Distrib Parallel Databases (2007) 21:193–225
DOI 10.1007/s10619-006-7007-3

Modeling of data center airflow and heat transfer: State of the art and future trends

Jeffrey Rambo · Yogendra Joshi

Published online: 19 January 2007

Abstract An assessment of the current thermal modeling methodologies for data centers is presented, with focus on the use of computational fluid dynamics and heat transfer as analysis tools, and model validation. Future trends in reduced or compact modeling of data center airflow and heat transfer are presented to serve as an overview of integrating rack-level compact models into full-scale facility level numerical computations. Compact models can be used to efficiently model data centers through varying model fidelity across length scales. Dynamic effects can be included to develop next-generation control schemes to maximize data center energy efficiency.

Keywords Thermal modeling · Data center · Reduced order models

1 Introduction

Data centers are computing infrastructure facilities utilized to provide a consistently reliable operating environment for servers, storage and networking devices, generically referred to as data processing equipment (DPE). The continuous operation of the DPE is critical to industries data centers serve, such as telecommunications, banking, and high-performance scientific computing. Thermal management issues associated with continually increasing power dissipation from DPE are further compounded by the vertical stacking of such equipment in 2-m tall racks in the data center environment. In 1990, a typical rack with a 1.07 m by 0.61 m (42 in by 24 in) footprint dissipated

Recommended by: Monem Beitelmal

J. Rambo
Shell Global Solutions (US), Inc. Houston, TX 77082

Y. Joshi (✉)
G. W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology,
Atlanta, GA 30332
e-mail: yogendra.joshi@me.gatech.edu

approximately 1 kW of power [1], while today's racks with the same footprint may dissipate up to 30 kW, based on current server heat loads. The large demand for computing power has caused data centers to grow as large as 9300 m^2 (∼100,000 ft^2), incorporating thousands of racks, and producing net facility power dissipation rates on the order of several MW. Data center practitioners measure power density as a heat flux based on the net power dissipated over the total footprint of the facility. Recent energy benchmarking studies have shown data centers are operating in the 270–800 W/m^2 (25–75 W/ft^2) range [2], and growth to 100 W/ft^2 average over the facility, with local regions exceeding 200 W/ft^2 is expected in the near future. These power densities are well beyond conventional HVAC loads for the cooling of similarly sized rooms such as auditoria and office spaces, which are typically 4–8 W/ft^2. Of the 11 different data center benchmarked in [2], the computing load consumed between 40–70% of the total power consumed and averaged 52%. Data center cooling loads including the chiller operation, chilled water pumping power and air conditioning unit blower power consumed an additional 34% on average. The remaining 14% can be attributed to power conversion losses, lighting and other operational overhead. Another benchmarking study showed that the computing, CRAC and chiller plant energy usage are roughly 1/3 of the total power consumption [3].

Data center floor space comes at a premium, with construction cost estimates as high as \$15,000/m^2 in metropolitan areas and annual operating costs ranging between \$500–\$1500/m^2 [4]. High-power dissipation racks would ideally be spread out to minimize the local power density, however the high construction costs requires data centers to maximize the floor space utilization by DPE. These cost constraints have caused data center designers and operators to pack as much computing power into the available space as possible, which combined with the DPE manufacturer trend of providing as much functionality into their equipment as possible, are causing severe thermal management issues.

The predominant cooling scheme for current data centers is to use the computer room air conditioning (CRAC) units to supply a raised floor plenum (RFP) underneath the racks with cold air. Perforated tiles are then located near the racks to deliver the cool supply air to the front of the DPE. The hot exhaust air from the racks is then collected from the upper portion of the facility by the CRAC units, completing the airflow loop, as shown in Fig. 1. Racks are typically arranged in rows with alternating

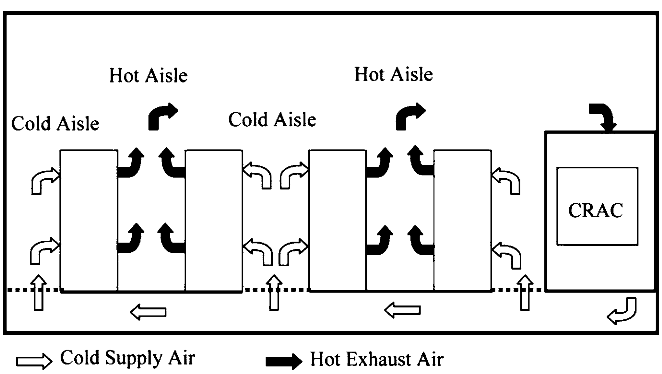

Fig. 1 Standard raised floor plenum and room return cooling scheme airflow schematics

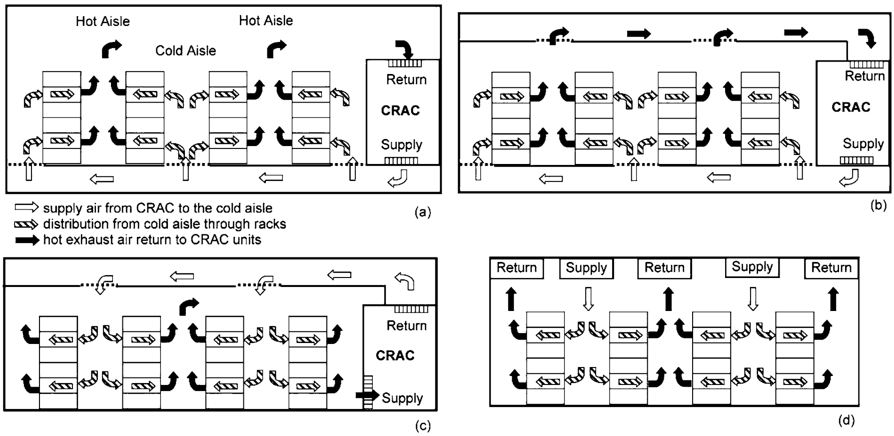

Fig. 2 Alternative cooling schemes employing (a) raised floor plenum (RFP) supply and standard return (SR), (b) RFP supply and overhead (OH) return, (c) OH supply with room return (RR) and (d) OH supply and return

airflow directions, forming 'hot' and 'cold' aisles [5], with the perforated tiles located in the cold aisles where the rack inlets face each other and the exhaust air is collected in the hot aisles. This hot aisle – cold aisle approach attempts to separate the supply from the exhaust air and increase the overall efficiency of the air delivery and collection from each rack in the data center.

Improved thermal efficiency may be possible with alternative cool air supply and hot exhaust return configurations. While any location for the CRAC unit supply and return can be modeled, only certain combinations are feasible due to mechanical constraints of the CRAC units without introducing an excessive amount of ductwork and are illustrated in Fig. 2.

2 Data center modeling objectives

Access to operating data centers is limited due to their high reliability constraints, and the large variations in data center architectures limit the extendibility of measurements to other facilities. As a result, most previous investigations on data center thermal characterization have involved computational fluid dynamics and heat transfer (CFD/HT) to predict airflow and heat transfer characteristics. A reliable simulation methodology for data centers allows thermal designers to identify potentially dangerous local hot spots, quickly evaluate alternatives in cooling systems and investigate next generation data center architectures.

Accurate data center airflow and heat transfer modeling can aid in the design of new facilities for energy efficiency, instead of drastic over-provisioning currently used to guard against power density increases. A lifecycle mismatch problem also occurs because most data center facilities are designed for 15–20 years of use [4], but the DPE becomes outdated in less than 2 years. With continual upgrades, the optimal arrangement of new DPE with higher power dissipation and flow rate requirements needs to be determined to mitigate its effect on neighboring DPE. Additional constraints imposed by the grouping of DPE by functionality and cabling requirements often conflict with thermal management strategies and data center managers need to provide localized supplemental cooling to high power racks [6].

The first published results for data center airflow modeling appeared in 2000 [7]. Since then, various modeling investigations have been presented, many employing CFD/HT to simulate data center airflow and heat transfer characteristics. The objective of this paper is to review data center modeling in an effort to highlight the wide range of techniques and assumptions that have been used in the thermal analysis of data centers. A critical review of the current state-of-the-art can aid in unifying the different modeling approaches and provide future directions for the development of data center thermal analysis tools.

3 Review of numerical modeling

The various modeling efforts have ranged from individual component modeling to rack and power layouts and can be classified into the following main categories:

1. Raised floor plenum (RFP) airflow modeling to predict perforated tile flow rates
2. Thermal implications of CRAC and rack layout and power distribution
3. Alternative airflow supply and return schemes
4. Energy efficiency and thermal performance metrics
5. Rack-level thermal analysis
6. Data center dynamics: control and lifecycle analysis

Previous thermal analysis efforts will be reviewed using the above partitioning and the exact nature of each category will be further explored in the respective subsection. Every attempt has been made to divide the literature as to minimize the overlap between the categories. Works with significant connections to multiple categories will be described in their primary category and noted in subsequent categories. A compendium of data center literature was presented by Schmidt and Shaukatullah [8], which compares the cooling, energy efficiency, humidity and contamination requirements between computer and telecommunication equipment rooms and also serves as a historical perspective considering the rapid growth in data center power density.

3.1 Raised floor plenum (RFP)

With a majority of operational data centers utilizing raised floor plena to deliver cool supply air to the racks, predicting the perforated tile flow rate distribution (PTFRD) is of great concern. RFP data centers almost universally use 0.61×0.61 m (2×2 ft) tiles to form the raised floor. Typical RFP depths vary between 0.305 m (12 in) and 0.914 m (36 in) across facilities with an industry trend to deeper plena in order to supply larger volumetric rates of air more uniformly to cool higher power dissipation DPE. The RFP is commonly used to route the various power and data cables as well as chilled water lines for the CRAC units, which can create substantial blockages and severely alter the PTFRD. Experimental measurements and CFD prediction of velocities and pressure distributions in RFP have shown that the airflow patterns can be very complex, especially when the exhausts from multiple CRAC units interact [9].

The plenum pressure distribution is important because pressure differences across the perforated tiles drive the flow from the RFP into the facility. In regions of high velocity, the static pressure may be lower that the facility pressure, causing reversed flow from the facility down into the plenum. Reversed flow commonly occurs when the perforated tiles are located in close proximity to the CRAC units where the blowers create high velocity exhaust jets. Note all investigations in this section consider the flow in the RFP only and no DPE equipment above the raised floor was modeled. Not modeling the DPE above the RFP means the back pressure for each tile needs to be specified and a constant value for all tiles has been prescribed in all the previous works. The pressure drop across the perforated tiles are much greater than any other pressure drop in the flow path and the pressure variations above the raised floor are very small relative to the pressure variations in the RFP, such that the pressure on the facility side of the perforated tile can be treated as a constant for all practical purposes.

The first modeling of data center RFP was performed by Kang et al. [7] who used flow network modeling [10] and CFD to predict perforated tile flow rates. Their design study showed that the flow distribution can be altered by varying the percent open area of the tiles. With no obstructions in the plenum, the only airflow resistance in the loop is caused by the pressure drop across the perforated tiles, making the assumption of a uniformly pressurized plenum valid. Specific details of each major investigation cited can be found in Table 2.

Schmidt et al. present a numerical model for a 0.284 m (11.2 in) deep plenum [11] and for a 0.292 m (11.5 in) deep plenum [12] for a 6.10 m $\times$ 20.1 m (66 ft $\times$ 20 ft) data center. The facility contains one CRAC unit near an artificially created wall and another unit located 15 tiles away in the middle of the plenum. Some of the models used turning vanes to direct the CRAC unit exhaust in various directions. A 2-dimensional (depth-averaged) computational model was developed and comparison with measurements showed the model predicts the general PTFRD trends but fails to capture sharp variations between neighboring tiles and in some cases does not resolve experimentally measured reversed flow. A fully 3-dimensional model of the same facility [12] showed improved agreement with experimental measurements for various perforated tile layouts, but also failed to predict the reversed flow in many tile configurations. In both papers, it is noted that a 0.1 m (3.9 in) diameter pipe and a 0.025 m (1.0 in) tall cable tray are located in the bottom of the plenum, although there is no discussion of how these obstructions are accounted for in the depth-averaged equations. Schmidt et al. [12] state that for plenum depths less than 0.3 m (11.8 m) deep, depth-averaged modeling is an adequate tradeoff between computational cost and accuracy.

Perforated tiles may be located near CRAC units to save data center floor space and the minimum distance between a row of perforated tiles and the front of a CRAC unit was numerically investigated in [13]. Special treatments were introduced to simulate the CRAC unit discharging into an infinitely large plenum without constructing an excessively large numerical model. The results showed that reversed flow may occur up to 4 tiles away from the CRAC unit, in agreement with the experimental measurements of [11, 12].

Radmehr et al. [14] experimentally investigated the leakage flow, or portion of total CRAC airflow that seeps through the seams of the raised floor tiles. Distributing the leakage flow uniformly throughout the perforated tile seams and modeling chilled water supply lines, Radmehr et al. [14] were able to produce predictions with an overall accuracy of 90%. Van Gilder and Schmidt [15] present a parametric study of plenum airflow for various data center footprints, tile arrangements, tile porosity and plenum depth. A vacant plenum was assumed, which makes the perforated tile resistance much greater than any other flow resistance in the plenum and allows the resulting perforated tile flow rate to scale with the net CRAC flow rate.

A systematic description of the boundary conditions, including the modeling of the perforated tiles and CRAC units was presented in [16] which used CFD to model the PTFRD in a large operational data center containing 11 CRAC units and over 200 perforated tiles. The results showed fairly good agreement with measured values and the computational techniques serve as the basis of a commercially available tool [17]. Karki and Patankar [18] present a 1-dimensional model to predict perforated tile flow rates, with somewhat limited applicability, but the analytical technique is able to provide additional insight beyond purely computational results.

3.2 Layout and power distribution

Layout type studies have generally sought optimal dimensions, including the hot and cold aisle spacing and facility ceiling height, to improve the thermal efficiency of standard room-level data center cooling. The relative orientation between the rows of racks and CRAC units has received attention in order to mitigate the effects of recirculation and short-circuiting, or cool supply air returning to the CRAC unit before it passes through any DPE. Another important layout concern is determining the optimal position to locate high power dissipation racks. This issue is continuously facing data center operators, as new equipment is upgraded due to the lifecycle mismatch between the facility and DPE. Locating higher-powered racks is also an important concern in day-to-day operations because data centers do not contain uniformly dissipating racks and even in a facility containing identically configured racks, the power dissipation would vary spatially over time due to different workloads imposed on the DPE.

Patel et al. [19] modeled a RFP data center containing 4 rows of 7 racks, each dissipating 12 kW. Further details regarding the model geometry and boundary conditions are summarized in Table 3. Moving a row of racks to perturb the hot and cold aisle spacing only a few inches showed a 15% change in load on the CRAC units. Rack power dissipation was also varied between rows by a factor of 4, which resulted in 30% variation in CRAC load. Sharma et al. [20] used a slightly different model geometry to evaluate the effect of hot and cold aisle width, top rack to ceiling height and distance between racks and facility walls. The rack to wall distance showed no effect, the ceiling height showed no effect in the room return scenario, but a lower distance between the top of the racks and the ceiling in an overhead return scheme improves the thermal performance. Improved thermal performance was demonstrated when the cold aisle width was increased in the room return case and when the hot aisle width was decreased in the overhead return case.

A novel concept of introducing cold supply air into the hot aisles to mitigate the effects recirculation and created a more uniform temperature throughout the facility was proposed by Schmidt and Cruz [21], although computational results showed this modification could diminish the basic cooling scheme's efficiency. The RFP plenum was not modeled and flow rates and temperatures entering the hot and cold aisles were prescribed as boundary conditions.

To better characterize the performance variations through a data center, a procedure to model individual servers within each rack was developed in [22]. Each rack contained 10 server models, each with their power dissipation and fan model to drive the flow in an attempt to mimic the convective processes of forced air-cooled DPE. The sub-models in [22] used a uniform heat flux surface to model the board-level power dissipation, but more detailed geometry and boundary conditions can be used as computational resources allow.

Each data center has a unique geometrical footprint and rack layout, therefore a common basis is needed to compare the thermal performance of various cooling schemes. A unit cell architecture of a RFP data center was formulated by considering the asymptotic flow distribution in the cold aisles with increasing number of racks in a row [23]. Preliminary analyses showed that at least 4 rows of racks are required to produce the hot aisle – cold aisle behavior, with 1 hot aisle and 2 cold aisles. The rack flow rates were successively halved, resulting in 1.16, 0.59, 0.29 and 0.15 m^3/s (2463, 1240, 620 and 310 CFM). The data show that flow rate of entrained air through the sides of the cold aisle is relatively independent of the number of racks in a row and increases with increasing rack flow rate. The air entrained through the top of the cold aisle increases with both the number of racks in a row and increasing rack flow rate. The unit cell study predicted trends in the distribution of airflow in the cold aisle for increasing number of racks. The findings indicate separate trends for the 'small' rack flow rate limit and the 'large' rack flow rate limit, with the large rack flow rate limit being of higher importance considering the rapid increase in power density. For 9 racks in a row, the portion of entrained air through the sides of the cold aisle has sufficiently approached its asymptotic limit for all rack flow rates, indicating that 9 racks is representative of a long row of racks and captures the behavior of racks at the end and middle of a row adequately for all flow rates. Using 7 racks in a row will also be sufficient to model current high power density data centers because the maximum flow rate of 2463 CFM is extremely high. Air temperature rises of 40 °C are allowable for extremely high-powered racks, and an overall energy balance indicates that each rack could dissipate up to 57 kW, which may be realized in several years if the current DPE heat loads persist at the current rate.

Schmidt and Cruz [24, 25] use the same model as in [21] to investigate the effect of PTFRD on uniformly dissipating racks. The PTFRD was varied by altering the perforated tile open area and RFP depth, resulting in more uniform rack inlet temperature in the vertical direction for specific combinations of tile open area, RFP depth and rack power dissipation.

A parametric study of the effect of ceiling height, RFP depth and perforated tile distance from CRAC on rack inlet temperature was presented by Bhopte et al. [26]. A multi-objective optimization study with equal weighting for the above 3 dimensional parameters showed the minimum rack inlet temperature occurred at the maximum plenum depth (4 ft), maximum ceiling height (20 ft) and perforated tiles located at the median value (8.75 ft), which need to be weighed against the financial constraints on data center space.

3.3 Alternative supply and return configurations

The earliest full-scale CFD/HT modeling of data centers by Patel et al. [1] examined a data center facility using an overhead supply and return air distribution system, that was first suggested in [27]. Comparison with measurements showed an average error of 14% when predicting rack inlet temperature. Shrivastava et al. [28] modeled 7 different permutations of floor, room and ceiling supply and return configurations for same data center geometry and rack layout as [21, 25] for fixed room geometry, uniform rack power and fixed CRAC conditions. Iyengar investigated various geometrical layouts for an overhead supply and room return data center with the same geometry as [21, 25, 28]. The dimensional parameters included supply diffuser height, arrangement and angle, return vent locations and ceiling height, in addition to the rack power dissipation and power distribution.

Various cooling schemes being investigated beyond the standard RFP design and the models presented in [1, 19, 20, 24, 29, 30] have used a variety of orientations between the CRAC units and racks. To develop a mechanistic understanding of convective processes in data centers, the global cooling scheme was divided into the processes of (1) the CRAC exhaust to the cold aisle, (2) cold aisle distribution of supply air through the racks and (3) the hot exhaust air return to the CRAC units [31]. Numerical modeling of various supply and return schemes, coupled with various orientations between the racks and the CRAC units, identified the causes of recirculation and non-uniformity in thermal performance throughout the data center. The parametric study presented in [32] is the first attempt to generally quantify these effects and following work using the same procedures was contributed in [28].

3.4 Energy efficiency and thermal performance metrics

The performance of the assorted data center models is assessed in various ways by different authors, with most authors reporting the maximum inlet temperature to each rack. This makes specific comparisons between various cooling scenarios difficult. Sharma et al. [20] introduced dimensionless numbers to quantify the effects of recirculation. These dimensionless numbers are arrived at by considering the ratios (cold supply air enthalpy rise before it reaches the racks) / (enthalpy rise at the rack exhaust), and (heat extraction by the CRAC units) / (enthalpy rise at the rack exhaust), which in practice require the air temperature evaluation at arbitrary points near the rack inlet and exhaust. Sharma et al. [33] later computed these dimensionless performance metrics in an operational data center by taking a air temperature measurements just below the top of each rack inlet and outlet. Norota et al. [34] used the statistical definition of the Mahalanobis generalized distance to describe the non-uniformity in rack thermal performance. Shah et al. [35–37] proposed an exergy-based analysis method that divides the data center into a series of subcomponents and then CRAC unit operation is optimized, given information regarding the rack power dissipation.

Thermal performance metrics for systems level electronics cooling based on the concept of thermal resistance (power dissipation / temperature rise) were formulated and applied to data centers in [32]. The metrics consider the spatial uniformity of thermal performance to characterize poor designs causing local hot spots. Entropy generation minimization was also proposed as a data center thermal performance metric because poor thermal performance is often attributed to recirculation effects. Since the mixing of hot exhaust air with the supply air in the cold aisles generates entropy, cold aisle entropy generation minimization was employed as a metric. The results presented in [32] show that using entropy generation minimization and thermal resistance with spatial uniformity considerations predict the same design as being the best.

3.5 Rack-level thermal analysis

With hot exhaust recirculation and cold supply air short-circuiting causing major inefficiencies in facility-level cooling schemes, recent trends have focused on rack-level thermal analysis to provide adequate cooling to the DPE. Using the thermal resistance and entropy generation metrics of [32], the optimal arrangement of servers and power dissipation profile for a forced air cooled rack was computationally investigated in [38]. Herrlin [39] also formulated rack-level thermal performance metrics based on rack inlet temperature exceeding desirable limits.

Heydari and Sabounchi [40] incorporated a model of an overhead rack-level supplemental refrigeration device and used CFD to compute the temperature distributions in a RFP data center. The results showed the airflow rate had a greater effect on rack thermal performance than refrigerant temperature or flow rate. Rolander et al. [41, 42] combined the reduced-order modeling approaches of [43, 44] to optimize the layout and inlet airflow rate to a sealed DPE cabinet with a single inlet and outlet. The results showed that rearranging the servers, or equivalently redistributing the heat load, could allow 50% more cabinet power dissipation than simply increasing the cooling supply flow rate.

3.6 Data center dynamics: control and lifecycle analysis

All of the above investigations are concerned with a nearly fixed geometry, number of CRAC units and number of racks, which may represent only a single point in time over a data center's life. To compensate for time varying heat loads, Boucher et al. [50] attempted to control rack inlet temperature by varying CRAC temperature and flow rate, as well as perforated tile open area. Sharma et al. [45] proposed redistributing the computing workload rather than the cooling resources and present CFD-based results and Patel et al. discusses the advantages for distributable cooling resources at various levels [46]. Other investigations have considered combined power and cooling systems and their viability over the facility lifecycle [47, 48].

Table 1 below summarizes data center thermal management studies that are numerical in nature including the use of CFD/HT or other modeling technique that numerically integrates the governing equations of fluid flow and heat transfer. Note some of the literature in Table 1 has been assigned multiple categories and any work comparing numerical simulation with measured values has been noted in the last column. Note the 4th category defined §3 has been split to segregate investigations concerning energy efficiency and the development of thermal performance metrics and the column 'modeling' has been added to catalog papers developing new, innovative or systematic computational approaches.

Table 1 Numerical modeling summary and categorization by scope

Author	Year	Ref	Alternative cooling	Control & lifecycle	Efficiency	Layout	Liquid cooling	Metrics	Modeling	Plenum	Rack	Measurements
Rambo & Joshi	2006	[32]				×		×	×			
Rambo & Joshi	2006	[31]	×						×			
Bhopte et al.	2005	[26]				×						
Herrlin	2005	[39]						×			×	
Iyengar et al.	2005	[29]	×			×						
Radmehr et al.	2005	[14]								×		
Rolander et al.	2005	[41]		×	×				×		×	
Sharma et al.	2005	[45]		×				×				
Shrivastava et al.	2005	[28]	×									
VanGilder & Schmidt	2005	[15]								×		
Heydari & Sabounchi	2004	[40]									×	
Rambo & Joshi	2004	[38]						×			×	
Rambo & Joshi	2004	[13]							×	×		
Schmidt & Cruz	2004	[24–25]				×						
Schmidt et al.	2004	[12]								×		×
Karki et al.	2003	[16]							×	×		×
Rambo & Joshi	2003	[22]				×			×			
Rambo & Joshi	2003	[23]				×			×			
Patel et al.	2002	[19]			×	×						
Schmidt & Cruz	2002	[21]				×				×		
Sharma et al.	2002	[20]				×		×				
Patel et al.	2001	[1]	×									×
Schmidt et al.	2001	[11]								×		×
Kang et al.	2000	[7]							×	×		

4 Model validation

Only limited validation studies of data center CFD/HT models have been performed due to the complexity of data center airflow and heat transfer, which is further compounded by each data center facility having its own unique layout, making a common basis of comparison difficult to achieve. A primary concern in data center CFD/HT validation is the appropriate resolution at which to validate numerical models. Even a small prototypical data center facility would provide an opportunity to take countless point-wise temperature, velocity and pressure measurements. Full-field visualization and measurement techniques, such as particle image velocimetry, tomographic inferometry and laser-induced fluorescence can provide very detailed descriptions of the velocity and temperature fields, although there are many challenges for using such systems in a data center environment beyond the colossal amount of data produced. With either point-wise or full-field measurement techniques, acquiring data at a fine enough resolution to validate CFD/HT models is a considerable undertaking.

4.1 Raised floor plenum (RFP) models

A majority of the results with experimental measurements are RFP investigations that used a flow rate measurement device, as outlined in [11], to measure the net flow through each perforated tile. While it may be the end goal to predict the PTFRD, comparing only the perforated tile flow rate with the CFD simulation does not ensure that the simulation has accurately captured the flow field between the CRAC exhaust and perforated tiles. Comparison with measurements in [11, 12, 15] shows all the CFD predictions have a root mean square error of at least 10% with specific locations exhibiting more than 100% error in some cases. A solution that matches the observed boundary conditions (perforated tile flow rates) does not guarantee that the solution over the domain is correct or even unique. With perforated tile flow rate prediction errors on the order of 10%, the error associated with the velocity field in the RFP may be significantly larger, rendering the physical arguments used to develop RFP design guidelines questionable.

The boundary conditions used to model the CRAC unit blowers supplying the RFP deserve some further investigation. Many of the RFP models fix the CRAC flow rate arguing the pressure drop across the CRAC unit is much greater than the sum of all the other pressure drops in the flow loop, which for a commercially available CRAC unit is approximately 500 Pa ($\sim$2.0 in H_2O). A simplistic flow network analysis of RFP airflow would place multiple CRAC units in parallel causing their aggregate resistance to be the same as a single resistance with their flow rates adding. Multiple perforated tiles are also in parallel, assuming a (unrealistic) uniformly pressurized plenum as in [7], with an average flow rate of 0.165 m^3/s (350 CFM), the aggregate resistance is only 12 Pa (0.05 in H_2O) from the Schmidt et al. correlation [12]. Also note no detailed measurements or model of a CRAC unit that would provide the details of the flow direction and distribution among exhausts of a single CRAC unit has been reported in the literature.

Figure 3 plots the blower curve (Δp [Pa] vs. Q [m^3/s]) for a Lau Industries Model A15–15A blower. The Liebert series FH529C CRAC unit, commonly used in RFP data centers, uses 2 such blowers to create a net flow rate of 5.85 m^3/s (12,400 CFM)

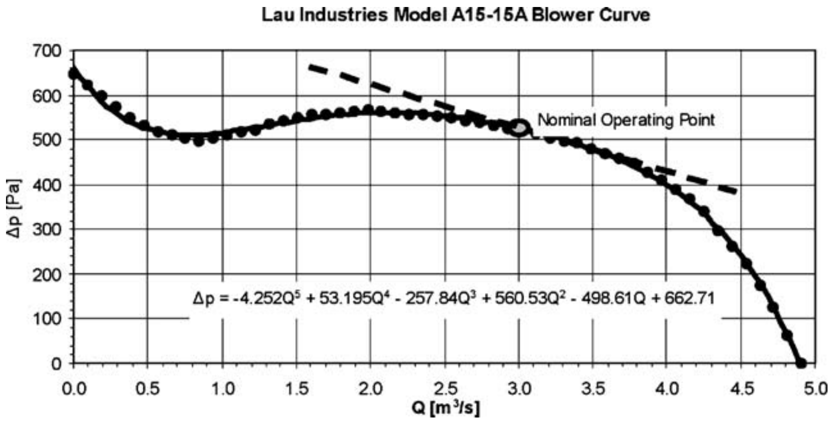

Fig. 3 Typical CRAC unit ΔP [Pa] versus Q [m^3/s] blower curve

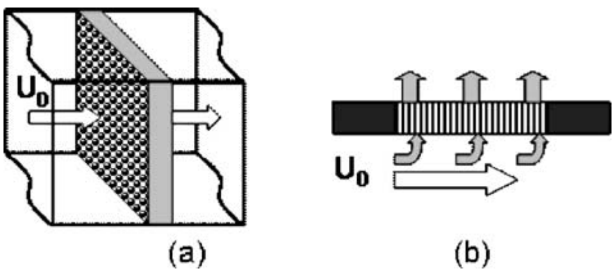

Fig. 4 Misaligned flow through perforated tile

against an external pressure of 75 Pa (0.30 in H_2O). Using a higher-order polynomial fit to differentiate the blower curve at the operating point shows $dp/dQ \approx -68$, which produces large changes in blower flow rate for small changes in system resistance. For example, an obstruction causing a 20 Pa pressure drop will change the flow rate by 0.294 m^3/s or $\sim$10% of the blower flow rate. For data centers with significant blockage effects, the CRAC unit should be modeled using the blower curve characteristics rather than using a fixed nominal flow rate to account for flow rate variations. Modeling the CRAC unit with a blower curve will increase the solution time of a CFD model because the velocity on the CRAC exhaust is computed according to the blower curve rather than being specified explicitly.

Perforated tiles are modeled as a lumped resistance using the relationship $\Delta p = ku^2$ where the coefficient (k) may be found through standard flow resistance handbooks, i.e. [49]. Experimentally determined resistance coefficients have been reported in the literature [12], although they have been determined in the same manner as [49]. Such data have been compiled for flow normal to perforated plates and screens fully spanning the cross section of a conduit, while the flow in a RFP data center is entrained through the perforated tiles rather than being directly aligned with it, see Fig. 4, and the flow straightening effect of the perforated tile is well understood or modeled. Perforated tiles used in RFP data centers are not homogeneous, rather they are a

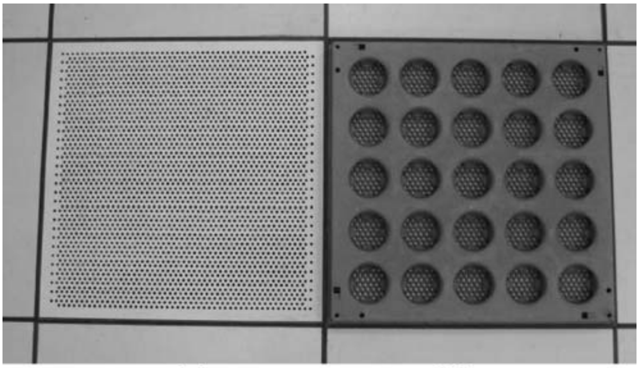

Fig. 5 Photographs for nominal 25% open area tiles, (a) facility side and (b) plenum side

combination of a thin perforated screen attached to a relatively thick structural support. Figure 5 shows top and bottom photographs of a commercially available perforated tile with nominal 25% open area, constructed of a combination of a 0.00635 m (1/4 in) thick perforated screen with measured 21.26% open area, attached to a 0.028575 m (1−1/8 in) thick support with a measured open area of 30.68%. The inhomogeneous construction of the perforated tiles suggests the pressure drop characteristics may be directionally dependent, resulting in a significantly different loss coefficient in reversed flow situations. The resistance coefficient is not particularly significant in the limiting case of all perforated tiles having the same % open area, but the coefficient magnitude determines the relative flow weighting when various % open area tiles, such as to achieve a uniform PTFRD [21, 24, 25].

Table 2 summarizes the details of previous numerical RFP investigations in terms of facility geometry, number and % open area of the perforated tiles, the number and individual flow rate of each CRAC unit and some information regarding the mesh size and commercial CFD code used to model the flow. (Note an entry of '???' indicates no data specific data was reported) All of the RFP models presented in Table 2 assume a 'vacant' condition for the plenum or no obstructions caused by chilled water piping, power and data cables, with the exception of [14] which included the effects of a chilled water pipe. A current outstanding issue is the modeling of obstructions in RFP. The development of a generalized model is difficult because each data center has a unique set of blockages. Correlations for the pressure drop across single and multiple tube banks typically assume the flow is normal to the tube axis, but RFP flows are complex, swirling type flows where the flow direction is not easily determined without detailed CFD results. Thus, some numerical modeling will be required to predict the airflow patterns in the RFP and the challenge is shifted to determining the appropriate level of modeling for obstructions. A main parameter that needs to be determined is the size below which obstructions have no significant effect on the PTFRD, which for a tubular obstruction could possibly be quantified by a blockage area ratio of tube diameter to plenum depth.

Table 2 Raised Floor Plenum (RFP) numerical modeling details summary

	Author, Year Reference	Radmehr et al., 2005 [14]	VanGilder and Schmidt, 2005 [15]	Rambo and Joshi, 2004 [13]	Schmidt et al., 2004 [12]	Karki et al., 2003 [16]	Schmidt et al., 2001 [11]	Kang et al., 2000 [7]
Data center layout [m]	Length	13.41	10 different layouts	6.40	20.00	39.60	20.00	7.27
	Width	4.88		4.27	6.06	22.30	6.06	7.87
	RFP depth	0.419	0.30–0.91	0.61	0.292	0.76	0.284	0.457
Perforated tiles	# Perforated tiles	2–20	28–187	5	30–76	352	30–76	14
	% Open area	16%	25%, 56%	25%	19%	25%	19%	25%, 60%
CRAC units	# CRAC units	1	2–18	1	2	11	2	2
	Ind. flow rate [m^3/s]	???	???	5.71	5.78	5.85	???	2.48
Modeling	# Grid cells	???	???	66,490	???	136,770	???	48,000
	# Grid cells / m^3	???	297–64,000	15,954	???	815	???	7,343
	Turbulence model	k-ε	k-ε	k-ε, k-ω	???	k-ε	???	???
	Commerical code	TileFlow	Flovent	Fluent	???	TileFlow	Compact	Compact

4.2 Above floor facility models

Regardless of whether the cold supply air is delivered through a raised floor plenum or overhead ductwork, the space that houses the actual racks is where the primary thermal management issues arise. Schmidt et al. [21, 24, 25] has modeled the RFP flow separately and input the perforated tile flow rate predictions as boundary conditions to an above floor model. Validation in data center models above the RFP has usually come in the form of point-wise temperature measurements at the inlet to individual racks at various vertical locations. These measurements have also shown significant errors in the numerical temperature predictions. The temperature field is closely coupled to the velocity field due to the nature of most turbulence models that are applicable to data center CFD/HT, making accurate velocity field predictions essential to accurate temperature predictions. The only work that reported velocity measurements is that of Patel et al. [1], which found relative prediction error of greater than 10% for all point-wise velocity measurements taken and showed 14% relative error in point-wise temperature measurements, (*cf.* Section 3.3).

Buoyancy effects have been included in some CFD/HT computations, although scaling arguments have shown that they are negligible in the rack exhaust [32] and experiments have shown linear changes in point-wise temperatures measurements for varying CRAC exhaust temperature [50], further confirming these estimates. Increasing rack flow rates reduces the effects of buoyancy relative to the flow inertia; however, servers are being designed to allow larger temperature rises across the rack, increasing the effects of buoyancy. Buoyancy effects may become significant far from the rack exhaust and the spatially vary effects of buoyancy needs to be validated because including these effects may significantly increase the computational time of CFD/HT simulation depending on the structure, the coupling of the momentum and energy equations and treatment of the turbulence model.

A general issue that arises in the characterization of functional data centers is determining the actual power dissipation from the DPE. Thermal analyses of electronic equipment usually use the maximum power dissipation from each component to determine the heat load, although in practice all the DPE in a data center may dissipate the maximum amount of power simultaneously less than 5% of the time. The discrepancy between the actual power dissipation and theoretical maximum power supported by the device has caused chip manufactures to produce 'thermal design power' guidelines that are often significantly less than the power dissipation obtained assuming all electric power is converted directly into thermal energy. Variations in DPE workload could cause a facility comprised of identically configured racks to exhibit widely different power dissipation levels, both between racks at any given instant in time as well as overall facility power dissipation over time.

Table 3 summarizes the modeling details for those investigations primarily concerned with the data center facility and not the RFP. The table is organized into facility geometry, rack and CRAC details (note the number in parenthesis indicates supplemental cooling flow rate above the CRAC unit) the cooling scheme and some details regarding the mesh size and CFD model construction. The abbreviation STD, RFP and OH refer to the standard CRAC return, raised floor plenum supply and overhead supply and returns, respectively. The em dash (—) indicates a range of parameters, while the plus sign (+) indicates a combination of different parameters. The cooling

Table 3 Data center facility modeling details summary

	Author, Year Reference	Rambo and Joshi, 2006 [32]	Rambo and Joshi, 2006 [31]	Bhopte et al., 2005 [26]	Iyengar et al., 2005 [29]	Sharma et al., 2005 [45]	Shrivastava et al., 2005 [24]	Heydari and Sabounchi 2004 [40]	Schmidt and Cruz, 2004 [24]	Rambo and Joshi, 2003 [22]	Rambo and Joshi, 2003 [23]	Patel et al., 2002 [19]	Sharma et al., 2002 [18]	Schmidt and Cruz, 2002 [19]	Patel et al., 2001 [1]
Facility geometry [m]	Length	7.62–10.77	7.62–10.77	11.58	13.40	11.70	13.42	17.07	13.40	11.68	10.76	9.14	11.67	13.40	7.62
	Width	7.62–11.68	7.62–11.68	6.10	6.05	8.50	6.05	9.75	6.05	9.43	5.40	6.09	8.53	6.05	5.49
	RFP Dipth	0.61	0.61	0.61	0.00	0.60	–	0.61	0.15–0.60	0.60	0.60	0.60	0.60	0.00	0.00
	Celling height	3.00	3.00	3.05	3.66–6.71	3.10	3.66	3.66	2.74	3.00	3.00	3.05	3.05	2.44–3.05	3.96
Rack	# Racks	28	28	24	20	28	20	???	20	28	20–52	28	28	20	18
	Power [kW]	3.2	3.2	4.5	4–36	15.75	12	???	4–12	4.23	0	12	12	4–2	14.4
	Flow rate [m^3/s]	0.585	0.585	0.327	0.325–1.464	???	0.685	???	0.325–0.975	1.260	0.146–1.162	0.850	0.680	0.325–0.975	1.133
CRAC	# CRACs	4	4	2	5	4	–	2(5)	2	4	2	4	4	2	18
	Ind. flow rate [m^3/s]	5.66	5.66	3.96	0.78–5.86	???	–	5.66(2.12)	???	5.70	5.70	5.70	4.72	3.25–9.75	2.12
Perforated tiles	# Perf tiles	28	28	24	0	???	20	???	20	5	20–52	20	28	20	0
	% Open area	25%	25%	30%	0%	???	–	32%	25%, 60%	25%	25%	40%	???	%	0%
Cooling scheme	Supply	RFP	RFP, OH	RFP, OH	OH	RFP	RFP, OH	RFP+OH	RFP	RFP	RFP	RFP	RFP	RFP	OH
	Return	STD. RR	STD, OH	STD, OH	RR	STD	RR. OH	STD+OH	STD	STD	STD	STD	STD, OH	STD	OH
	Orientation	N-A	N-A	N	N-A	N+A	N	N	A	N	N	N	N+A	A	–
Global parameters	Power density [W/ft^2]	88–109	88–109	142	92–825	412	275	???	92–275	100	0	561	314	92–275	576
	Mrack / Mcrac	0.72	0.72	0.99	0.6–1.0	???	1.00	???	0.8–1.2	1.55	0.26–5.29	1.04	1.01	0.8-.2	0.53
	Qrack / Qcrac	0.24	0.24	0.57	0.16–1.44	1.16	–	???	–	0.31	0.00	0.88	0.88	–	0.15
Modeling	# Grid cells × 1000	1.500	1500	28	150	434	150	???	135	573	110	71	450	135	850
	# Grid cells / m^3	27656	33447	108	276–506	1179	505	???	608	1445	526	349	1239	546–682	5131
	Turbulence model	Various	k-ε	k-ε	k-ε	???	k-ε	k-ε	k-ε	k-ε	Various	???	???	k-ε	k-ε
	Commerical code	Fluent	Fluent	Flotherm	Flotherm	???	Flotherm	Flovent	Flovent	Fluent	Fluent	Flovent	Flovent	Flotherm	Flovent

scheme orientation 'N' indicates the CRAC unit exhaust direction is normal to the cold aisle while 'A' indicates the CRAC unit exhaust direction is aligned down the cold aisle. The global metric m_{rack}/m_{CRAC} indicates the ratio between the net rack flow rate and net CRAC flow and serves as a measure of recirculation while the Q_{rack}/Q_{CRAC} metric indicates the ratio of the new power dissipated by all the racks in the facility to the net cooling capacity of all the CRAC units, which is used to indicate the level of cooling provisioning.

4.3 Turbulence modeling

Data center airflows are turbulent throughout the facility based on the velocity and lengths scales, although the Reynolds number can vary drastically and the flow patterns exhibit a wide range of characteristics. In the absence of detailed experimental measurements, most CFD studies have employed the standard k-ε turbulence model, *cf.* Table 3. This model is the most basic of the differential closure models for Reynolds-averaged Navier-Stokes (RANS) equations. It is most broadly applicable because it involves the least number of assumptions. However, applicability does not imply accuracy and the standard k-ε model has been demonstrated to produce extremely large errors in a number of simple flows [51].

Specifying the boundary conditions and turbulence model validation are the primary issues in the simulation of complex industrial flows. Wall functions are used to compute the boundary conditions for the k and ε equations at solid surfaces, but such boundary conditions need to be prescribed for various inlet and outlet situations. Without detailed measurements, the profiles or even area-weighted average values of k and ε are not known and need to be manufactured. Correlations exist for k and ε given the turbulent intensity and hydraulic diameter, assuming fully developed conduit flow upstream of the imposed boundary condition.

4.4 Grid resolution and numerical accuracy

Some of the numerical models in Table 3 employ surprisingly coarse grid sizes, measured by the average grid cell density computed as the ratio of total number of grid cells to RFP model volume. The grid cell density used in [52] has been estimated based on minimum and maximum allowable grid sizes quoted in the paper. For perspective on grid cell density, a mesh with 1000 cells per m^3 has a characteristic grid cell of 0.1 m (4 in) per side. The smallest features that need to be resolved control the mesh size in CFD/HT, and local mesh refinement techniques exist to resolve sharp gradients without using uniformly fine grids. In the code used in [53], the specialized CFD solver for RFP modeling, uses a default grid spacing of 0.305 m (12 in) in the horizontal plane and 0.152 m (6 in) in the vertical direction (across plenum depth). Obstructions in the RFP below the grid cell size can be modeled by specifying a reduction in flow area through the appropriate grid cells, similar to the modeling of fluid flow through porous media.

The effect of each modeling simplification on perforated tile flow rate prediction accuracy needs to be validated in order to determine the detail at which the RFP geometry needs to be resolved. Only after these blockage effects have been quantified and reliable models have been developed can CFD models be used to simulate the airflow in real-world data center RFPs. Dependable and robust modeling techniques can greatly aid in the development of RFP design guidelines, because actual RFP facilities that can be used for design studies are rare and detailed measurements are highly time-consuming.

Grid sizes employed in many facility level simulations are seen in Table 3. The investigations [1, 22, 31, 32] model details at the rack level, which greatly increases the number of grid cells. Bhopte et al. [26] also modeled some details within the rack and used a far coarser mesh. While most CFD/HT investigations report some sort of grid convergence study, the criteria are often not specified and using such metrics as maximum temperature and velocities may not be sufficient to guarantee that reasonable estimates of all the velocity and temperature profiles have been computed.

Data center designers and managers want to quickly assess the thermal performance by using simple numerical models with short solution times. This goes against the CFD/HT fine grid requirements for modeling the complex airflows and heat transfer processes in data centers. Excessively coarse grids may result in serious errors, so there needs to be an acceptable trade-off between the levels of detail modeled and model size for quick analysis times.

4.5 Prototypical flow comparison

A generic data center model can be decomposed in a series of prototypical flows [23]. Any turbulence model employed for data center simulation must produce accurate predictions in these regions, while still being robust enough to produce reasonably accurate predictions in the remaining portions of the model. The CRAC exhaust region in a raised floor plenum is akin to a confined impinging jet, the cold aisle is similar to flow in a channel with distributed suction, and the hot aisle is essentially a bank of jets directed at each other.

The length scales of RFP flows are similar to those of other indoor airflows, especially the modeling of large office spaces or auditoria. Indoor airflows typically fall into the regime of turbulent mixed convection. However, the flow in data center RFPs is at a Reynolds number (*Re*) significantly greater than indoor airflows. The literature for high *Re* cavity flows typically considers compressible and supersonic flows for acoustic modeling, making existing CFD modeling studies of incompressible flows at high *Re* in cavity-type geometries scarce.

A CRAC flow rate of $5.75\,\mathrm{m}^3/\mathrm{s}$ is representative of most RFP investigations, as shown in Table 3. Note all flow rates are referenced to the density of air taken at 15 °C, the common outlet temperature of CRAC units. Commercially available CRAC units typically use 2 blowers without exhaust areas of approximately 0.117 m^2 (1.25 ft^2) for a hydraulic diameter based Reynolds number ($Re = uD_h/\nu$) of ∼530,000. Recent impinging jet investigations [53–57] have considered target plate distances (H) between 2 to 50 jet nozzle hydraulic diameters (D_h), $2 \leq H/D_h \leq 50$. Based on CRAC exhaust of a single blower, the aspect ratio for RFP data centers between 0.305 m (12 in) to 0.914 m (36 in) is $0.941 \leq H/D_h \leq 2.82$, putting the geometry in the lower range of prototypical studies. The standard k-ε model has been shown to fail for predicting impinging jet velocity profiles, wall shear stress and heat transfer rates, and the previous studies [53–57] use the k-ω model [51] to predict the flow field for $Re \leq 110,000$. The CRAC exhaust *Re* is well beyond the range of other numerical investigations, requiring the validation of the standard k-ε model for RFP predictions. Although very detailed velocity profiles in the CRAC exhaust region may not appear to be of primary concern to data center designers, a number of recent studies [13, 58] have focused on the behavior of perforated tiles in close proximity to the CRAC units, and especially the prediction of reversed flow.

The hot-aisle cold-aisle approach has gained wide acceptability in data center layouts, making this rack configuration possibly the only common feature between data centers. This makes it a good candidate for validation studies, since the hot- and cold-aisle flow patterns strongly influence the overall thermal performance of the data center. The flow in a cold aisle is similar to flow in a channel with wall suction, or fluid pulled out of the bulk flow through perforations in the wall. The flow is independent of the channel orientation, either vertical or horizontal, because both buoyancy effects and the weight of the fluid are negligible. Limiting cases of 'low' and 'high' suction rates need to be considered. The low suction rate would be applicable to low flow rate racks that do not draw a large portion of the bulk and corresponds to a more prototypical flow where literature may be available. The high suction rate limit is more applicable to data centers where all of the bulk flow is entrained through the sidewalls. The hot aisle is also similar to a channel flow, but with injection or fluid being added to the bulk flow through the walls. While there are some differences between the flow in data center cold- and hot-aisles and channel flows with suction and injection, an experimentally validated model of such flows would provide guidance in selecting appropriate turbulence models for data center airflow simulations.

5 Rack-level compact modeling

Modeling approaches for racks in data center CFD/HT simulations have varied significantly depending on the purpose of the investigation. The level at which data center managers are ultimately responsible for the DPE is currently a source of debate. Some argue that the DPE manufacturer has designed an adequate thermal management system for the components contained within the racks, and the data center designer is only responsible for meeting the airflow rate and inlet temperature requirements. Considering the high reliability expected from data processing environments and noting that an equipment failure is a failure regardless of the root cause, others argue that some level of modeling of the individual servers is required to fully understand data center thermal management.

In the recent analysis of data center airflows, researchers have either modeled the rack in a black-box fashion with prescribed flow rate and temperature rise, or as a box with fixed flow rate and uniform heat generation [1, 6]. Subdividing the rack into a series of server models was introduced in [22] where each server contained an induced-draft fan model at the exhaust, a lumped flow resistance at the inlet to model the pressure drop, and a uniform heat flux over the bottom surface to model the power dissipation from a number of components. The reasoning behind subdividing the rack into multiple server models is that it is the heat generated at the chip level inside the servers that causes the need for a data center cooling scheme and the convective processes at the server level should be captured in some manner.

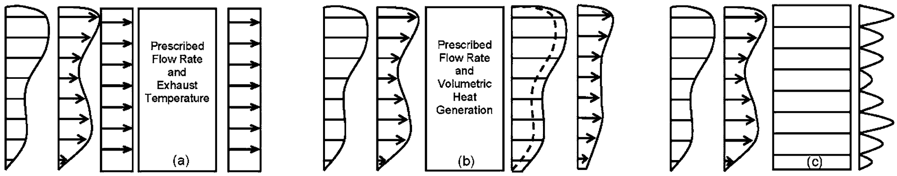

Fig. 6 Rack-level modeling approach comparison for a given inlet velocity and temperature profile (flow is left to right) showing (a) uniform velocity and temperature exhaust profiles for a prescribed rack model, (b) slight velocity change and constant (T(x) + ΔT) change in temperature profile for uniform volumetric heat generation models and (c) localized velocity and temperature profiles for server-level modeling

The rack-level model employed depends upon the objective of the simulation. Figure 6 illustrates the differences between the different rack modeling approaches. Black-box type models are essentially boundary conditions that prescribe outlet conditions on the velocity and temperature field at the rack exhaust, given a set of conditions over the rack inlet. The most common approach is based on an overall energy balance ($Q = \dot{m}c_p\Delta T$, where Q is the rack power dissipation rate, $\dot{m}$ is mass flow rate, c_p is the specific heat of air and ΔT is the fluid bulk temperature rise between the rack inlet and outlet). Here the rack flow rate $\dot{m}$ is computed based on the rack power dissipation rate (Q), and the specified fluid bulk temperature rise ΔT. The velocity and temperature profiles are then specified as uniform values over the rack exhaust, because no specific information regarding the flow details inside the rack have been computed. This approach requires the least computational effort, but nearly all the details of how the rack interacts with the data center facility are lost. Another black-box type approach coarsely computes the flow through the rack and models the power dissipation rate as a uniform volumetric energy source, $q''' Q/V_{\text{rack}}$. This procedure will produce a profile at the rack exhaust, which is entirely dependent on the rack inlet profiles because the rack model does not significantly alter the velocity field.

The rack flow rate is fixed for both fixed temperature rise and uniform power dissipation black-box rack models. The rack sub-modeling of [22] used pressure-velocity relationships to model server-level fans and up to 10% variations in net rack flow rate were observed in a CFD/HT study of 28 racks using the same flow rate model designed for a high rack flow rate of 1.260 m^3/s (∼2400 CFM). Specifying the rack flow rate also smears out the effect of inhomogeneous rack populations since racks in operational data centers often contain a variety of DPE, each with their own flow rate characteristics. Fixing the net rack flow rate cannot determine if higher flow rate DPE is consuming most of the cooling supply and starving the lower flow rate DPE of adequate airflow. Another emerging issue is that DPE manufacturers are beginning to use variable speed fans in order to maintain constant component temperatures inside the servers. The rationale behind this considers DPE workload variations and aims to not waste cooling supply air for a nominally high power dissipation server that is not being used and therefore not outputting any heat.

Modeling individual servers can describe power dissipation and flow rate variations at the rack level, but such an approach significantly increases the CFD/HT model size. Limits on available computing have severely restricted the server-level details in the past, reducing server-level models to boxes with constant heat fluxes on the bottom surface and a fan model specified over the entire exhaust region. This is still


(a) (b)

Fig. 7 Photographs of actual compute racks, showing (a) individual server fan and (b) rack cabling

a significant departure from reality because fans and vents only occupy a portion of the rear of the server, with the rest dedicated to cable connections, see Fig. 7a. The cabling in the rear of rack may also present a significant blockage and reduce the DPE flow rate, see Fig. 7b.

The level of detailed modeled at the rack depends on the purpose of the investigation and available computing resources. Black-box type models may be satisfactory if the purpose of the data center model is to investigate airflow and heat transfer characteristics far from the rack, such as a CRAC unit efficiency investigation, where the racks only act as the thermal load. More details at the rack level are required if the scope of the investigation concerns DPE thermal management because the interface between the rack and facility is what drives the need for data center cooling schemes. (If the racks did not exhaust to a common data center environment, then there would be no need for CRAC units.)

5.1 Compact models

Compact models offer an acceptable tradeoff between modeling details and computational expense, although the term 'compact model' has received wide use without a consensus on the definition. Here, 'compact model' will be defined as a model that uses a number of internal states to produce pre-defined output data, given a prescribed set of inputs and parameters. The use of internal states is the key differentiating factor between compact and black-box (or commonly 'lumped') models. Lumped models are input—output functions obtained from physical conservation laws, or curve fits to experimental data. For example, fan and resistance models that specify a pressure-velocity relationship are common lumped models that compute the velocity through a surface as a function of the local upstream and downstream pressure. Figure 8a illustrates a lumped model for the temperature difference between the chip and ambient of an electronic package.

Compact models in the form of thermal resistance networks have been widely used in electronics cooling [59, 60] although the model accuracy is strongly subject to how the resistance network is constructed and how multidimensional effects are treated with equivalent 1-dimensional resistances. Figure 8b shows a resistance network type

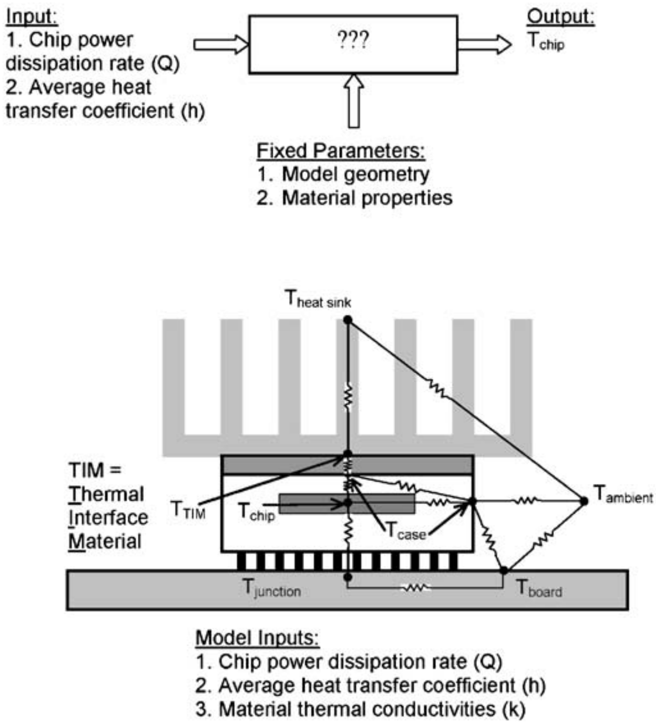

Fig. 8 Two models of an electronic package, (a) lumped model and (b) resistance network compact model

compact model of an electronic package, which computes the temperature at a number of discrete locations (i.e. internal states), such as junction and case temperatures.

Compact models use slightly more degrees of freedom (DOF) than lumped models, but provide internal states based on additional details that allow for further examination of the model. An important process in modeling is system identification, or the rigorous experimentation with unknown (black-box) or partially characterized (gray-box) systems to efficiently and accurately determine the underlying behavior. At the other end of the DOF spectrum are detailed CFD/HT models and finite element analyses, which are subsets of distributed parameter modeling which aims to approximate the full-field system behavior. The process of taking a model from a large number of DOF, either from detailed numerical simulations or full-field experimental measurements, to a model involving significantly fewer DOF is termed *model reduction*. A number of tools exist for reducing the number of internal states of large linear systems, such as those resulting from discretizing differential equations [61]. As an example, the proper orthogonal decomposition has been used to accurately reduce detailed turbulent CFD/HT models by a factor of 10^4 and 10^5 [41, 42, 44, 62, 63]. Figure 9 illustrates this taxonomy of efficient modeling procedures by comparing the level of description and model size, measured in DOF.

The rack level is the appropriate scale for compact modeling in data centers because interface between the racks and the data center facility is the key focus for data center thermal management. Also, the uniqueness in layout and cooling scheme at the facility

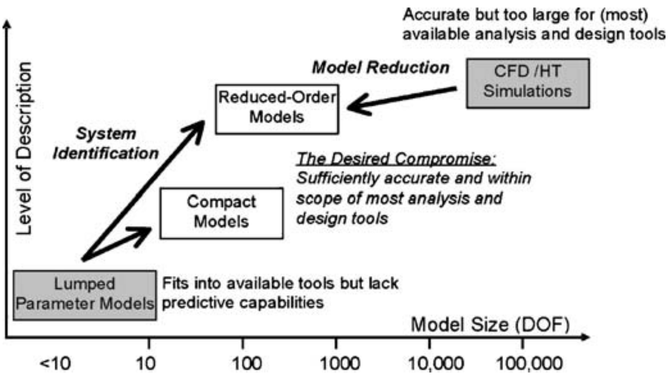

Fig. 9 Modeling description level and DOF taxonomy

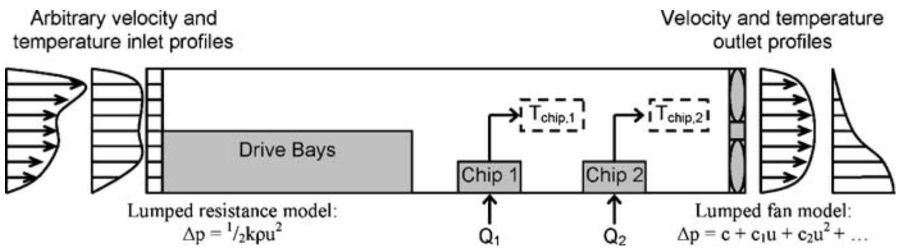

Fig. 10 Server thermal model sketch for compact model development

level makes developing generalized models above the rack level very difficult. To provide the broadest description as possible, rack level compact models will actually be an agglomeration of server-level compact models to account for heterogeneous rack populations. The end goal of the compact model is for it to be integrated with full scale CFD/HT models of data center facility to efficiently produce detailed descriptions at the rack level without significantly increasing the computational cost over black-box type rack models.

Figure 10 shows a server model used to describe the compact model development procedure. The main features of the model include some flow obstructions created by drive bays and power supplies that provide a representative system pressure drop. A number of power dissipating components are also modeled with arbitrarily complex details including multiple materials and chip-level thermal management devices such as heat sinks. The model includes an induced-draft fan to pull the air through the server box and may also include a lumped resistance to simulate the pressure drop through screens at the front of the server. Modeling the server inlet and outlet as discrete vents and fans is a significant improvement over the previous attempts to subdivide the rack into a series of server models because those models required the inlet and exhaust to be smeared over the entire face of the server inlet and outlet due to computation limits.

The model thermal inputs are inlet temperature and component-wise power dissipation rates. The flow inputs may vary depending on the scope of the compact model, with the simplest being a fixed velocity that also fixes the outlet velocity by continuity. More advanced strategies would use a fan model to drive the flow against the server system resistance accounting for the inlet and outlet static pressures. Since the ultimate goal is for the compact model to be integrated into the full CFD/HT computation of data center facility, detailed velocity, pressure and temperature profiles are available from the neighboring grid cells in the computational mesh. This means that the compact models can be formulated to accept and output either profiles or area-averaged quantities. The model parameters are the geometry and material properties, which may range over specified values, if the model is constructed to accommodate such variability.

The specific details of compact model construction are beyond the scope of this review, but suffice it to say a wide range of tools is available. A straightforward approach would be to construct a physical or numerical model of the server and correlate the maximum component temperature as a function the server flow rate, component power dissipation and geometrical properties. Alternatively, reduced-order modeling techniques could be used to quickly approximate the velocity and temperature fields and a post-processing layer can be added to extract maximum chip temperatures [32, 41, 42]. More detailed analytic approaches would involve control volume methods coupled with heat transfer correlations for flow over multiple heated blocks, which is the prototypical geometry for electronics thermal management and has been studied extensively in the literature.

5.2 Model integration

Rack-level compact models interact with the CFD/HT simulation as input—output functions, but unlike black-box models, they output their internal states such as a list of component temperatures and energy usage data for thermal performance monitoring. If the server compact model has a fixed flow rate, the CFD/HT simulation only interacts with the compact model through prescribed boundary conditions, but the model still resides in the category of compact models if it outputs internal states. A compact model that determines the server flow rate from a fan model coupled to the CFD/HT mesh needs to be solved iteratively with the velocity and pressure fields, because the server flow rate is a function of the inlet and outlet pressure fields neighboring the compact model.

Server-level compact models can be stacked up to form a rack-level compact model, with 4 such models used here for illustrative purposes. The compact models are vertically located on top of each other do not directly interact by exchanging mass or momentum, i.e. the top and bottom of the servers are treated as adiabatic solid walls. This is typical of most air-cooled racks that draw air in from the front and exhaust it through the rear. The server-level models indirectly interact as the output of server or rack may eventually lead to the input of another model through the computational mesh, resolving recirculation effects. Direct compact model interaction can be accounted for by specifying additional heat fluxes on the top or bottom surface of the server-level models. A major issue is coupling the discrete exhaust velocity and temperature profiles from each compact model because the exhaust of typical rack-mounted DPE does align flush with the back of the rack to allow for power and data cabling. The mixing of the individual exhaust jets can be accounted for with smoothing functions that conserve mass, momentum and energy across the shear layers where the individual jets merge. Cables blockages (*cf.* Fig. 7b) can be included through lumped pressure drop models.

The ability to couple compact models at the rack level extends far beyond corrections for shear layers and cable blockages and allows users to model rack-level cooling devices such as integrated chilled water loops [64, 65] and passive airflow management schemes [66, 67]. Heat transfer effects from chilled water flow in pipes can be characterized analytically and do not need to be solved for. A chilled water heat exchanger compact model would feature the pressure drop and heat transfer rate through the device. Additional internal states could be used to monitor local humidity levels to guard against condensation. A review of the current state of the art of liquid cooling for data centers has been presented by Beaty [68].

Coupling compact models requires the suitable passing of information between models. The communication protocols between the models are the physically conserved quantities of mass, momentum and energy. Mass and energy conservation are achieved through matching the appropriate mass flow and heat transfer rates, respectively. Conservation of momentum can be expressed as a pressure-velocity relationship, $\Delta p = f(u)$, in a 1-dimensional sense for scalar compact model outputs. Additional methods are available to account for boundary profiles [62]. These conserved quantities form the basis for coupling facility wide models, such as racks directly ducted to the RFP [69], and incorporating more advanced models for data center air handling devices [70].

Compact modeling at the rack level can provide multiple levels of description for data center airflow and heat transfer characteristics depending on the details included in the compact model. The rack model presented in Fig. 11 uses 4 server-level sub-models, which may be adequate to investigate alternative cooling schemes. It also fits into other data center thermal profiling efforts [30, 45] which monitor the temperature at the rack inlet at 4 different vertical locations. Maintaining libraries of server-level compact models from various manufacturers would not require developing individual models from scratch because most manufacturers have performed thermal analyses of their equipment to ensure its reliable operation and determine the require inlet temperature and flow rate requirements. Many such models are CFD/HT-based due to the widespread use of commercial codes, making detailed flow and heat transfer information available for compact model development.

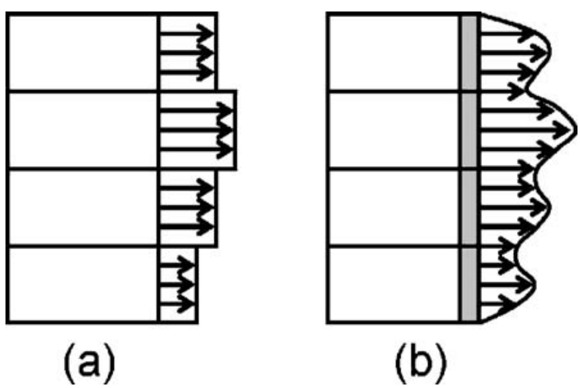

Fig. 11 Effect of rear rack cabling on compact model output

Again, the language of physically conserved quantities can form the basis to develop a common reporting system for different manufacturers to share information without having to disclose proprietary information.

6 Data center dynamics

Most data center thermal analyses have focused on a single operating state of the facility, or considered the data center over its entire lifecycle, but the overall thermal efficiency can be improved by considering the intermediate timescale of hour-by-hour and day-to-day operation. Some racks and possibly entire data center facilities experience a drastic change in computing load over the course of several hours and days depending on the nature of computing being performed. Recent studies have suggested dynamically partitioning the computing load to maximize cooling system efficiency [45, 50], although this may require radical changes in current computing infrastructures. However, dynamically adjusting the cooling capacity to meeting continually changing power dissipations at the rack level can lead to large savings for the data center end-user.

There is a sizable increase in computational expense associated with simulating dynamical phenomena with CFD/HT because of large model sizes needed to accurately resolve spatial gradients and the large number of time steps needed to accurately resolve temporal gradients. The use of compact models is deeply rooted in dynamical investigations, especially for control to efficiently predict internal states of complex systems based on measured quantities at specified locations. A simple data center control scheme could sense rack inlet temperatures and adjust CRAC flow rates and output temperature to eliminate local hotspots. More efficient control schemes would be at the rack level with variable chilled water flow rates to rack-level heat exchangers and a relatively few number of CRAC units to provide supplemental cooling to the facility for human comfort. Current research issues concerning the type and placement of sensors need to be addressed.

With chip temperatures being the ultimate concern in data centers, compact models can play an important role in developing efficient dynamic multi-scale simulation tools, with resolution down to components within individual servers. A large portion of the existing literature in compact modeling is aimed at dynamic modeling and implementing such models in control schemes. Dynamic compact models at the rack level can be integrated with full-scale data center CFD/HT simulations to develop control scheme that resourcefully use variable speed fans at the server and rack level in order to efficiently accommodate power dissipation variations and avoid drastic over-provisioning of cooling resources. Since dynamic compact models can be integrated out in time given a set of initial conditions, it may be possible to implement compact models and the full scale CFD/HT at different time scales. Dynamic compact model development can extend beyond individual servers and include rack-level liquid cooling schemes and supplementary cooling units at the rack level.

7 Summary

To review the existing literature, Table 4 summarizes all the literature presented in Section 3–5 directly related to data center thermal management, organized in reverse

Table 4 Data center thermal management literature summary

Author	Year	Ref	Alternative cooling	Control and lifecycle	Efficiency	Layout	Liquid cooling	Metrics	Modeling	Plenum	Rack	Measurements	Computational
Karki and Patankar	2006	[18]							×	×			
Rambo and Joshi	2006	[32]				×		×	×				×
Rambo and Joshi	2006	[31]	×						×				×
Rolander et al.	2006	[42]		×	×				×		×		
Bhopte et al.	2005	[26]				×							×
Herrlin	2005	[39]						×			×		×
lyengar et al.	2005	[29]	×			×							×
Karlsson and Moshfegh	2005	[3]			×							×	
Patel et al.	2005	[46]		×	×								
Radmehr et al.	2005	[14]								×			×
Rambo and Joshi	2005	[38]						×			×		×
Rolander et al.	2005	[41]		×	×				×		×		×
Schmidt et al. ***	2005	[4]						×					
Schmidt and Ivenqar	2005	[30]				×						×	
Shah et al.	2005	[37]			×			×				×	
Shah et al.	2005	[36]			×			×	×				
Shah and Krishnan	2005	[48]		×									
Sharma et al.	2005	[45]		×				×					×
Shrivastava et al.	2005	[28]	×										×
VanGilder and Schmidt	2005	[15]								×			×
Webb and Nasir	2005	[64]					×					×	
Beaty	2004	[68]					×						
Boucher et al.	2004	[50]		×								×	
Furihata et al.	2004	[66]			×							×	
Heydari and Sabounchi	2004	[40]									×		×
Hwang et al.	2004	[69]									×	×	

(*Continued on next page*)

Table 4 (Continued)

Author	Year	Ref	Alternative cooling	Control and lifecycle	Efficiency	Layout	Liquid cooling	Metrics	Modeling	Plenum	Rack	Measurements	Computational
Rambo and Joshi	2004	[13]							×	×			×
Schmidt and Cruz	2004	[24, 25]				×							×
Schmidt et al.	2004	[12]								×		×	×
Sharma et al.	2004	[33]						×				×	
Wang	2004	[67]									×	×	
Wilson et al.	2004	[65]					×				×		
Karki et al.	2003	[16]							×	×		×	×
Koplin	2003	[9]								×		×	
Norota et al.	2003	[34]			×			×				×	
Rambo and Joshi	2003	[22]				×			×				×
Rambo and Joshi	2003	[23]				×			×				×
Shah et al.	2003	[35]			×			×	×				
Anton et al.	2002	[70]							×				
Herold and Radermacher	2002	[47]		×									
Patel et al.	2002	[19]			×	×							×
Schmidt and Cruz	2002	[21]				×				×			×
Schmidt and Shaukatullah***	2002	[8]											
Sharma et al.	2002	[20]				×		×					×
Sullivan	2002	[5]				×							
Baer	2001	[6]					×						
Patel et al.	2001	[1]	×									×	×
Schmidt et al.	2001	[11]								×		×	×
Stahl and Belady	2001	[27]	×									×	
Kang et al.	2000	[7]							×	×			×

***Review paper with historical prespective

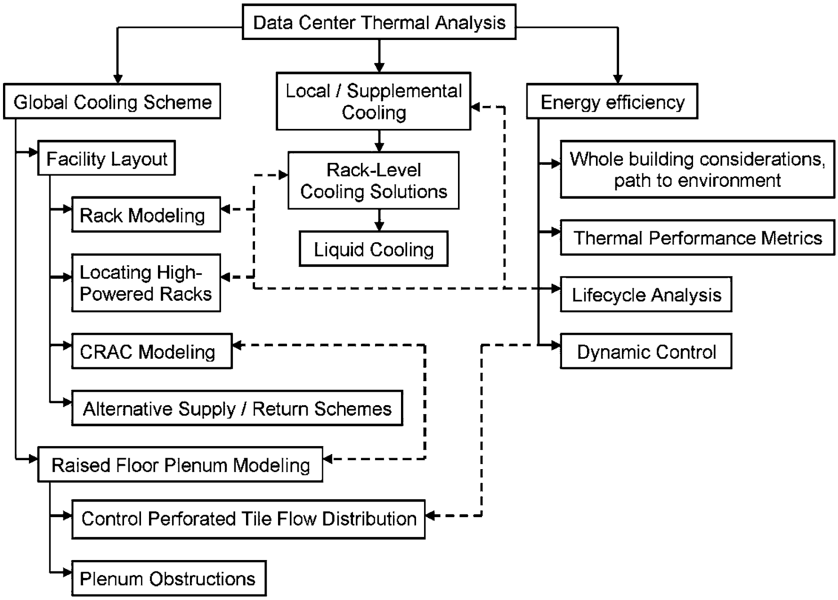

Fig. 12 Data center thermal management organizational chart

chronological order. Table 4 utilizes the subdivisions of Section 3 and extends the summary presented in Table 1 and includes the additional column of 'Computational' to highlight those investigations, which are predominantly computational in nature, and also includes 2 review papers.

In addition to a brief review of numerical modeling and possible validation studies for data center airflows and heat transfer, this paper presents a number of future directions to increase data center energy efficiency. Figure 12 presents a data center thermal analysis organization chart that establishes a hierarchy of research directions and the interconnections between different topics. The three main thrusts are evaluation of global cooling schemes, investigating the effects of local and supplemental cooling schemes, and improving energy efficiency at all levels of the data center. All topics in the organizational chart are directed at thermal optimization and thus no specific optimization block is included as it is the underlying theme of the entire hierarchy.

Acknowledgements The authors acknowledge the support of this work through the Consortium for Energy Efficient Thermal Management (CEETHERM), a joint academic and industry initiative focusing on current and next-generation thermal issues in systems level electronic architectures.

References

1. C.D. Patel, C.E. Bash, C. Belady, L. Stahl, and D. Sullivan, "Computational fluid dynamics modeling of high compute density data centers to assure system inlet air specifications," Presented at Proc. IPACK'01 – The Pacific Rim/ASME International Electronics Packaging Technical Conference and Exhibition Kauai, HI, 2001.

2. “Data Center Energy Benchmarking Case Study.” California: Lawrence Berkeley National Laboratory and Rumsey Engineers, Inc., 2003.
3. J.F. Karlsson and B. Moshfegh, “Investigation of indoor climate and power usage in a data center,” Energy and Buildings, vol. 37, pp. 1075–1083, 2005.
4. R.R. Schmidt, E.E. Cruz, and M. Iyengar, “Challenges of data center thermal management,” IBM Journal of Research and Development, vol. 49, pp. 709–723, 2005.
5. R.F. Sullivan, “Alternating cold and hot aisles provides more reliable cooling for server farms,” The Uptime Institute, 2002.
6. D.B. Baer, “Emerging cooling requirements and systems in telecommunications spaces,” Liebert Corporation 2001.
7. S. Kang, R.R. Schmidt, K.M. Kelkar, A. Radmehr, and S.V. Patankar, “A methodology for the design of perforated tiles in a raised floor data center using computational flow analysis,” presented at ITHERM 2000 – Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, Las Vegas, NV, 2000.
8. R.R. Schmidt and H. Shaukatullah, “Computer and telecommunications equipment room cooling: a review of the literature,” presented at ITHERM 2002 – Eight Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, San Diego, CA, 2002.
9. E.C. Koplin, “Data center cooling,” ASHRAE Journal, vol. 45, pp. 46–53, 2003.
10. C. Belady, K.M. Kelkar, and S.V. Patankar, “Improving productivity in electronic packaging with flow network modeling (FNM),” Electronics Cooling, vol. 5, pp. 36–40, 1995.
11. R.R. Schmidt, K.C. Karki, K.M. Kelkar, A. Radmehr, and S.V. Patankar, “Measurements and predictions of the flow distribution through perforated tiles in raised floor data centers,” Presented at IPACK’01 – The Pacific Rim / ASME International Electronics Packaging Technical Conference and Exhibition, Kauai, HI, 2001.
12. R.R. Schmidt, K.C. Karki, and S.V. Patankar, “Raised-floor data center: perforated tile flow rates for various tile layouts,” Presented at ITHERM 2004 – Ninth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, Las Vegas, NV, USA, 2004.
13. J. Rambo and Y. Joshi, “Supply air distribution from a single air handling unit in a raised floor plenum data center,” Presented at ISHMT/ASME’04 –Joint Indian Society of Heat and Mass Transfer – American Society of Mechanical Engineers Heat and Mass Transfer Conference, Kalpakkam, India, 2004.
14. A. Radmehr, R.R. Schmidt, K.C. Karki, and S.V. Patankar, “Distributed leakage flow in raised-floor data centers,” Presented at IPACK’05 – International Electronic Packaging Technical Conference and Exhibition, San Franciso, CA, 2005.
15. J.W. Van Gilder and R.R. Schmidt, “Airflow uniformity through perforated tiles in a rasied-floor data center,” Presented at IPACK’05 – International Electronic Packaging Technical Conference and Exhibition, San Francisco, CA, 2005.
16. K.C. Karki, A. Radmehr, and S.V. Patankar, “Use of computational fluid dynamics for calculating flow rates through perforated tiles in raised-floor data centers,” International Journal of Heating, Ventilation, Air-Conditioning and Refrigeration Research, vol. 9, pp. 153–166, 2003.
17. Innovative Research Inc., TileFlow v. 3.0: A Simulation Tool for Airflow and Temperature Distribution in Raised-Floor Data Centers. Plymouth, MN, 2006.
18. K.C. Karki and S.V. Patankar, “Airflow distribution through perforated tiles in raised-floor data centers,” Building and Environment, vol. 41, pp. 734–744, 2006.
19. C.D. Patel, R. Sharma, C.E. Bash, and A. Beitelmal, “Thermal considerations in cooling of large scale high compute density data centers,” Presented at ITHERM 2002 – Eight Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, San Diego, CA, 2002.
20. R.K. Sharma, C.E. Bash, and C.D. Patel, “Dimensionless parameters for the evaluation of thermal design and performance of large-scale data centers,” AIAA, vol. AIAA 2002–3091, 2002.
21. R.R. Schmidt and E. Cruz, “Raised floor computer data center: effect on rack inlet temperatures of chilled air exiting form both the hot and cold aisles,” Presented at ITHERM 2002 – Eight Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, San Diego, CA, 2002.
22. J. Rambo and Y. Joshi, “Multi-scale modeling of high power density data centers,” Presented at IPACK’03 – The Pacific Rim / ASME International Electronics Packaging Technical Conference and Exhibition, Kauai, HI, 2003.
23. J. Rambo and Y. Joshi, “Physical models in data center airflow simulations,” Presented at IMECE’03 – ASME International Mechanical Engineering Congress and R&D Exposition, Washington D.C., 2003.

24. R.R. Schmidt and E. Cruz, "Cluster of high powered racks within a raised floor computer data center: Effects of perforated tiles flow distribution on rack inlet air temperature," Presented at IMECE'03 – ASME International Mechanical Engineering Congress and R&D Exposition, Washington D.C, 2003.
25. R.R. Schmidt and E.E. Cruz, "Cluster of high-powered racks within a raised-floor computer data center: Effect of perforated tile flow distribution on rack inlet air temperatures," ASME Journal of Electronic Packaging, vol. 126, pp. 510–518, 2004.
26. S. Bhopte, D. Agonafer, R.R. Schmidt, and B. Sammakia, "Optimization of data center room layout to minimize rack inlet air temperature," Presented at IPACK'05 – International Electronic Packaging Technical Conference and Exhibition, San Francisco, CA, 2005.
27. L. Stahl and C. Belady, "Designing an alternative to conventional room cooling," Presented at INTELEC'01 – 24th Annual International Telecommunications Energy Conference, 2001.
28. S. Shrivastava, B. Sammakia, R.R. Schmidt, and M. Iyengar, "Comparative analysis of different data center airflow management configurations," Presented at IPACK'05 – International Electronic Packaging Technical Conference and Exhibition, San Francisco, CA, 2005.
29. M. Iyengar, R.R. Schmidt, A. Sharma, G. McVicker, S. Shrivastava, S. Sri-Jayantha, Y. Amemiya, H. Dang, T. Chainer, and B. Sammakia, "Thermal characterization of non-raised floor air cooled data centers using numerical modeling," Presented at IPACK'05 – International Electronic Packaging Technical Conference and Exhibition, San Francisco, CA, 2005.
30. R.R. Schmidt and M. Iyengar, "Effect of data center layout on rack inlet air temperatures," Presented at IPACK'05 – International Electronic Packaging Technical Conference and Exhibition, San Francisco, CA, 2005.
31. J. Rambo and Y. Joshi, "Convective transport processes in data centers," Numerical Heat Transfer A – Applications, accepted for publication 2006.
32. J. Rambo and Y. Joshi, "Thermal modeling of technology infrastructure facilities: A Case Study of Data Centers," in The Handbook of Numerical Heat Transfer: Vol II, W.J. Minkowycz, E.M. Sparrow, and J.Y. Murthy, Eds. New York: Taylor and Francis, 2006, pp. 821–849.
33. R. Sharma, C.E. Bash, C.D. Patel, and M. Beitelmal, "Experimental investigation of design and performance of data centers," Presented at ITHERM 2004 – Ninth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, Las Vegas, NV, 2004.
34. M. Norota, H. Hayama, M. Enai, T. Mori, and M. Kishita, "Research on efficiency of air conditioning system for data center," Presented at INTELEC'03 – 25th International Telecommunications Energy Conference, Yokohama, Japan, 2003.
35. A. Shah, V.P. Carey, C.E. Bash, and C.D. Patel, "Exergy analysis of data center thermal management systems," Presented at IMECE'03 – ASME International Mechanical Engineering Congress and R&D Exposition, Washington, D.C, 2003.
36. A.J. Shah, V.P. Carey, C.E. Bash, and C.D. Patel, "Exergy-based optimization strategies for multi-component data center thermal management: Part I: Analysis," presented at IPACK'05 – International Electronic Packaging Technical Conference and Exhibition, San Francisco, CA, 2005.
37. A.J. Shah, V.P. Carey, C.E. Bash, and C.D. Patel, "Exergy-based optimization strategies for multi-component data center thermal management: Part II: Application and validation," Presented at IPACK'05 – International Electronic Packaging Technical Conference and Exhibition, San Francisco, CA, 2005.
38. J. Rambo and Y. Joshi, "Thermal performance metrics for arranging forced air cooled servers in a data processing cabinet," ASME Journal of Electronic Packaging, vol. 127, pp. 452–459, 2005.
39. M.K. Herrlin, "Rack cooling effectiveness in data centers and telecom central offices: The rack cooling index (RCI)," ASHRAE Transactions, vol. 111, pp. 725–731, 2005.
40. A. Heydari and P. Sabounchi, "Refrigeration assisted spot cooling of a high heat desnity data center," Presented at ITHERM 2004 – Ninth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, Las Vegas, NV, 2004.
41. N. Rolander, J. Rambo, Y. Joshi, and F. Mistree, "Towards sustainable design of data centers: Addressing the lifecycle mismatch problem," Presented at IPACK'05 – International Electronic Packaging Technical Conference and Exhibition, San Franciso, CA, 2005.
42. N. Rolander, J. Rambo, Y. Joshi, F. Mistree, and J.K. Allen, "Robust design of turbulent convective systems using the proper orthogonal decomposition," ASME Journal of Mechanical Design: Special Issue Robust and Risk Based Design, accepted for publication 2006.
43. J. Rambo, "Reduced-order modeling of multiscale turbulent convection: Application to data center thermal management," Ph.D. Thesis, Department of Mechanical Engineering, Georgia Institute of Technology, 2006.

44. J. Rambo and Y. Joshi, “Reduced order modeling of steady turbulent flows using the POD,” Presented at ASME Summer Heat Transfer Conference, San Francisco, CA, 2005.
45. R.K. Sharma, C.E. Bash, C.D. Patel, R.J. Friedrich, and J.S. Chase, “Balance of power: Dynamic thermal management for internet data centers,” IEEE Internet Computing, vol. 9, pp. 42–49, 2005.
46. C.D. Patel, C.E. Bash, R. Sharma, A. Beitelmal, and C.G. Malone, “Smart chip, system and data center enabled by advanced flexible cooling resources,” Presented at Twenty First Annual IEEE Semiconductor Thermal Measurement and Management Symposium, 2005.
47. K.E. Herold and R. Radermacher, “Integrated power and cooling systems for data centers,” Presented at ITHERM 2002 – Eight Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, San Diego, CA, 2002.
48. A.J. Shah and N. Krishnan, “Life evaluation of combined heat and power alternatives in data centers,” Presented at IEEE Interantional Symposium on Electronics and the Environment 2005.
49. E. Fried and I.E. Idelchik, Flow Resistance: A Design Guide for Engineers. New York: Hemisphere Publishing, 1989.
50. T.D. Boucher, D.M. Auslander, C.E. Bash, C.C. Federspiel, and C.D. Patel, “Viability of dynamic cooling control in a data center environment,” Presented at ITHERM 2004 – Ninth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, Las Vegas, NV, 2004.
51. S.B. Pope, Turbulent Flows. New York: Cambridge University Press, 2000.
52. J.W. VanGilder and R.R. Schmidt, “Airflow uniformity through perforated tiles in a rasied-floor data center,” Presented at IPACK’05 – International Electronic Packaging Technical Conference and Exhibition, San Francisco, CA, 2005.
53. Q. Chen and V. Modi, “Mass transfer in turbulent impinging slot jets,” International Journal of Heat and Mass Transfer, vol. 42, pp. 873–887, 1999.
54. S. Maurel and C. Solliec, “A turbulent plane jet impingment nearby and far from a flat plate,” Experiments in Fluids, vol. 31, pp. 687–696, 2001.
55. B. Merci, J. Vierendeels, C. DeLanghe, and E. Dick, “Numerical simulation of heat transfer of turbulent impinging jets with two-equation turbulence models,” International Journal of Numerical Methods for Heat and Fluid Flow, vol. 13, pp. 110–132, 2002.
56. T.H. Park, H.G. Choi, J.Y. Yoo, and S.J. Kim, “Streamline upwind numerical simulation of two-dimensional confined impinging slot jets,” International Journal of Heat and Mass Transfer, vol. 46, pp. 251–262, 2003.
57. S.H. Seyedein and A.S. Majumdar, “Modeling of a Single Confined Turbulent Slot Jet Impingment using Various k-ε Turbulence Models,” Applied Mathematical Modeling, vol. 18, pp. 526–537, 1994.
58. J. Rambo, G. Nelson, and Y. Joshi, “Airflow distribution through perforated tiles in close proximity to computer room air conditioning units,” Presented at ASHRAE Symposia, 2006, submitted.
59. A. Bar-Cohen and W. Krueger, “Thermal characterization of chip packages – evolutionary design of compact models,” IEEE Transactions on Components, Packaging, and Manufacturing Technology – Part A, vol. 20, pp. 399–410, 1997.
60. C.J.M. Lasance, “Highlights from the european thermal project PROFIT,” ASME Journal of Electronic Packaging, vol. 126, pp. 565–570, 2004.
61. B. Shapiro, “Creating reduced-order models for electronic systems: An overview and suggested use of existing model reduction and experimental system identification tools,” in Thermal Challenges in Next Generation Electronics Systems, Y.K. Joshi and S.V. Garimella, Eds. Rotterdam, Netherlands: Millpress, 2002, pp. 299–306.
62. J. Rambo, “Reduced-order modeling of multiscale turbulent convection: Application to data center thermal management” Ph.D. Thesis, Department of Mechanical Engineering, Georgia Institute of Technology, 2006.
63. J. Rambo and Y. Joshi, “POD analysis of steady flows with parametric boundary conditions,” SIAM Journal of Scientific Computing, under review.
64. R.L. Webb and H. Nasir, “Advanced technology for server cooling,” Presented at IPACK’05 – International Electronic Packaging Technical Conference and Exhibition San Francisco, CA, 2005.
65. M. Wilson, J. Wattelet, and K. Wert, “A thermal bus system for cooling electronic components in high-density cabinets,” ASHRAE Transactions, pp. 567–573, 2004.
66. Y. Furihata, H. Hayama, M. Enai, and T. Mori, “Improving the efficiency of cooling systems in data center considering equipment characteristics,” Presented at INTELEC – 26th Annual International Telecommunications Energy Conference, 2004.
67. D. Wang, “A Passive solution to a difficult data center environment problem,” Presented at ITHERM 2004 – Ninth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, Las Vegas, NV, 2004.

68. D.L. Beaty, “Liquid cooling: friend or foe?,” ASHRAE Transactions, pp. 643–652, 2004.
69. Y. Hwang, R. Radermacher, S. Spinazzola, and Z. Menachery, “Performance measurements of a force-convection air-cooled rack,” Presented at THERM 2004 – Ninth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, Las Vegas, NV, 2004.
70. R. Anton, H. Jonsson, and B. Palm, “Modeling of air conditioning systems for cooling of data centers,” Presented at ITHERM 2002 – Eight Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, San Diego, CA, 2002.