Proceedings of IPACK'01
The Pacific Rim/ASME International Electronic Packaging Technical Conference and Exhibition
July 8-13, 2001, Kauai, Hawaii, USA

IPACK2001-15622

Computational Fluid Dynamics Modeling of High Compute Density Data Centers to Assure System Inlet Air Specifications

Chandrakant D. Patel, Cullen E. Bash and Christian Belady

Hewlett-Packard Company
1501 Page Mill Road
Palo Alto, CA, 94304-1126
Phone (650) 857-7140, Fax (650) 236-9675, chandrakant_patel @hp.com
Phone (650) 236-2748, Fax (650) 236-9675, cullen_bash @hp.com
Phone (972) 497-4049, Fax (650) 236-9675, christian_belady@hp.com

Lennart Stahl and Danny Sullivan

Emerson Energy Systems
P.O.Box 851678
Richardson, Texas 75085-1678
Phone (972) 583-5572, Fax (972) 583-7999, lennart.stahl@emersonenergy.com
Phone (972) 583-7573, Fax (972) 583-7999, danny.sullivan@emersonenergy.com

Abstract

High power dissipation from microprocessors, support chips, memory chips and mass storage has resulted in large overall power dissipation from computer systems. The deployment of these computer systems in large numbers and in very dense configurations in a data center has resulted in very high power densities at room level. These computer systems are deployed in a rack. A standard 2-meter high rack can accommodate an equivalent of 40 thin desktop systems. If the maximum power dissipation from each system is 300W, a single rack in a data center can be assumed to dissipate 12 KW. A data center can have hundreds of these 12 KW racks. Due to such high heat loads, designing the air conditioning system in a data center using simple energy balance is no longer adequate. Moreover, the data center design cannot rely on intuitive design of air distribution. It is necessary to model the air flow and temperature distribution in a data center. In this paper, a computational fluid dynamics model of a prototype data center is presented to make the case for such modeling. The model is compared with experimental results from the prototype data center.

Key words: computational fluid dynamics, electronics cooling, data center, internet data center, DataCool

Introduction

The continued development of the microprocessor has led to significant increases in input/output pads, high frequency interconnects and power dissipation[1]. System level performance requirements have resulted in computer systems that feature increasing numbers of microprocessors in very close proximity. A typical microprocessor system board contains one or more CPU(central processing unit) with associated cache memory, support chips, and power converters. The system board is typically mounted in a chassis containing mass storage, input/output cards, power supply and cooling hardware. Several such systems, each with maximum power dissipation of 300W, are mounted in a rack. The rack used in today's data center is an Electronics Industry Association (EIA) enclosure, 2 meters (78 in) high, 0.61 meter (24 in) wide and 0.76 meter (30 in) deep. A standard 2 meter rack has an available height of 40 U, where U is 44.4 mm (1.75 in). Recent market forces have driven the computer manufacturers to produce 1 U high systems. Therefore, a rack can accommodate 40 of these systems. If the power dissipation from each system is 300W, a single rack in a data center can be assumed to dissipate 12 KW.

The purveyor of computing services, such as an internet service provider, installs these rack based systems in a data center. In order to maximize the compute density per unit area of the data center, there is tremendous impetus to maximize the number of systems per rack, and the number of racks per data center. If 80 half U systems were accommodated per rack the power dissipation will reach 20KW per rack for a system assumed to dissipate 250 W. This is akin to an auditorium with 200 people per seat, assuming 100 W of power dissipation per person.

The total power dissipation from a 2-meter rack in 1990 was approximately 1 KW. A decade later, for the same rack footprint, the power has gone up ten fold. Based on the extrapolation described earlier, for future half U systems in a rack, the power can be expected to be twenty times as high as it was ten years ago. Several references [2][3][4] describe the increase in heat load in data centers. Techniques such as the creation of "hot aisles" and "cold aisles" are proposed[5]. A numerical model of a computer room is presented by Schmidt[2]. In this paper, we expand upon the modeling work by comparing results in a prototype data center built explicitly for modeling and metrology.

Nomenclature

C_x = $m_x c_{p,x}$, capacity of fluid x
c_p = heat capacity
$\dot{m}$ = fluid mass flow rate (synonymous with m)
Q = rate of heat exchange
T = temperature
ε = Q/Q_{max}, heat exchanger effectiveness

Subscripts

c = cold fluid
exp = experimental
h = hot fluid
in = in/fluid inlet (context sensitive)
min = minimum
max = maximum
out = out/fluid outlet (context sensitive)
sim = simulation (numerical results)

Motivation

Designing the air conditioning system in a high end data center using a simple energy balance is no longer adequate i.e. summing up the maximum power dissipation from the racks and sizing the sensible air conditioning capacity will not suffice. Moreover, the data center design cannot rely on intuitive distribution of air. The fluid mechanics and heat transfer processes inside the data center must be understood. To this end, it is necessary to model the air flow and temperature distribution in the data center. As noted in the introduction, there are many references on heat load issues in data centers. However, there is a lack of published work on computational fluid and heat flow modeling in data centers.

The purpose of this paper is to present a computational fluid dynamics (CFD) model of a data center. We have modeled a prototype data center constructed with a unique arrangement of heat exchangers called DataCool [6]. The modeling results of a given configuration are compared to the temperature and flow data taken from the prototype room. Our modeling focus is on determining inlet air temperature to the systems. Equally important motivation is the determination of "hot spots" in the room. Specifically, for hot locations near the system inlet at various heights.

Data Center Cooling Design

Figure 1 shows a simplified representation of a traditional data center with under floor cool air distribution. The hot exhaust air from the racks is cooled by recirculating the air through modular air conditioning units located within the room. A rack can be assumed to exhaust air with a 15 °C rise with respect to the inlet air. One can use the energy equation to get the minimum mass flow needed for a given temperature rise.

$$Q = m c_p (T_{c,out} - T_{c,in}) \qquad (1)$$

Using Eq. 1, the mass flow, and hence the volumetric flow rate for a given temperature rise can be calculated. For a 12 KW rack, built up with 1U systems, the volumetric flow rate required for a 15 °C rise is approximately 0.68 m^3/sec (1440 CFM) at sea level. Air moving devices such as axial fans deliver the required volumetric flow needed to sustain the given temperature rise. However, with the advent of slim 1U servers, the air movers tend to be small 40 mm axial fans. The densely packed 1U systems have significant flow resistance and coupled with a lack of space for accommodating the air moving devices, it may not be uncommon to observe racks that exhaust air with a higher temperature difference. The rack inlet air temperature specification is typically 30 ^{o}C at an altitude of 3000 meters. Therefore, the examination of data center cooling must begin with the premise that this inlet air specification must be maintained for each rack.

As shown in Fig. 1, the modular air conditioning (AC) units cool the recirculated exhaust hot air from the racks. A refrigerated or chilled water cooling coil in the AC unit cools the air to a temperature of approximately 10 ^{o}C to 17 ^{o}C. A typical 3 m by 0.9 m by 1.8 m modular AC unit has a maximum sensible heat removal capacity of 95 KW. The cool air is re-circulated back to the racks through vented tiles in the raised under-floor plenum. The air movers in a modular AC unit have a volumetric delivery of approximately 5.7 m^3/sec (~12,000 CFM). The air movers pressurize the plenum with cool air. The cool air enters the data center through vented tiles near the inlet of the racks. A properly devised and applied vent tile allows the air to be delivered with adequate momentum to reach the inlet of the systems located at the top of the racks.

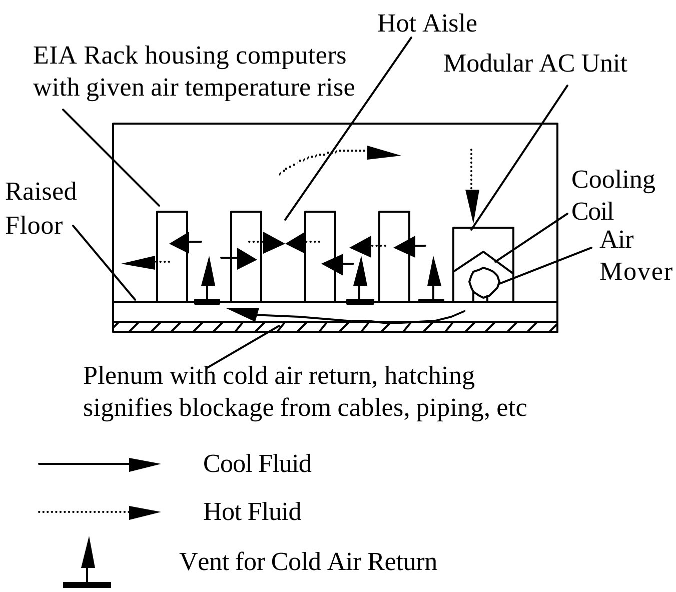

Figure 1. Typical Data Center Configuration

Upon further examination of Fig. 1 in detail, even in its simplest form without any three dimensional effects, we assert that temperature gradients and flow patterns need to be analyzed for a given layout of equipment to assure appropriate air inlet specifications to the systems. The exhaust air from the racks and the inlet air to the racks has to be managed in such a way that the cascading effect of pre-heated air does not result in violating inlet air specifications to the systems. In many cases, the preferred orientations of the rack may be such that hot air is exhausted in a common aisle – labeled "hot aisle" in Fig. 1 [5]. This would prevent the hot rack exhaust air from being drawn into the neighboring racks. This type of arrangement is useful in meeting the inlet air specifications. However, such intuitive equipment layout changes do not alleviate the need of numerical modeling to produce a thermally appropriate layout of a data center.

Alternate Approach in Data Center Cooling

Figure 2 is a simplified representation of an alternate data center cooling arrangement using modular heat exchangers in the ceiling. The product, called DataCool[6][7], was developed by Emerson Energy Systems for HP to address an industry need for a breakthrough increase in data center cooling capacity.

As shown in Fig. 2, the hot air exhaust from the racks is recirculated and cooled by a distributed set of air to liquid heat exchangers in the ceiling. The advantage of this approach is the proximity of the heat exchangers to the racks. All the rack cooling is localized with this scheme. Unique mechanical design ideas have been implemented in the heat exchangers to help in directing air flow to and from the racks e.g. the ability to shift the intake and exhaust section.

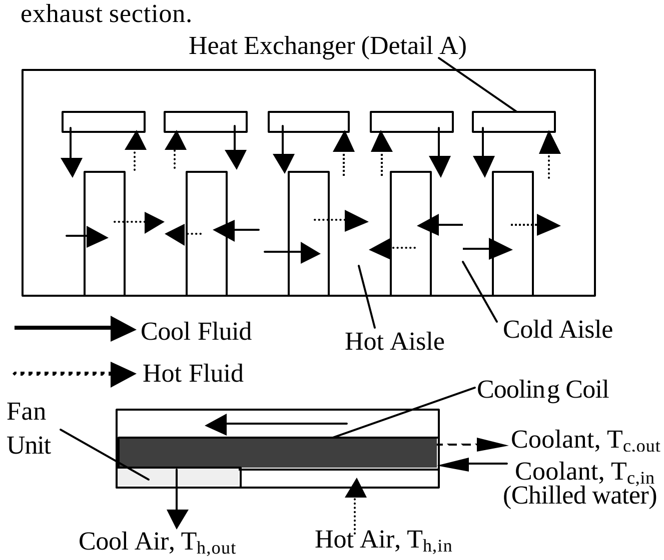

Detail A. Heat Exchanger Block Diagram

Figure 2. Alternate Data Center Cooling Design

Additionally, the modular heat exchangers offer the flexibility to scale the cooling as needed. It also saves revenue generating floor space and the raised floor can be solely used for cable distribution. Regardless of these obvious advantages, we feel CFD modeling is a must to understanding the optimal layout of a room. This paper will present modeling and metrology from an example DataCool configuration. Optimal configurations will be the subject of subsequent papers.

Prototype Data Center Geometry

A prototype data center was built expressly for the purpose of studying DataCool at Hewlett-Packard's Richardson Site. Figure 3 is an image of the prototype data center equipped with DataCool heat exchangers with rack A6 shown for reference (see Fig. 10). The room serves as a demonstration vehicle for the cooling system's capabilities and as an experimental test bed that can be utilized for parametric studies and numerical model verification. The latter is possible because the room was designed to enable changes to all significant thermal parameters such as rack heat dissipation and airflow, rack position, heat exchanger airflow etc. This unprecedented work allowed a completely controlled environment for modeling and metrology.

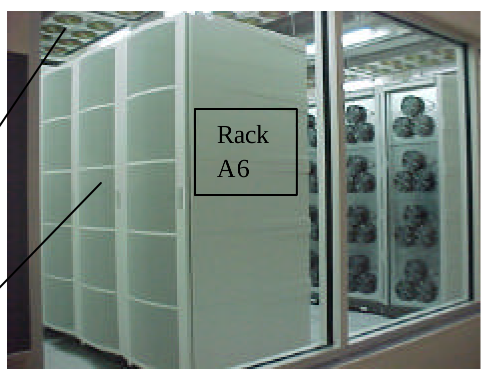

Figure 3. Prototype Data Center

The room measures 7.62 m (25 feet) by 5.49 m (18 feet) in the x-z plane and has a 3.96 m (13 feet) ceiling. There are three rows of six 2-meter racks oriented in the z-direction. Three rows of three DataCool units are placed longitudinally with the rows of racks and are utilized to cool the racks. As shown in Fig. 4, the fan trays in the heat exchanger units are oriented such that the exhaust of each tray is located on the intake side of the racks.

Each rack is partitioned into 4 compartments. Each compartment is equipped with three tube-axial AC fans arranged in parallel that deliver a combined flow rate of 0.28 m^3/s (600 CFM), resulting in a rack flow rate of 1.13 m^3/s (2400 CFM). The three rows of racks are arranged in a front-to-front, back-to-back orientation such that hot and cold aisles are created as shown on Fig. 4. Heaters within each compartment are utilized to vary the compartment load from 0 to 3600 Watts for a total rack load of 14,400 Watts. With the given geometry, airflow patterns within the room are such that cool exhaust air from the heat exchangers is ejected downward on the intake side of each rack. Heat is added to the cool air within each rack and is exhausted resulting in a temperature rise of approximately 11 ^oC. The temperature rise across the rack was set at 11 ^oC in the configuration that was tested, even though a 15 ^oC is typical. The hot rack exhaust air is drawn into the heat exchanger and cooled by chilled water cooling coil in each unit.

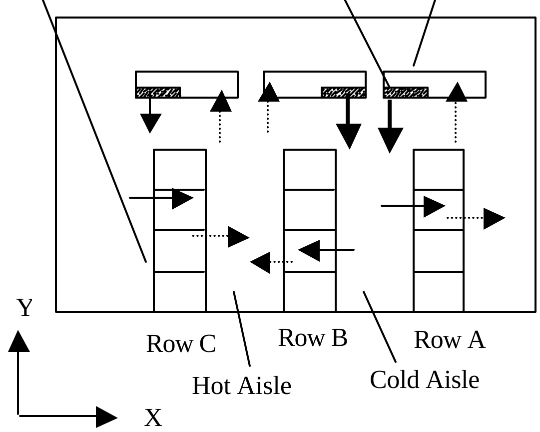

Figure 4. Simplified Diagram of the Prototype Data Center

CFD Model Construction

The modeling was conducted using computational fluid dynamics tool called Flovent [8]. In the model, the numerical computational domain was the overall room. As shown in Fig. 5, the racks were modeled as "enclosures" with an inset rectangular block called a "cuboid". The cuboid had four sets of "recirculating openings" as shown by arrows labeled 1 to 4 in Fig. 5 [8]. Each recirculating opening pair was assigned a flow of 0.28 m^3/s (600 CFM). Alternating pairs were assigned a heat load of either 0 or 3600 W such that only two compartments within a rack were at full power, and each rack was dissipating 7200 W. The rack, thus defined, was arrayed across the room with the geometry as defined in the prototype data center.

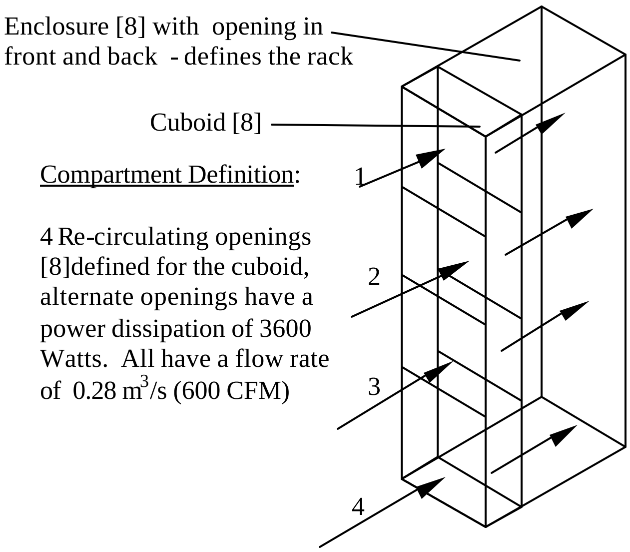

Figure 5. Simplified Definition of a Rack in the Model

The DataCool heat exchangers were modeled as shown in Fig. 6. The heat transfer attributes were identified in terms of the effectiveness, ε, of the cooling coil. The following were the key attributes:

- Heat exchanger effectiveness, ε
- Mass flow through each heat exchanger unit
- Temperature of the coolant to each heat exchanger unit, $T_{c,in}$

The heat transferred to the coolant is given by Eq. 2[9].

$$Q_{hex} = \varepsilon(mc_p)_{min}(T_{h,in} - T_{c,in}) \qquad (2)$$

where ε is the effectiveness of the heat exchanger, (mc_p) is the capacity of the fluid, the subscript *min* refers to the fluid (hot air from room or cooling fluid) with the minimum capacity, and $T_{c,in}$ is the inlet temperature of the cooling fluid[9] . In our example the hot air from the room, drawn through each heat exchanger, is the one with minimum capacity. Figure 6 shows the heat exchanger fluid flow direction. The fluid flow in the heat exchanger is either parallel flow or counter flow.

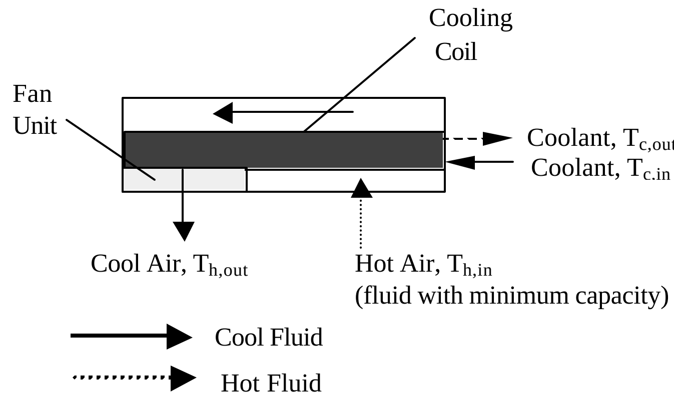

Figure 6. Heat Exchanger Definition

Figure 7 is an image of the DataCool model and Fig. 8 shows the deployment of DataCool in the prototype data center. The heat exchanger 3D geometry was created using the Enclosure, Cuboid and Volume Resistance object types from Flovent [8]. A recirculating opening [8] was applied with following heat exchanger characteristics:

- Heat exchanger effectiveness calculated using approach shown by Bash [10]: 0.60 - 0.53 (range in effectiveness due to flow arrangement, middle row has an effectiveness of 0.53)
- Air flow rate through each DataCool Unit, $T_{h,in}$, $T_{h,out}$: 2.12 m^3/s (4500 CFM)
- Inlet coolant Temperature, T_{cin}: 17 ^oC

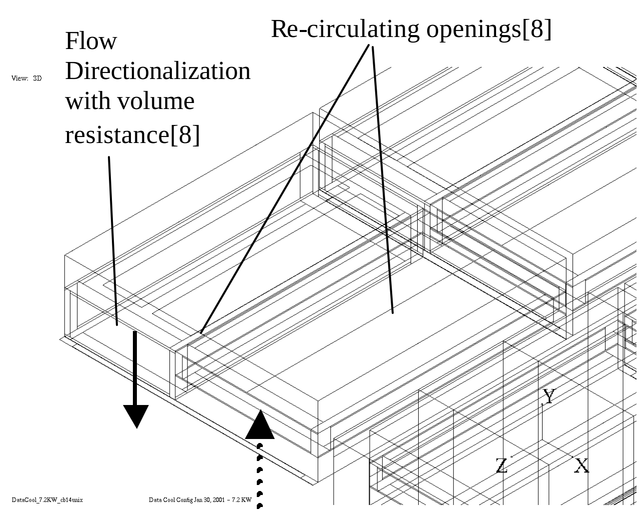

Figure 7. Representation of the Heat Exchanger in the Model

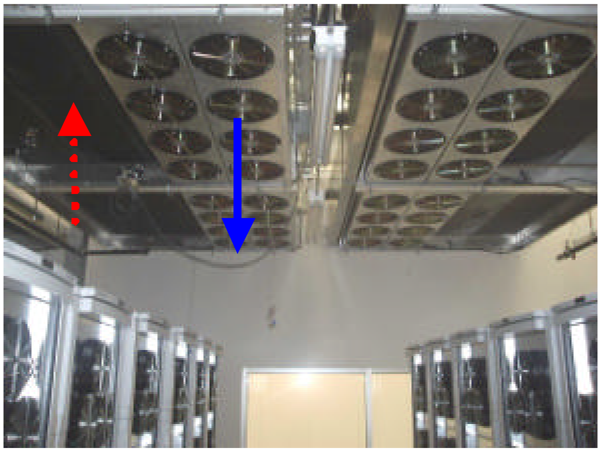

Figure 8. Image of the Heat Exchanger Modules in the Prototype Data Center

The racks and heat exchangers, thus defined, were arrayed across the room to form the room model. Figure 9 shows the overall model of the room. Moderate detail was added to the ceiling plenum and included a matrix of I-beams that affected the distribution of flow within the plenum. The room was modeled as adiabatic with no-slip boundary conditions. The revised k-epsilon model was used to account for the large scale turbulence within the room. 850,000 grid cells were arrayed across the solution domain.

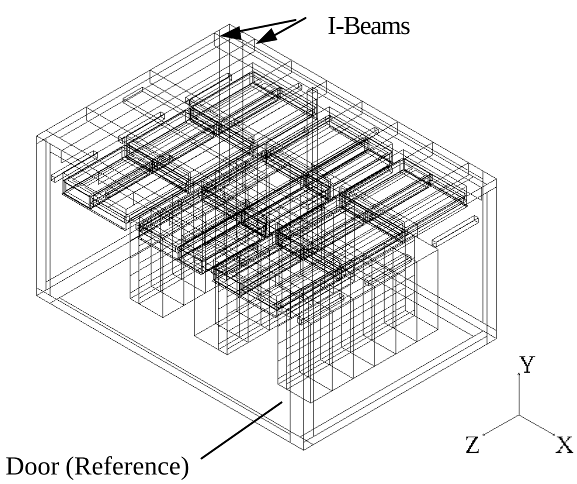

Figure 9. Model of the Prototype Data Center

CFD Modeling Assumptions and Key Measures

The CFD modeling was conducted with intent on gaining understanding of flow patterns and establishing a maximum value of inlet air temperature into the compartments modeled in the rack. The DataCool heat exchangers were allowed to operate based on the attributes defined in the earlier section. The model calculated $T_{h,in}$ and $T_{h,out}$ – the terminal air temperatures into and out of the heat exchangers. With the average air terminal temperature into the heat exchanger, one could now determine the heat extracted by each heat exchanger unit. The sum of heat extracted by all the heat exchangers should equal the heat dissipated in the room. Such an energy balance was used as a preliminary check of the modeling.

Results

The results of the simulation of the prototype data center with the example configuration are reported. The simulation results are compared with measurements obtained from the prototype data center test bed.

Figure 10 is a plan view of the room. Locations in which comparisons were made are given numerical designations on Fig. 10. In addition, racks are labeled for a subsequent comparison. Heights of 0.5 m, 1.0 m, and 2.0 m off of the floor are considered.

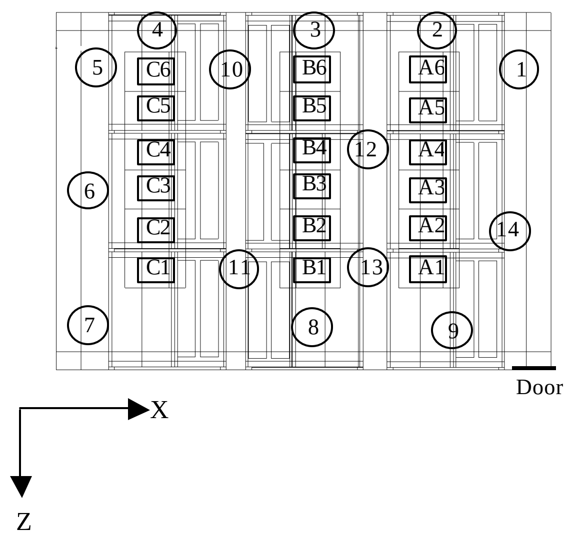

Figure 10. Plan view of room showing comparison locations.

Figures 11-13 display the results of the plan view comparisons at the indicated heights. At Y=0. 5 m, both the experimental and numerical results show hot spots in the room in areas 1, 2, 4 and 10. Thermal gradients within the room are also in general agreement, with absolute values showing less agreement. Locations 2 and 3 in particular exhibit disagreement in absolute value as well as trends. Average disagreement between numerical and experimental results is 12% at this height. The disagreement is calculated by considering the difference between the air temperature and liquid inlet temperature of the experimental and numerical results at the indicated point in the room by using the following equation:

$$Error = \frac{\left|(T_{sim} - T_{c,in}) - (T_{\exp} - T_{c,in})\right|}{(T_{\exp} - T_{c,in})} \qquad (3)$$

Similar results are observed at Y=1.0 m with discrepancies in absolute values instead occurring at points 2 and 4 and an average disagreement of 7%. Similar agreement is had at Y=2.0 meters with a disagreement of 11%. The primary areas of disagreement between the numerical and experimental results are most likely a result of simplifications made to the model in combination with the removal of incidental physical detail like tubing and windows. Windows in particular affect heat loss and local flow conditions at points 2 and 4 above Y = 1.0 m and help explain the error in absolute value and trend in those areas, especially at location 2. An additional contributor was found to be caused by the orientation of racks A5 and A6 – which were rotated counterclockwise by approximately 10° resulting in additional hot air moving to position 3. This was accounted for in the model by angling the flow exiting A5 and A6 with a planar resistance.

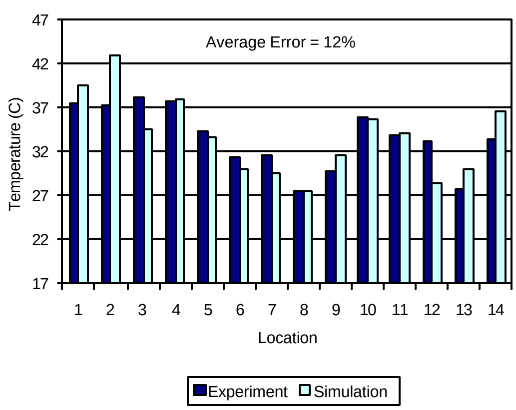

Figure 11. Temperature map at Y=0. 5 m.

A major result of the analysis, which is also in agreement with experiment, is that the portion of the room opposite the door (Locations 1 - 5 & 10 in Fig. 10) is hotter than that near the door, especially at Y = 1.0 m. This is due to the asymmetric spacing of the rack rows in the z-direction and may not be obvious, without the aid of analysis, to the individual charged with designing the room layout and cooling infrastructure.

Figures 14 - 16 compare experimental and numerical inlet temperatures for all components in the room. In compiling the numerical results, the maximum inlet temperature to each component was chosen. In the figures, components within a rack are numbered from top to bottom. (The designation of Fig. 5 is used where position 1 is the top most component, position 4 the bottom most.) The dark horizontal line on each figure can be used as a guide to quickly determine which components are near or over specification (in this case 30° C).

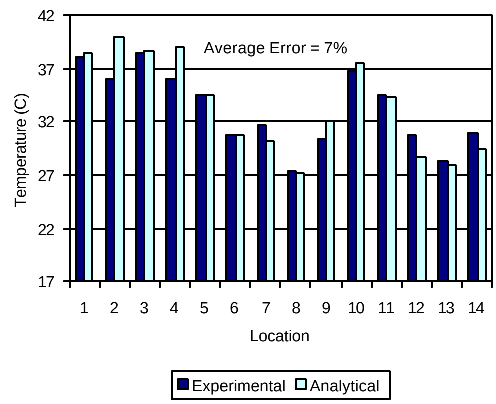

Figure 12. Temperature map at Y=1.0 m.

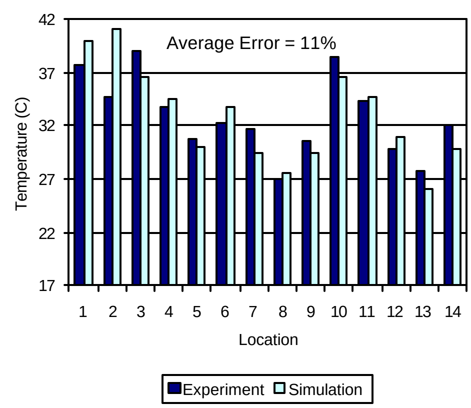

Figure 13. Temperature map at Y=2.0 m.

The results shown in Figs. 14 - 16 indicate that while absolute values show a disagreement of between 11 and 17% when averaged over each row, the simulation adequately captures the experimental pattern. The simulation also does a good job of estimating the number of components in a row that have exceeded the specified inlet temperature, even though the location of the components in the row may be incorrectly predicted. The largest areas of disagreement typically occur in the top most components near the middle racks of each row directly underneath the DataCool outlet. In these areas, especially in the aisle between rows A and B, the model predicts a cool blanket of air extending from the DataCool outlet to the level of position 1 on Fig. 5. While this trend towards cooler components was observed in the test bed, the magnitude was smaller due to increased mixing of the cool air with the surrounding warm air near the DataCool exhaust. Both inlet temperatures and rack level thermal gradients correlate moderately well with experiment and accurately indicate where improvements in the thermal management of the room can be had.

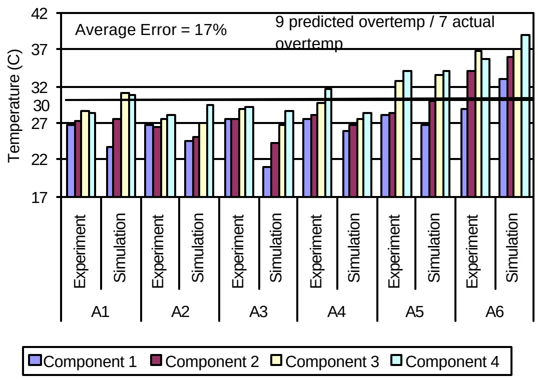

Figure 14. Row A rack inlet temperatures.

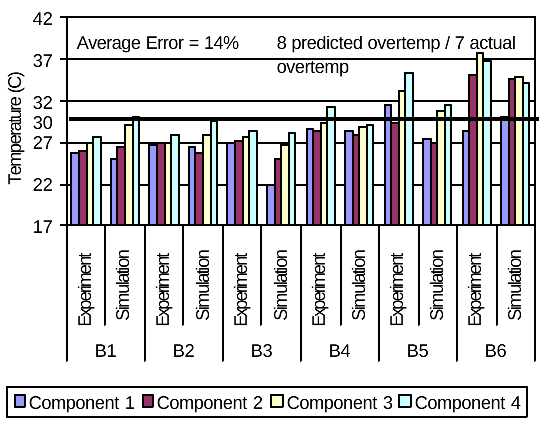

Figure 15. Row B rack inlet temperatures.

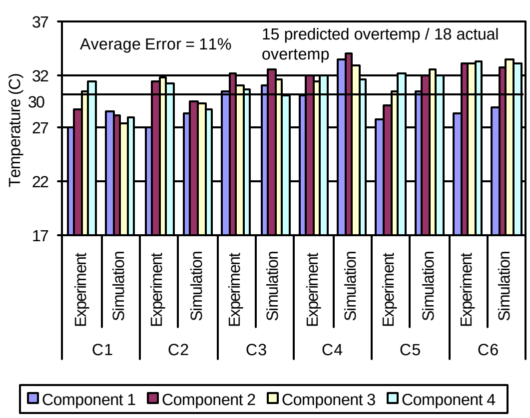

Figure 16. Row C rack inlet temperatures.

Figure 17 displays the velocity vectors from the simulation at Y=1.0 m. The local region around rack A6, the hottest rack predicted by the simulation and the second hottest measured in the experiment, is enlarged. Note the circulation pattern from the rack exhaust to the inlet as well as the stagnation point near the wall on the inlet side of the rack. Experimental measurements confirm the existence and location of each of these features. It is this circulation pattern that is responsible for the high temperatures at the inlet to A6. This could not have been determined intuitively. In addition, Fig. 17 compares the simulated velocity with experimental results at three locations around A6. Good agreement is observed.

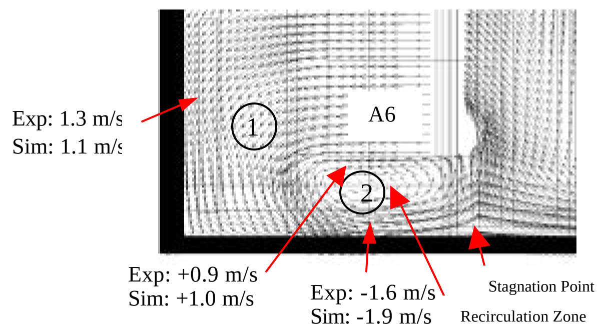

The numbered locations correspond to locations shown in Figure 10.

Figure 17. Velocity map at Y=1.0 m near rack A6.

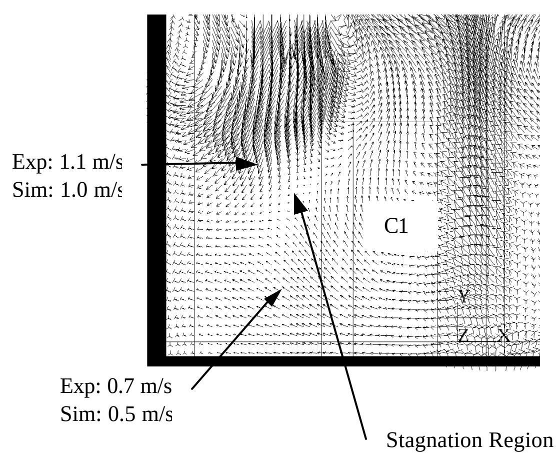

Figure 18. Velocity map at Z = 4.9 m near rack C1.

Figure 18 shows a cross section of the flow field at Z = 4.9 meters near rack C1. A diagonal stagnation region was predicted in the model at Y = 1.4 meters. This region was verified to exist at a location of around Y = 1.5 meters in the test bed. In addition, airspeed measurements were taken on either side of the stagnation region show good agreement with numerical results.

Summary and Conclusions

This paper initially outlines the power dissipation in a data center in terms of computer rack power. Computer racks of defined geometry are assumed to occupy a data center. These racks contain a number of computer systems. Typically, forty 250 W systems are said to occupy a rack. In the future, 80 of these systems are envisioned in a rack. Thus, the rack power dissipation is assumed to be 10 KW today and 20 KW in the future. In addition to this high power dissipation at the rack level, the internet data center of today has an immense physical density of racks. References that detail the data center power density are cited[2][3][4][5].

Upon examining a typical data center configuration, the authors assert that one must treat the data center itself as a computer. The racks are akin to power sources in computers. As we have in the past spent considerable effort in thermal modeling of computer systems, so must we now for data centers. Our traditional system definition has now expanded to include the room or the data center. Therefore, it is critical that we explore new ideas and technologies to match the demands of future data centers.

In response to that need, an alternative cooling arrangement called DataCool was developed by HP and Emerson. This arrangement, with heat exchangers in the ceiling, offered the advantage of space savings over conventional deployment of modular air conditioning units in a data center. A prototype data center at Hewlett-Packard Richardson site, designed specifically as a test bed for data center cooling design, was used to show that computational fluid dynamics modeling is very useful in predicting problem areas. Principally, the authors used the system inlet air temperature as the key measure of proper data center thermal design. To that end, CFD modeling was found to be useful in predicting the inlet temperatures of the systems. Also noteworthy was the fact that the modeling was useful in predicting "hot spots" in the room that are not otherwise obvious by examining the room.

The authors know of no other instance where such an effort has been undertaken to develop a controlled prototype data center for examining a data center cooling system and for data center thermal modeling and metrology. Indeed, these continuing efforts, coupled with innovations in data center design such as DataCool, are imperative to accommodate future high compute and power densities.

In a continuing effort to address this area, we plan to publish subsequent papers that will detail DataCool product design and additional modeling and metrology work that has been conducted.

References

[1] Patel, C.D., 2000, "Enabling Pumped Liquid Loop Cooling: Justification and the Key Technology and Cost Barriers", International Conference and Exhibition on High Density and Systems Packaging, Denver, pp 145-152

[2] Schmidt, R., 1997, "Thermal Management of office Data Processing Centers", Proceedings of the Interpack 1997, pp 1995- 2010

[3] The Uptime Institute, 2000, "Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment", The Uptime InstituteVersion1.0, ww.upsite.com/whitepapers.html

[4] W.P. Turner and E.C.Koplin, M.N., 2000, "Changing Cooling Requirements Leave Many Data Centers At Risk", The Uptime Institute, www.upsite.com/whitepapers.html, Computer Site Engineering

[5] R.F.Sullivan, "Alternating Cold and Hot Aisles Provides More Reliable Cooling for Server Farms", The Uptime Institute, http://www.upsite.com/whitepapers.html

[6] DataCool is a registered trademark of Emerson Energy Systems,Inc., Richardson, Texas 75085 http://www.emersonenergy.com/us

[7] DataCool brochure, "DataCool the New Thinking in Environmental Control, The Power to Cool the Most Extreme Environments", Emerson Energy Systems, Richardson, Texas 75085

[8] Flovent version 2.1, 1999, Flomerics Ltd., 81 Bridge Road, Hampton Court, Surrey, KT8 9HH, England

[9] Ozisik, M.N., 1985, "Heat Transfer A Basic , McGraw-Hill, Inc., Mexico

[10] Bash, C.B., 2000, " A Hybrid Approach to Plate Fin-Tube Heat Exchanger Analysis", Proceedings of the International Conference and Exhibition on High Density Interconnect and Systems Packaging, pp. 40-48, Denver, CO