers.

In summary, this problem has been analysed using a commercial FEM package that offers ease of setting up the problem, but uses the largest amount of computational time. An opensource FVM package required more time input in setting up the problem, but solves the problem with less computational time than the FEM package. The LBM code developed by the authors and that operates on GPUs is more challenging than both FEM and FVM to setup this problem, but offers the fastest time to compute.

Acknowledgments

The authors would like to thank 4Energy for both inspiration and part funding of this research and Bios IT for the loan of the direct liquid cooled K40 GPUs. We thank the European Union FP7 Regions of Knowledge (Grant 320013) for funding Greg de Boer. Morgan Tatchell-Evans would like to thank the Engineering and Physical Sciences Research Council for funding (Grant EP/G036608) and Dan Burdett acknowledges the funding from aql.

References

1. Mitchell-Jackson, Jennifer, Jonathan G. Koomey, Michele Blazek, and Bruce Nordman. "National and regional implications of internet data center growth in the US." *Resources, conservation and recycling* 36, no. 3 (2002): 175-185.

2. J. Glanz, "THE CLOUD FACTORIES - Power, Pollution and the Internet," *The New York Times*, 22 9 2012. [Online]. Available: http://www.nytimes.com/2012/09/23/technology/data-centers-waste-vast-amounts-of-energy-belying-industry-image.html?pagewanted=all&_r=1&.
3. Garimella, Suresh V., et al. "Technological drivers in data centers and telecom systems: Multiscale thermal, electrical, and energy management." *Applied energy* 107 (2013): 66-80.
4. Koomey, J., 2011. Growth in data center electricity use 2005 to 2010. A report by Analytical Press, completed at the request of The New York Times, p.9
5. Liam Newcombe, Mark Acton, John Booth, Sophia Flucker, Paul Latham, Steve Strutt, and Robert Tozer. 2012 Best Practices for the EU Code of Conduct on Data Centres. European Commission, Joint Research Centre, 2012.
6. Yogendra Joshi and Pramod Kumar. Introduction to Data Center Energy Flow and Thermal Management. In Yogendra Joshi and Pramod Kumar, editors, *Energy Efficient Thermal Management of Data Centers*, pages 1–38. Springer US, 2012
7. Emad Samadiani, Jeffrey Rambo, and Yogendra Joshi. Numerical modeling of perforated tile flow distribution in a raised-floor data center. *Journal of Electronic Packaging*, 132(2):021002, 2010.
8. Jayantha Siriwardana, Saman K Halgamuge, Thomas Scherer, and Wolfgang Schott. Minimizing the thermal impact of computing equipment upgrades in data centers. *Energy and Buildings*, 50:81–92, 2012.
9. Kyosung Choo, Renan Manozzo Galante, and Michael M. Ohadi. Energy consumption analysis of a medium-size primary data center in an academic campus. *Energy and Buildings*, 76(0):414 – 421, 2014.
10. S.K. Shrivastava, B.G. Sammakia, R. Schmidt, M. Iyengar, and J.W. VanGilder. Experimental-Numerical Comparison for a High-Density Data Center: Hot Spot Heat Fluxes in Excess of 500 W/FT². In *Thermal and Thermomechanical Phenomena in Electronics Systems, 2006 ITherm '06. The Tenth Intersociety Conference on*, pages 402–411, May 2006.
11. Alissa, H. A., K. Nemati, B. Sammakia, K. Ghose, M. Seymour, and R. Schmidt. Innovative approaches of experimentally guided CFD modeling for data centers. In *Thermal Measurement, Modeling & Management Symposium (SEMI-THERM), 2015 31st*, pp. 176-184. IEEE, 2015.
12. Gray, D.D. and Giorgini, A. The validity of the Boussinesq approximation for liquids and gases. *International Journal of Heat and Mass Transfer*, 19(5), pp.545-551, 1976.
13. Karki, K.C., Radmehr, A. and Patankar, S.V. Use of computational fluid dynamics for calculating flow rates through perforated tiles in raised-floor data centers. *HVAC&R Research*, 9(2), pp.153-166, 2003.
14. Jacobs, C.T., Jammy, S.P. and Sandham, N.D. OpenSBLI: A framework for the automated derivation and parallel execution of finite difference solvers on a range of computer architectures. *Journal of Computational Science*, 18:12-23, 2017.
15. Weller, H. G., Tabor, G., Jasak, H., & Fureby, C. A tensorial approach to computational continuum mechanics using object-oriented techniques. *Computers in physics*, 12(6), 620-631, 1998
16. X. He and L.S. Luo. Lattice Boltzmann Model for the Incompressible Navier-Stokes Equation. *Journal of Statistical Physics*, 88(3-4):927–944, 1997.
17. Prabhu Lal Bhatnagar, Eugene P Gross, and Max Krook. A model for collision processes in gases. I. Small amplitude processes in charged and neutral one-component systems. *Physical review*, 94(3):511, 1954
18. Zhaoli Guo, Baochang Shi, and Chuguang Zheng. A coupled lattice BGK model for the Boussinesq equations. *International Journal for Numerical Methods in Fluids*, 39(4):325–342, 2002.
19. Pavlo, Pavol, George Vahala, and Linda Vahala. Preliminary results in the use of energy-dependent octagonal lattices for thermal lattice Boltzmann simulations. *Journal of statistical physics* 107, no. 1 (2002): 499-519.
20. Guo, Zhaoli, Baochang Shi, and Chuguang Zheng. A coupled lattice BGK model for the Boussinesq equations. *International Journal for Numerical Methods in Fluids* 39, no. 4 (2002): 325-342.
21. Du Hong-Yan, Chai Zhen-Hua, and Shi Bao-Chang. Lattice Boltzmann study of mixed convection in a cubic cavity. *Communications in Theoretical Physics*, 56(1):144, 2011
22. S. Hou, J. Sterling, S. Chen, and G.D. Doolen. A Lattice Boltzmann Subgrid Model for High Reynolds Number Flows, volume 6 of *Fields Institute Communications*, pages 151–166. AMS, Providence, 1996.
23. Dick, E. Introduction to Finite Element Methods in Computational Fluid Dynamics. *Computer Fluid Dynamics*, Wendt, J.F. (Ed.), Springer, pp. 235-274, 2009.
24. Versteeg H.K. and Malalasekera, W. An introduction to Computational Fluid Dynamics the Finite Volume Method, 2nd Ed., Prentice Hall, 2007.
25. Summers, Jonathan, Nik Kapur, and Harvey Thompson. "Design of data centre rack arrangements using open source software." *Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2013 29th Annual IEEE*. IEEE, 2013.
26. Delbosc, N., Summers, J.L., Khan, A.I., Kapur, N. and Noakes, C.J. Optimized implementation of the Lattice Boltzmann Method on a graphics processing unit towards real-time fluid simulation. *Computers & Mathematics with Applications*, 67(2), pp.462-475, 2014
27. Delbosc, N. Real-Time Simulation of Indoor Air Flow using the Lattice Boltzmann Method on Graphics Processing Unit (Doctoral dissertation, University of Leeds), 2015.

28. Niemann, J., Brown, K. and Avelar, V. Impact of hot and cold aisle containment on data center temperature and efficiency. Schneider Electric Data Center Science Center, White Paper, 135, pp.1-14, 2011.