Spring 1-1-2011

Validation of Simplified Rack Boundary Conditions for Numerical Data Center Models

Knud A. Hermansen
University of Colorado at Boulder, knudahermansen@gmail.com

VALIDATION OF SIMPLIFIED RACK BOUNDARY CONDITIONS FOR NUMERICAL DATA CENTER MODELS

by

KNUD A. HERMANSEN

B.S., University of Maine, 2003

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirement for the degree of
Master of Science
Department of Civil, Environmental and Architectural Engineering
2011

This thesis entitled:
The Development Of Simplified Rack Boundary Conditions For Numerical Data Center Models
written by Knud A Hermansen
has been approved for the Department of Civil, Environmental and Architectural Engineering

_____
John Zhai, Ph.D

_____
James VanGilder, P.E.

_____
Michael J. Brandemuehl, Ph.D., P.E.

*Date*_____

The final copy of this thesis has been examined by the signatories, and we
Find that both the content and the form meet acceptable presentation standards
Of scholarly work in the above mentioned discipline.

Abstract

Hermansen, Knud A. (M.S., Civil, Environmental and Architectural Engineering)

The Development Of Simplified Rack Boundary Conditions For Numerical Data Center Models
Thesis directed by Associate Professor John Zhai

As cloud computing and computational needs grow, data centers will continue to become a larger part of our energy load. Proper design and layout is crucial to efficient energy use in data centers. Modeling the rack is one of critical pieces in this design. Often this is done as a black box rather than modeling the rack in detail. Modeling a computer rack as a black box has been done in numerous data center studies, but rarely has it been validated against experimental temperature and velocity data. This study looks at two simplified rack models and compares them against a rack composed of four 10U server simulators. The first model is an open box model that has a heating and fan element and allows air to flow through the rack. The second model is a black box model that allows no flow through the rack and imposes a constant pressure boundary across the inlet and exhaust. The model adds the enthalpy generated by the rack load to the upwind cells at the rack inlet plate to generate the exhaust temperature profile. The models were tested across a range of airflows and rack loads. Average agreements were found to be within 3°C and 0.2 m/s over all experiments. An interesting finding of this study was the importance of correctly capturing the boundary conditions at the perforated floor tile. Modeling the perforated floor tile as a nozzle using the momentum method described in ASHRAE RP-1009 was found to produce acceptable results for airflow at the area between the rack and the perforated floor tile.

Key Words: Computational Fluid Dynamics, Data Centers, Numerical Modeling

ACKNOWLEDGMENTS

This project has had the contributions of numerous people who contributed a portion of their time and expertise to make this research possible.

First I need to thank Saleh Al-Saadi. He did the programming for the black box model on which much of the last part of this research was based. His long hours in learning how to do the ground programming for Phoenix helped make the conclusion of this project possible.

Haidong Wang and Dr. Ling Wei did much of the background research and gathered together much of the information on numerical modeling related to data centers. Their efforts greatly accelerated the start-up of this project. Both Haidong Wang and James McNeil's expertise in numerical modeling were valuable in creating the models that were used in this project. James was also extremely generous in his time and setting up the cluster to provide the computational power needed to run all of the models necessary for this study.

The set-up and running of the experiments also owes its success to many people. Jim Bright, the student lab manager for the Larson Lab at the time of this report, put in many long hours in addition to the requirements imposed by his class-work to ensure that the lab was maintained and ran efficiently. Mike Elliot from the Integrated Teaching and Learning Laboratory (ITLL) provided technical training in the welding of the thermocouples and donated the time and welding equipment required to fabricate all of the thermocouples used in the experiments. Tom Bowen, the lab safety manager for the Department of Civil, Environmental and Architectural Engineering, and Mike, from Encore Electric, worked with this research team to get the needed 208V power to the server simulators. Both Eric Tu and Eric Wilson donated materials and HOBO expertise that contributed to the set-up of the lab space and experiments. Larry Young and Vern Shultz were extremely helpful product representatives who helped to provide accurate equipment that met the needs of the research.

None of this research would have been possible were it not for the server simulators donated by James VanGilder and the American Power Corporation.

Finally, I would like to thank my advisor, John Zhai, for being available for weekly meetings and almost any other time I needed him for questions. His insightful questions and drive helped keep this project and myself on track.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

1 INTRODUCTION

A 2007 study by the EPA found that 1.5% of the nation's energy use was currently being used by data centers (consuming approximately 61 billion kWh)[1]. As current computing trends place more emphasis on cloud computing and data storage as well as the development of more compact and powerful servers, this percentage is only likely to increase. It is a safe to assert that energy use in data centers is an important field that requires attention and research if we are to have any hope of moving towards net-zero energy or "green" buildings.

Designing for energy efficiency in data centers is trickier than in office buildings because data centers are often considered mission critical facilities. This means that deviations from operating windows are avoided at almost any cost. ASHRAE's Thermal Guidelines for Data Processing Environments gives 64.4 to 80.6 °F (18-27 °C) as the recommended inlet air temperature range for Class I servers[2]. Operators of data centers are loath to do anything that could result in high inlet air temperatures or disrupt data center operations.

While the undisrupted operation of data centers is the primary concern, it is increasingly recognized that this must be balanced with the increasing power required to run these data centers in an era of rising energy costs and green building and operating requirements. The layout and design of a data center can make a significant difference in its energy use and the consequences of improper data center design can be dramatic. Cooling energy in poorly designed data centers can constitute up to 50% of their entire energy use[3]. Current standards, such as ASHRAE's *Thermal Guidelines for Data Processing Environments* and Pacific Gas and Electric's *High Performance Data Centers: A Design Guidelines Sourcebook* provide guidance on best practices[4,5]. However, even with these guidelines an iterative approach is often required to optimize performance because each data center is unique. Variables such as power distribution within racks, locations of high powered racks, depth of floor plenum, location of Computer Room Air Conditioning (CRAC) units, and numerous other variables can all be adjusted and each can have significant effects on cooling energy requirements.

Computational Fluid Dynamics (CFD) plays an important role in aiding the design and management of data centers. CFD is more economical than iterative design experimentation and testing and provides a wealth of information over the entire domain as opposed to temperature measurements at a few points. Its use in indoor environments has been well proven[6,7]. Nowadays its use in data center design is fairly common and a wealth of CFD tools are available[8].

[1] (EPA, 2007)
[2] (ASHRAE, 2008)
[3] (Uptime Institute, 2011)
[4] (ASHRAE, 2008)
[5] (PG&E, 2006)
[6] (Zhai, 2006)

While the use of CFD modeling is common in data center design, there are some important issues that need to be resolved. Due to the fact that many data centers are considered mission critical facilities, it is often hard to generate real data sets against which to compare models[9]. This also makes it extremely hard to test hypotheses under varying conditions since few data center operators are willing to allow their data centers to undergo any experimentation which could potentially disrupt operations.

One modeling issue that several studies find important, but hard to test, is the effect of rack modeling detail[10]. The issue can be seen as quantifying the tradeoff between calculation time and accuracy. The ideal solution for simplicity is a black box where room inputs are put into the front (inlet) of the rack and the added enthalpy and air velocity is generated at the back of the rack (outlet).

This thesis project seeks to generate a set of simple guidelines for producing models and to verify the room-space accuracy for temperature and velocity of these models against experimental data.

[7] (Chen & Zhai, 2004)
[8] (Rambo & Joshi, 2007)
[9] (Shrivastava, Iyengar, Sammakia, Schmidt, & VanGilder, 2006)
[10] (Zhang, VanGilder, Iyengar, & Schmidt, 2008)

2 LITERATURE REVIEW

2.1 Summary of sources

The literature review covered over 100 sources to include over fifteen primary authors. Figure 2–1 shows the data bases that were searched as well as the keywords that were used. Table 2–1 summarizes the major primary authors found in the literature review. Majority of all publications reviewed were from 2001 through 2009.

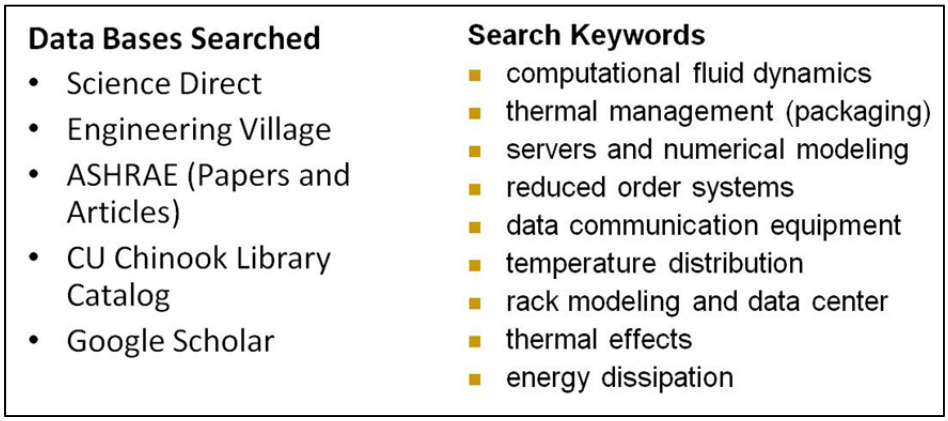

Figure 2–1: Literature Review Initial Search Parameters

Table 2–1: Summary of Primary Authors Used in Literature Review

Years	Authors	# of Publications	Affiliations
2001-2007	Roger Schmidt	13	IBM Corp., Poughkeepsie, NY, USA
2003-2007	Jeffrey Rambo	6	G. W. Woodruff Sch. of Mech. Eng., Georgia Institute of Technology, Atlanta, GA
2005-2008	Magnus K. Herrlin	4	ANCIS Incorporated,San Francisco, CA
2006-2009	Saurabh K. Shrivastava,	4	American Power Conversion (APC-MGE) Critical Power and Cooling Services, Billerica, MA, United States
2006-2008	Qinghui Tang	4	Arizona State Univ., Tempe, USA
2003-2005	Don Beaty	3	DLB Associates, Ocean, New Jersey, USA
2003-2007	Kailash C. Karki	3	Innovative Research, Inc., Plymouth, MN , United States
2005-2008	Vali Sorell	3	Syska Hennessy Group Inc., Charlotte, NC
2007-2009	XuanhangZhang, Madhusudan Iyengar	3	American Power Conversion (APC), Systems and Technology Group, Poughkeepsie, NY, United States
2006	Siddharth Bhopte	2	Department of Mechanical Engineering, State University of New York at Binghamton,USA
2008	Hendrik Hamann	2	International Business Machines Corporation, T. J. Watson Research Center, Yorktown Heights, New York, USA
2001-2002	Chandrakant D. Patel	2	Hewlett-Packard Labs., Palo Alto, CA, USA
2007-2008	Joseph F. Prisco	2	IBM, Rochester, MN
2005-2007	Amir Radmehr	2	Innovative Research, Inc.,Plymouth, MN, USA
2006-2007	James W. VanGilder	2	American Power Conversion, Billerica, MA

2.2 Typical Racks and Data Center Design

2.2.1 Servers and Racks

The Electronic Industries Alliance (EIA) establishes formal standards for 19-inch racks in their EIA 310-D (92) & EIA 310-E (2005). The standards specify rack and server dimensions to allow for the mounting of servers into standard racks. Servers come in multiples of rack units, where one rack unit, U, is 1.75 in (44.5 mm). Currently, the most popular sizes on the market are 1U, 2U and 4U[11]. A typical rack can accommodate 42 U of servers. While racks are starting to push 30kW many data centers still have racks in the 4-12 kW range.

2.2.2 Data Center Layout and Design

The standard practice for data centers is to set up the racks to create hot and cold aisles. This is done by facing the intakes of racks towards each other (or the outlets towards each other) so that any aisle is either all hot exhaust or all rack intakes. The cool air can either be distributed from the floor (most common) or from the ceiling. Figure 2–2 from (Rambo & Joshi, 2007) shows typical air distribution schemes in data centers.

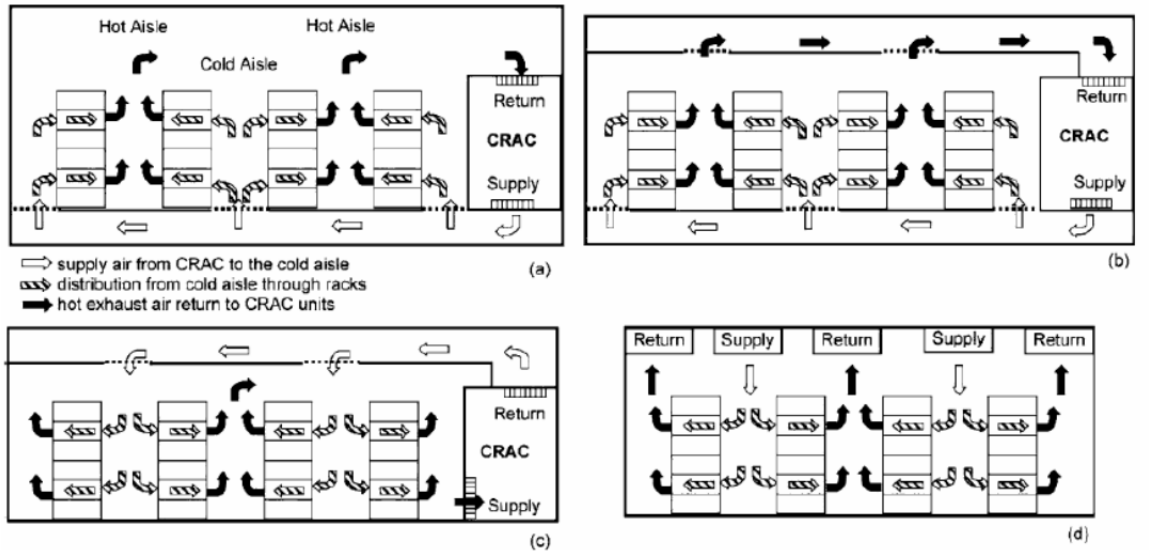

Figure 2–2: Typical Air Distribution Schemes in Data Centers[12]
(a) raised floor plenum supply and standard return, (b) raised floor plenum supply and overhead return, (c) overhead supply with room return and (d) overhead supply and return

For most data centers the air is delivered through a raised floor plenum with a depth from 12-18 in (30.5 – 45.7 cm). The perforated floor tiles may have porosities ranging from 20-40%. Typical airflow per perforated floor tile is around 500 cfm (but with significant variability possible).

[11] (42U, 2011)
[12] (Rambo & Joshi, 2007)

Table 2–2 shows some typical dimensions in data centers. Some of these are standardized and have little variability (such as rack and server widths and heights). Other dimensions such as length of the rack rows and aisle widths can have significant variability from one data center to the next.

Table 2–2: Typical Dimensions in Data Centers

	Typical Dimension
Row of racks	16 ft+ (4.88 m)
Aisle	2 to 8 ft (0.61 to 2.44 m)
Rack Width	2 ft (61 cm)
Rack Height	6.56 ft (2 m)
Ceiling height	8 to 10 ft (2.44 to 3.04 m)
Server Width	19 in (48.3 cm)
Server Height	1.75 in (44.5 mm)

2.3 Data Center Modeling

Simulating a computer rack is a modeling problem that has been tackled at many different levels. At one end of the spectrum are extremely detailed models that look at airflow and heat transfer within the server itself. At the other end are black box models which do not model any detail within the rack. For these varying levels of detail some representative papers are presented.

2.3.1 Detailed Models

Zhang et. al. (2008) examined the effects of modeling detail at three levels. They broke down the distinction in modeling between the rack and server simulators. Their most detailed modeled was called the Detailed Rack with Detailed Server Simulators (DR-DSS). In this approach all the fans, heating elements and all internal obstructions were modeled. At an intermediate level they modeled the server simulators in crude detail but maintained a high level of detail for the rack. At the simplest level they created a black box. Figure 2–3 shows representations of these three racks. They found that rack modeling detail had little influence on predicted temperatures. The same study also found rack level detail did not appear to influence the results associated with different room sizes.

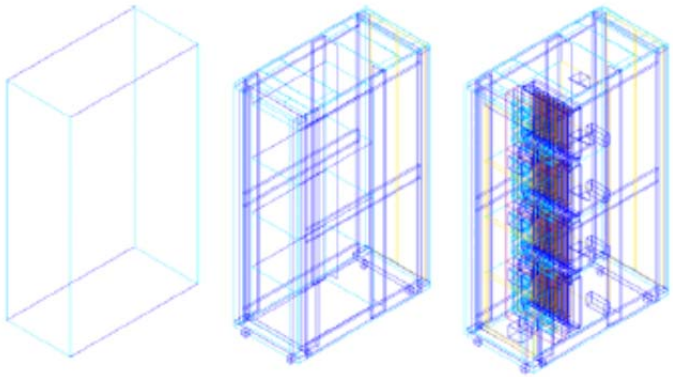

Figure 2–3: Levels of Rack Modeling Detail[13]
Left to Right: Black Box, Detailed Rack with Crude Server Simulator (DR-CSS), Detailed Rack with Detailed Server Simulator (DR-DRR)

[13] (Zhang, VanGilder, Iyengar, & Schmidt, 2008)

Rolander (2006) also used a detailed rack model to evaluate the effectiveness of a turbulent convective cooling system that delivered air from the bottom of an enclosed rack. Each server was modeled as a solid block with two chips protruding on the top (Figure 2–4). This particular paper used a 2-D model although it cited other equally detailed work done in 3-D models. This paper did not compare the modeling results against experimental data but instead compared different modeling approaches to show an improvement in heat dissipation.

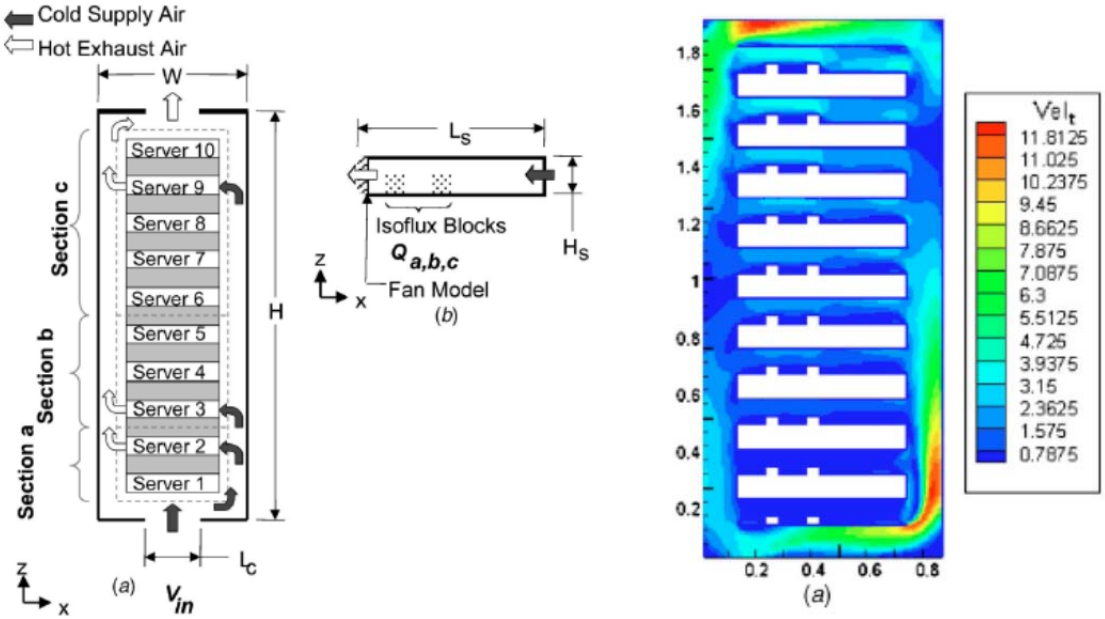

Figure 2–4: Detailed Rack Model[14]

Nelson (2007) conducted a detailed analysis of an APC server simulator of similar construction to the one used in this research. He looked extensively into modeling the fans, heating coils and all obstructions within the server simulator. He was able to obtain extensive data on all characteristics of the server simulator such as pressure drops through all components and detail on the fans. Overall he achieved reasonable and accurate results for this single server simulator.

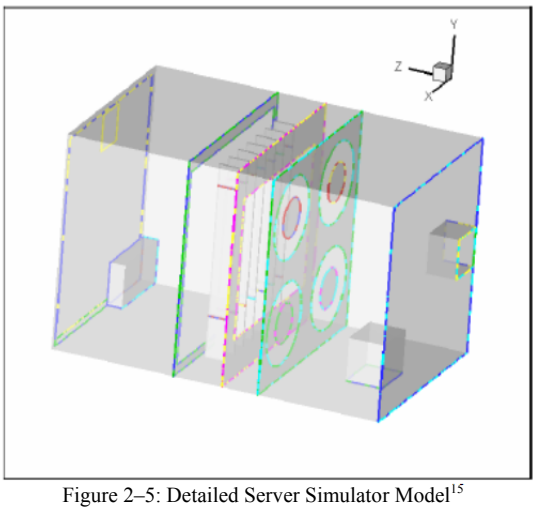

Figure 2–5: Detailed Server Simulator Model[15]

[14] (Rolander, Rambo, Joshi, Allen, & Mistree, 2006)
[15] (Nelson, 2007)

2.3.2 Intermediate Models

Herrlin, MK. (2005) used CFD to evaluate how a rack cooling index measured the effectiveness of various cooling schemes. He modeled the horizontal sections within the rack but still maintained a fairly coarse level of detail.

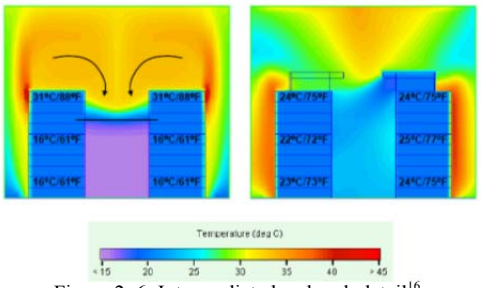

Figure 2–6: Intermediate level rack detail[16]

2.3.3 Black Box Models

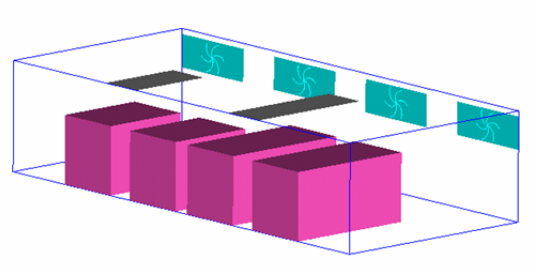

Figure 2–7: Black box model used[17]

Udakeri (2008) used black box models of racks in his study comparing overhead and underfloor air distribution systems. They imposed a uniform airflow rate across the rack and added the specified heat flux across the outlet of the rack. However, this study contrasts the two cooling system models and does not compare the accuracy to measured results.

Bhopte et. al. (2006) looked at the effects of plenum depth, floor tile placement and ceiling height on rack inlet air temperature. They examined a multivariable optimization for all three aspects. This study did not go into details about the boundary conditions for the black box rack models but they can be assumed to be similar to Udakeri (2008). Once again this was a comparison study and the results were not compared against experimentally measured conditions.

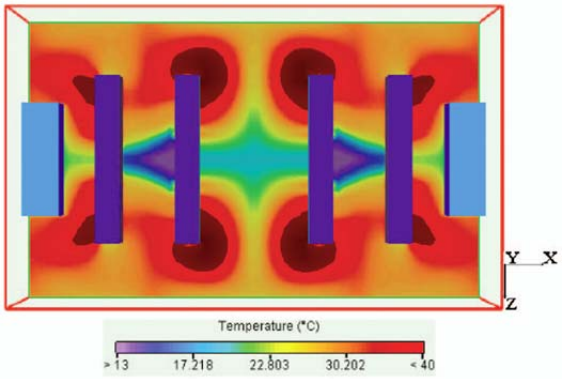

Figure 2–8: Black box model used in Bhopte et. al. (2006)

Those studies that did use experimental data from actual data centers all tended to use black box models. Shrivastava (2006, 2009) looked at the IBM facility in Poughkeepsie, NY. Prisco (2007) characterized the San Diego Super Computer (SDSC) data center located on the University of California San Diego campus. Shrivastava (2008) also used a genetic algorithm to achieve very accurate results at the data center level although this algorithm had to be trained with multiple data sets from the data center and was therefore only good for that particular data center.

[16] (Herrlin, 2005)
[17] (Udakeri, Mulay, & Agonafer, 2008)

2.3.4 Modeling Conclusions

Nearly all data center level models discussed used the standard κ-ε turbulence models, with the Boussinesq approximation in a RANS set of equations. The general trend is that for most of these studies CFD is still used as a comparison tool against which to optimize different cooling approaches both at the data center level and the rack and server level. Very few papers were able to compare their conclusions against measured data. The grid resolution practical for room dimension simulations still makes the black box the ideal approach for modeling a rack. Clearly there is still a need to quantify the accuracy of the black box modeling approach and to codify standards for creating such models.

3 LABORATORY EXPERIMENTS

3.1 Layout

3.1.1 Larson Lab

The Larson Lab at the University of Colorado at Boulder consists of a full-size commercial HVAC system, four representative commercial building zones (two zone simulators and two full-sized zones) and web-based control systems (Figure 3–9). The HVAC systems consists of an Outdoor Air Conditioning System (OACS), a main Air Handling Unit (AHU), a chiller plant (with ice storage), boiler, return fan and all of the associated ducting.

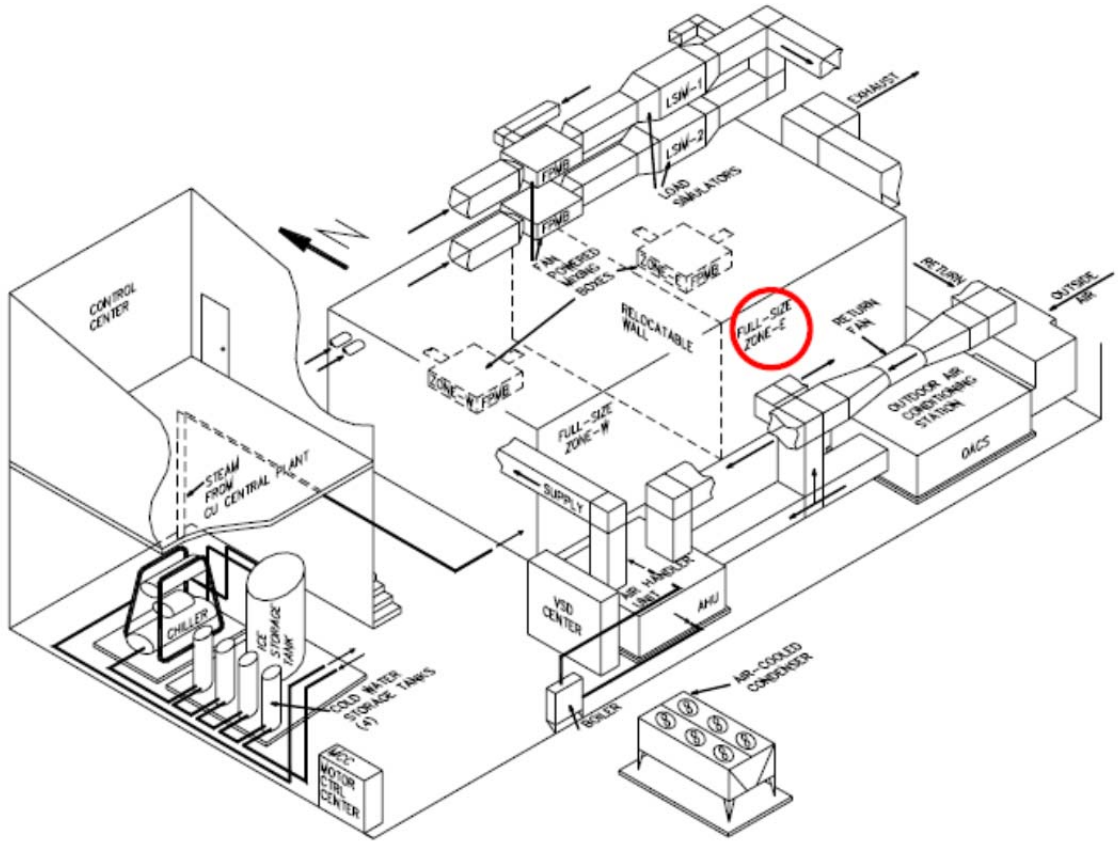

Figure 3–9: Larson Lab

3.1.2 Test Chamber

To study the airflow characteristics through the rack, a test chamber was constructed in the east zone full-size test chambers (Figure 3–9). The full sized zones consist of R-50 walls, a 15 in (38 cm) raised floor plenum and a suspended ceiling. Each zone has a fan-power mixing box (FPMB) capable of handling 1200 cfm (0.57 m^3/s).

The test chamber for the rack was built within the east zone. A partition was created with wood framing and thick black plastic sheeting for the front and left walls. All directions given in the test chamber refer to the orientation on one entering the test chamber from the door (Figure 3–10). The door itself was wood-framed with clear plastic to allow viewing into the chamber. The edges around the door were sealed with tape. The dimensions of the finished chamber were 8 ft wide, 11 ft deep and 8 ft in height (2.44 m x 3.35 m x 2.44m). Figure 3–11 and Figure 3–12 show the test chamber from two viewpoints.

Figure 3–10: Test Chamber Entrance

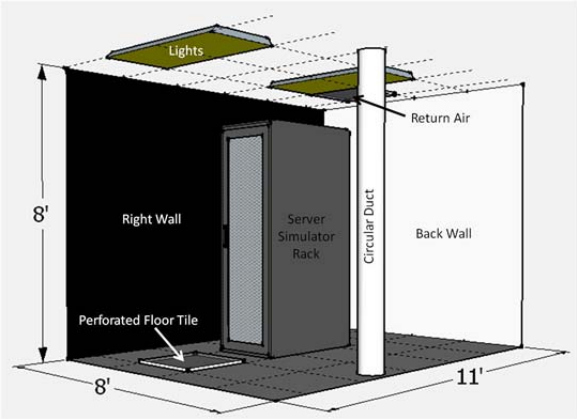

Figure 3–11: Test Chamber sketch, Front Right View

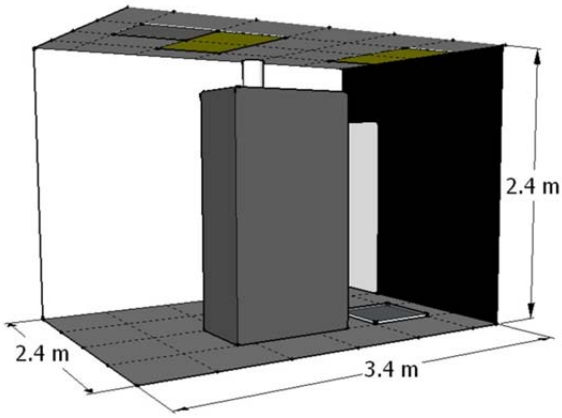

Figure 3–12: Test Chamber sketch, Back Left View

In order to reduce the complexities associated with modeling leakage, every effort to seal the test chamber for air infiltration / exfiltration. The floor tiles were all sealed around the edges with tape. Plastic sheeting was secured over the entire ceiling to minimize air exchange through the suspended ceiling (Figure 3–13).

Figure 3–13: Sealing the test chamber ceiling and floor

3.1.2.1 Air Supply

The supply air was delivered through a perforated floor tile. The dimension were a two foot by two foot (61 cm x 61 cm). It was located one foot (30.5 cm) from the front wall and two feet (61 cm) from the right wall.

The perforated floor tile itself was the AF200 CS-1003 from ASM products. It had a twenty-five percent open area. This percentage was chosen to ensure a high velocity and give a sufficient throw to get cool air up to the server simulators at the top of the rack (Figure 3–14).

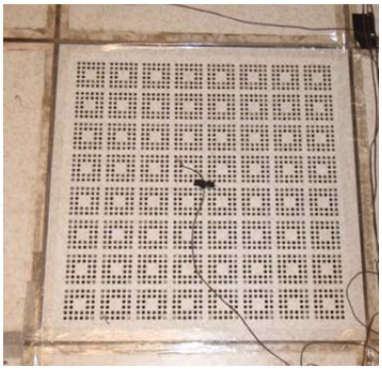

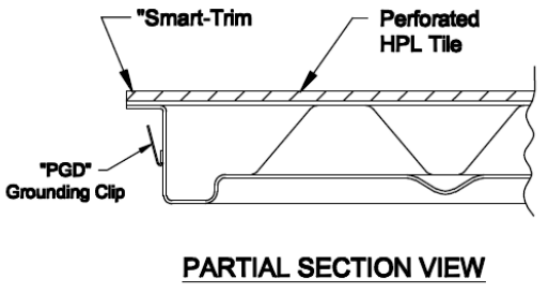

[18]

Figure 3–14: Perforated Floor Tile

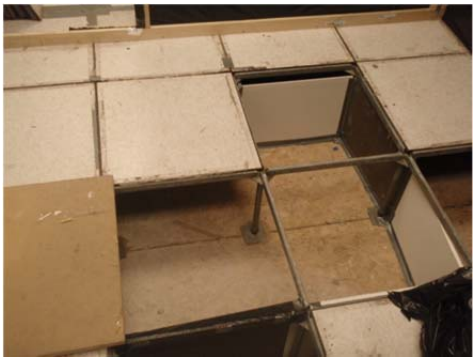

Figure 3–15: Raised Floor Plenum Partition Construction

Both the supply and return air are delivered through ducts that run through the suspended ceiling. Since most data centers use under-floor air distribution, the supply air was ducted down from the ceiling to the under-floor air distribution system (UFADS) via a circular duct which can be seen in Figure 3–11. Since the test chamber was only a partition of the entire East Zone and this team did not want to deal with leakage across the floor of the entire zone, part of the raised floor plenum was partitioned in order to make it easier to seal and maintain air pressure through the perforated floor tile (Figure 3–15).

[18] (ASM Modular Products, 2010)

The primary concern in partitioning the raised floor plenum was that this would create an uneven air distribution over the perforated floor tile. In order to address this concern the team designed the partition so that the edges of the partition were always at least one tile's length away from the perforated floor tile (Figure 3–16). The air distribution was tested using hot-wire anemometer measurements taken at multiple measuring points across the tile space and found to be uniform with maximum variations across the tile of around five percent.

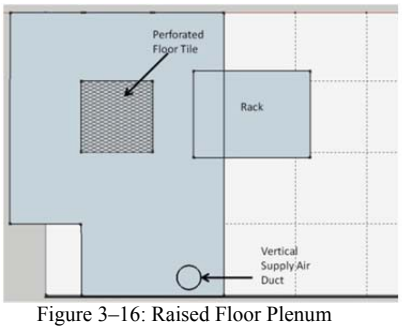

Figure 3–16: Raised Floor Plenum Partition Space

The return air grille was also two foot by two foot (61 cm by 61 cm) and was located on the ceiling, one foot (30.5 cm) from the back wall and 31.5 in (80 cm) from the left wall. It had a open area of approximately 56%.

3.1.3 Rack, Server Simulators and Major Sub-Components

While there are many components to the rack and the server simulators, this report focuses on the details used in generating the CFD models. The simplifications of the physical rack in the model will be discussed in greater detail in *Section 4: CFD Modeling*.

3.1.3.1 Rack


Figure 3–17: Rack with Server Simulators (Front View)

The rack is 24 in (58.4 cm) wide, 3 ft 6 in deep (1.067 m) and 7 ft (2.134 m) in height. The rack rests on wheels or adjustable supports that put the bottom of the first server simulator at approximately 3.5 in (9 cm) above the floor. For this report the rack refers to the entire cabinet that encloses the server simulators. The front and rear grills (rack doors) had approximately a 60% porosity (Figure 3–17).

The server simulators were 10U (17.4 in or 44.2 cm) in height and the rack was capable of mounting four server simulators. Each server simulator was approximately 19 in wide (48.3 cm) and 27.5 in (69.9 cm) deep. When mounted in the rack there was approximately a 3 in (7.6 cm) gap between the front of the server simulators and the front grille of the rack and a 1 ft (30.5 cm) gap between the back of the server simulators and the back grille of the rack.

3.1.3.2 Fans

The server simulators had four six-inch fans which were mounted onto a plate (Figure 3–18). When the server simulator was mounted in the rack, the fan plate was approximately 10.6 in (29.4 cm) from the front door of the rack. Each fan was an AC fan, model number AA1751HB-AT, manufactured by ADDA Corporation. Each individual fan drew around 37.5 W of power.

For each server simulator, the set of four fans could be set at a value of 1 through 10. This was done via a dial on the server simulator (Figure 3–19). This dial did not lock into place for each number, so setting the fan to a certain value was be done by eye and had some inherent uncertainty.

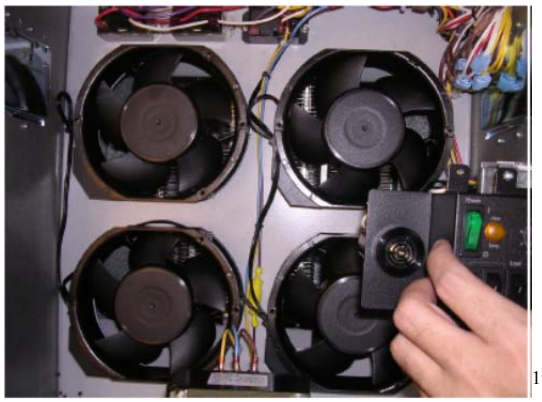 [19]

Figure 3–18: Server Simulator Fans

Figure 3–19: Fan Control Knob

The volumetric flow rates for given dial setting are shown in
. The airflow per simulator indicates the airflow generated by all four fans in a given server simulator. The airflow per rack gives a sense of what the airflow through the entire rack would be if all server simulators were put on the same given fan setting.

Table 3–3: Server Simulator and Rack Volumetric Flow Rates by Setting (IP)[20]

Fan Speed Setting	Airflow / simulator [cfm]	Airflow / rack [cfm]	Fan Speed Setting	Airflow / simulator [cfm]	Airflow / rack [cfm]
1	200	800	6	320	1280
2	215	860	7	380	1520
3	230	920	8	480	1920
4	245	980	9	590	2360
5	280	1120	10	650	2600

[19] (Nelson, 2007)
[20] Adapted from Server Simulator User's Manual, www.apc.com

Table 3–4: Server Simulator and Rack Volumetric Flow Rates by Setting (SI)

Fan Speed Setting	Airflow / simulator [cfm]	Airflow / rack [cfm]	Fan Speed Setting	Airflow / simulator [cfm]	Airflow / rack [cfm]
1	0.09	0.38	6	0.15	0.60
2	0.10	0.41	7	0.18	0.72
3	0.11	0.43	8	0.23	0.91
4	0.12	0.46	9	0.28	1.11
5	0.13	0.53	10	0.31	1.23

3.1.3.3 Heating Elements

When the server simulators were mounted in the rack, the heating elements were located approximately17.3 in (43.9 cm) behind the front grille of the rack. The power for these heating elements was provided by a three-phase, 208-V circuit. Each server simulator contained seven heating fin elements (Figure 3–20). The heating elements were controlled by five switches located on the front on each server simulator which were located below the fan control knob (Figure 3–21). These switches allowed the power for each server simulator to be set anywhere from zero to 5.75 kW at 250 W increments.

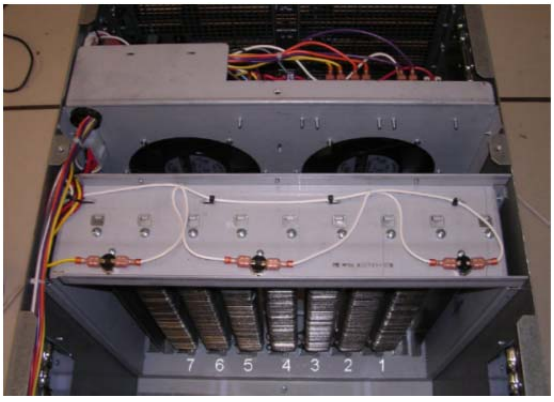

Figure 3–20: Server Simulator Heating Elements[21]

Figure 3–21: Heating Element Control Switches

It is important to note that each switch controls one or two of the heating elements so they do not create a uniform heat over the cross sectional area of the server simulator, but instead create vertical bands depending on which switch is used. Figure 3–22, taken from (Nelson, 2007), shows the breakdown of which heating element is activated by which switch. As one can see, it is also not always symmetrical.

[21] (Nelson, 2007)

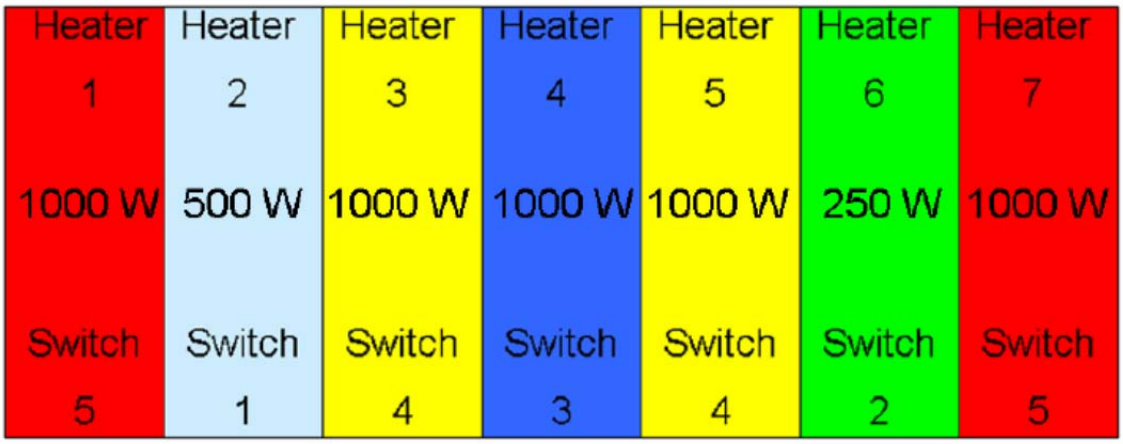

Figure 3–22: Heating Distribution in Server Simulator by Switch[22]

3.2 Data Acquisition

The data acquisition needed for this project could be broken into two categories:

1. Inputs required for the boundary conditions of the CFD model.
2. Parameters required to validate the CFD models.

Addressing the second point first, for the CFD model, following values were initially gathered:

- Supply air flow rate and temperature
- Surface temperatures
- Rack fan speed and power output

In order to validate the CFD model this team chose to consider temperature and velocity pole measurements. Additionally, this team also measured relative humidity in order to assure that test conditions were within normal operating limits as specified in ASHRAE's Thermal Guidelines for Data Processing Environments (ASHRAE, 2004).

A summary of the data acquisition devices, their purpose in the experiments and their accuracies can be seen in

[22] (Nelson, 2007)

Table 3–5.

Table 3–5: Measurement Devices, Purposes and Accuracies

Equipment	Purpose	Accuracy
Thermocouples, Type T	Temperature grid	± 0.9 °F (5 °C) or 4%
HOBO loggers	Ambient temperatures and RH	± 0.63 °F (0.4 °C), ± 2.5% RH
Alnor Flow Hood	Measure supply air volumetric airflow	± 3%
Fluke Ti-30 Thermal Imager	Surface temperature measurements	± 3.6 °F (2 °C)
Fluke Power Quality Analyzer	Measure power drawn by server simulators	± 2.5%
TSI Hot-Wire Anemometer	Measuring air velocities	± 3% over ±3 ft/min (±0.015 m/s),

3.2.1 Temperature

3.2.1.1 Thermal Imaging Camera

This team used Fluke Thermal Imaging Cameras (Ti 30 and Ti 32) (Figure 3–23) to measure surface temperatures of the rack, heating elements (Figure 3–24) and the room (ceiling, floor and walls). These temperatures were used to determine the placement of the thermocouples as well as to verify their readings.

Figure 3–23: Ti 30 Thermal Imaging Camera

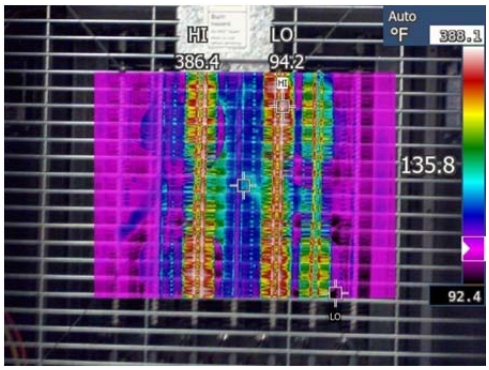

Figure 3–24: Thermal Image taken of the heating coils

It was initially decided to get surface measurements in all experiments since these can potentially be an important aspect of CFD models. The rack was set to multiple loads and run for three hours to ensure that surface temperatures had reached equilibrium. Then a series of images were taken of the walls, ceiling, floor and rack surfaces (Figure 3–25). Surfaces were broken down by regions of similar temperatures (within 4 °C (7.2 °F)) and thermocouples were assigned to each region (Figure 3–26).

The figures below show the process for the back wall. As expected there was a hot plume behind the rack and the regions to either side were cooler. The two side regions were determined to be close enough in temperature that only one thermocouple was assigned to get a temperature for both cooler regions. All other surfaces were broken down in a similar fashion.

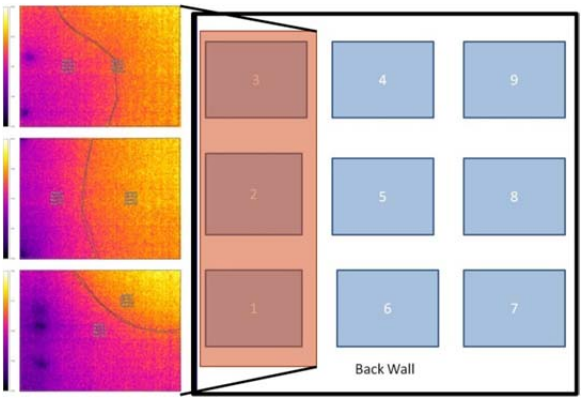

Figure 3–25: Surface Temperature Map of Back Wall

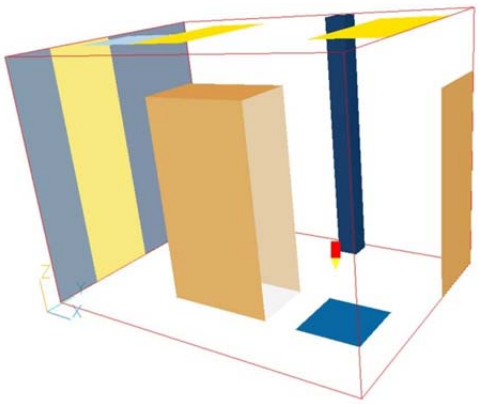

Figure 3–26: Thermocouple zones used in CFD model

4.2.1.2 Thermocouples

In addition to the thermal camera, this team looked at three different courses of action for taking space and surface temperature measurements – thermocouples, RTD (resistance temperatures detectors) and thermistors. This team decided to go with thermocouples. Thermocouples are more rugged and have a greater range than thermistors and are less expensive than RTDs. Their disadvantage compared to thermistors and RTDs is that they can be relatively less accurate[23]. However, their accuracy is within 0.9 °F (0.5) ±4% which is acceptable for what this research team needed.

This team used 32 type-T thermocouples to obtain both space and surface temperature measurements. Type-T thermocouples are a copper (positive) and copper-nickel alloy (negative) wire combination. One end of this wire pair was welded together. The two different metals generate different voltages when subjected to a temperature gradient. This is then measured with a data acquisition device (DAQ). It is important that the DAQ have a cold-junction compensation to negate the voltage differences caused at the junctions between the wires and the contacts in the DAQ.

This team chose the National Instruments 16-channel, 24-bit C Series for thermocouples (NI 9213). This provided both the required cold-junction compensation as well interfacing with LabView, the data acquisition software that National Instruments provides to the University of Colorado. Two modules were purchased to give the required 32 channels.

Initially 17 of the 32 thermocouples were placed on surfaces throughout the room.

[23] (National Instruments, 2010)

Table 3–6 shows the surface and general purpose of the surface thermocouples.

Table 3–6: Initial surface thermocouple placement

Region	Zone
Floor	Temp of cooler region behind computer
	Avg. Floor temp
Left Wall	Avg. Lower wall temp
	Avg. Upper wall temp
Back Wall	Temp of warmer wall region behind computer
	Avg temp of back wall not behind computer
Right Wall	Temp of warmer upper region of right wall
	Temp of vertical supply air duct
	Temp of cooler lower region of right wall

Region	Zone
Front Wall	Avg. Temp of door
	Avg Temp of front wall
Rack	Cooler region of rack side
	Warmer region of rack side
	Avg. temp of top surface of rack
Ceiling	Avg. temp of warmer region of ceiling
	Avg. temp of cooler region of ceiling
	Avg. temp of lights

In addition to the surfaces, thermocouples were initially placed in the supply and return air as well as on three movable poles. The poles were approximately 6.5 ft (2 m) in height and constructed of half-inch PVC. The thermocouples were placed to protrude perpendicularly approximately 6 in (15.2 cm) from the poles (Figure 3–27). Two of these poles had five thermocouples while the third pole got the remaining three thermocouples.

While there were only three poles available, each pole was given enough latitude in the length of the thermocouple wires to allow it to be moved around the room. It was assumed that while the room was not completely symmetrical in relation to the rack (due to a limitation on where the supply and return air perforated tiles could be placed) that measurements taken on one side would give a reasonable confirmation as to the accuracy of temperature predictions in the space on both sides of the rack. Therefore the initial pole positions were chosen to focus on the temperatures in front of and behind the rack. Figure 3–28 shows the initial placement of the thermocouple poles and Table 3–7 explains the reasoning behind why each pole position was chosen.

Figure 3–27: Thermocouple Pole

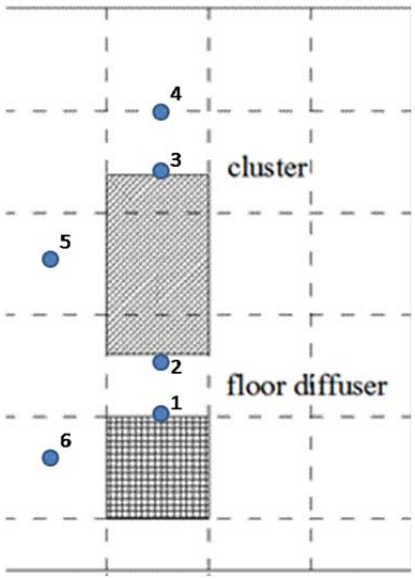

Figure 3–28: Initial Thermocouple Pole Placement

Table 3–7: Purpose of Initial Thermocouple Poles

Number	Purpose
1	Collect temperatures directly influenced by the supply air from the perforated floor tile
2	Collect temperatures entering the rack
3	Collect temperatures directly leaving the rack
4	Collect temperatures underneath the return air
5	Collect temperatures to the side of the rack
6	collect temperatures near the front of the room that have less influence from the rack, supply air or recirculation

3.2.2 Airflow

3.2.2.1 Volumetric airflow

This team used an Alnor Balometer Capture Hood (Model EBT721) to measure the volumetric flow rate from the perforated floor tile (Figure 3–29). The flowhood directs the air from the perforated floor tile across a manifold. The manifold simultaneously measures the total and static pressure at multiple points across the array. The flowhood has an automatic density correction and back-pressure compensation. The later refers to an adjustable flap that is used to correct for the pressure drop imposed by the flow hood and manifold.

Figure 3–29: Balometer Capture Hood

3.2.2.2 Point Velocity Measurements

This team was also interested in taking spot velocity measurements of airflow between the perforated floor tile and the rack inlet as well as between the rack outlet and the return air grill.

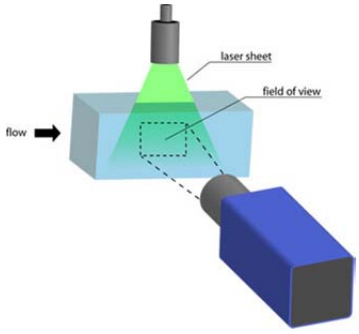

Figure 3–30: PIV[24]

Particle Image Velocimetry (PIV) was initially considered. PIV uses a laser to illuminate a plane within a fluid field that has been seeded with a visible substance (often a long-lasting stage fog). A camera takes multiple shots at a know time interval and calculates the point velocities based on the detected movement of the seeded particles. PIV is attractive because it is able to instantly measure the velocity across an entire plane within the experiment. It had to be abandoned due to space constraints within the room. It was not possible to mount the laser and camera at a distance that would allow for measurements.

Hot wire anemometers were chosen instead. Hot-wire anemometers heat a wire to a certain temperature above ambient. As air flows across the wire it cools down and the resistance changes. This team used the VelociCalc 9545 from TSI. It was chosen for its portability, ease of use and datalogging.

[24] (Tampere University of Technology, 2006)

The principle drawback to using a hot-wire anemometer is its accuracy when related to measuring the low velocities found in indoor environments. This issue is not necessarily the accuracy of the instrument (±3% or 3 fpm / 0.015 m/s) but the fact that the probe must be aligned with the direction of the flow. At low velocities or where there is turbulent and / or recirculating current (such as above the perforated floor tile near the ceiling) it can often be hard to get this aligned. Where velocities are low any minor variations can result in large perecent differences in accuracy.

With this caveat in mind the anemometer poles were placed primarily on the axis along the centerline of the rack (Figure 3–31) where airflow from the perforated floor tile and rack fan would provide airflows of sufficient magnitude to make reasonable measurements for comparison.

The anemometer poles were, like the thermocouple poles, constructed out of half-inch PVC. A double pole was used for the anemometers to allow for the probe to be secured in a horizontal position (Figure 3–32). The poles on the anemometers were extended approximately 22.5 in (57.1 cm) from the edge of the poles. The set-up allowed for the anemometers poles to be placed at multiple heights. Data was gathered at 60 Hz.

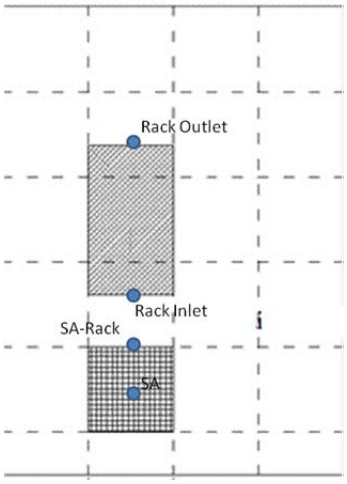

Figure 3–31: Anemometer Pole Locations

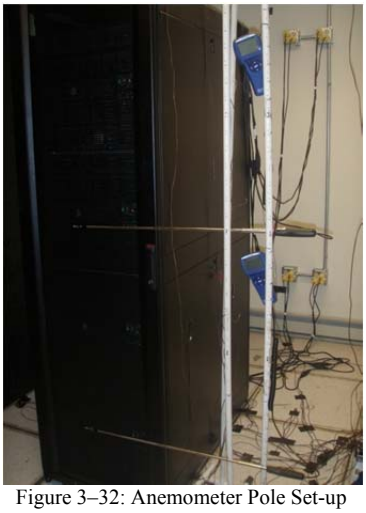
Figure 3–32: Anemometer Pole Set-up

Table 3–8 shows the purpose behind each of the locations that were chosen for the anemometer poles.

Table 3–8: Purpose and Location of Anemometer Poles

Pole	Location	Purpose
SA	Center of perforated floor tile	Measure decay of velocity from the perforated floor tile
SA-Rack	Edge of perforated floor tile closest to rack	Measure the initial influence of the rack on the air from the perforated floor tile
Rack Inlet	1 in (2.5 cm) from the front grill of the rack	Measure the velocities of air going into the rack
Rack Outlet	1 in (2.5 cm) from the back grill of the rack	Measure the velocities of air exiting the rack

3.2.3 Secondary Measurements

3.2.3.1 Relative Humidity

While not a variable that was used in the boundary conditions or validation, relative humidity was also monitored on all of the initial experiments to see where test conditions fell relative to the ASHRAE recommended guidelines for humidity in data centers. HOBO data-loggers were used to record this data. The HOBO loggers used were capable of recording temperature and relative humidity. Data was logged at 15-minute intervals.

For most experiments the relative humidity was found to be between 21 and 35% over the course of the test. This was due to both the cool temperature of the supply air provided through the perforated floor tile as well as the naturally dry conditions that exist in Boulder, Colorado. All relative humidity readings were within the allowable class I operating environment window as specified in ASHRAE's Thermal Guidelines for Data Processing Environments[25].

3.2.3.2 Rack Power

In order to confirm the amount of power drawn by the rack, the power drawn was analyzed at the panel. Each server simulator was on its own 40-A / 2-pole breaker which provided 208-V 60-Hz power. A Fluke 43B Power Quality Analyzer was used to measure the power draw of each server simulator and various settings throughout its range. Two clamps were put on each terminal to get the voltage and the ring clamp was put around one of the wire to get the current (Figure 3–33). All server simulators were found to draw the expected power (as put into the front control panel and accounting for the power drawn by the fans). The power factor was between 0.96 and 0.99 as one would expect of a primarily resistive load.

Figure 3–33: Power Quality Analyzer

3.3 Experiments

The experiments first focused on quantifying parameters of repeatability and time required to achieve steady state results. Then the research team looked at ensuring that the supply air boundary condition was adequately classified before finally moving onto examining the rack under various conditions.

[25] (ASHRAE, 2008)

3.3.1 Supply Air Quantification

In order to accurately model the performance of the perforated floor tile it was first necessary to quantify the following variables that categorize its performance:

- throw
- terminal velocity
- decay constant
- turbulence intensity

For indoor air distribution, an air jet is normally divided into four zones[26]:

1. Core
2. Transitional
3. Main
4. Terminal

For the purposes of quantifying this perforated floor tile, this team was primarily concerned with the transition region between the main and terminal jet. ASHRAE Fundamentals defines throw as:

> Horizontal or vertical axial distance an airstream travels after leaving an air outlet before the maximum stream velocity is reduced to a specified terminal velocity (e.g., 50, 100, 150, or 200 fpm), defined by ASHRAE Standard 70[27].

In these experiments terminal velocity was chosen as the average room velocity – or the point at which there was no noticeable effect from the jet. Also from chapter 33 of ASHRAE Fundamentals, we have the formula for the decay constant in zone 3, the main zone in IP units, Equation (3–1). If one converts this to the formula for SI units and solves for the decay constant K, then we get Equation (3–2).

IP
$$\frac{V_X}{V_0} = \frac{1.13 \cdot K \cdot \sqrt{A_0}}{X} \tag{3–1}$$

SI
$$K = \frac{V_X \cdot X}{V_0 \cdot \sqrt{A_0}} \tag{3–2}$$

Where:

- V_X is the velocity at point X (in this case the terminal velocity) [fpm, m/s]
- X is the height above the jet [ft, m]
- V_0 is the initial velocity in the main zone [fpm, m/s]
- A_0 is the effective area of the perforated floor tile [ft^2, m^2]

[26] (ASHRAE, 2005)
[27] (ASHRAE, 2005)

For this experiment, X could be assumed to be equal to or close to the height of the room and V_X can be assumed to be zero or close to zero. V_0 can be measured by getting an average velocity measurement across the perforated floor tile using procedure described in ASHRAE's Standard 70-2006 *Method of Testing for Rating the Performance of Air Outlets and Inlets*. The most difficult value to get for this formula is the effective area, A_0. (Chen & Srebric 2001) note that this is not necessarily the actual area. However, for this type of "diffuser", the value for the perforated floor tile can be assumed to be close to the effective area or within 90%[28].

Using Equation (3–2) with values from experimentation for X, V_0, V_X and A_0, gives a decay constant, K, between 1.5 and 1.8. This value was examined for each experiment based on the velocity measurements over the perforated floor tile. There is some inherent uncertainty is some of these values but they proved close enough to give good temperature agreements.

3.3.2 Steady State Time, Variability and Repeatability

Before validation experiments could be conducted, it was first necessary to determine at what point (if any) the experiment would reach steady state and how repeatable the measurements were. Testing on multiple rack conditions indicated that this stabilization time could be anywhere from 60 minutes (Figure 3–34) for rack with high heat fluxes to 120 minutes (Figure 3–35) for racks with lower heat fluxes and often showed show variability between sensor location. To be safe all experiments were run for 120 minutes before taking temperature readings.

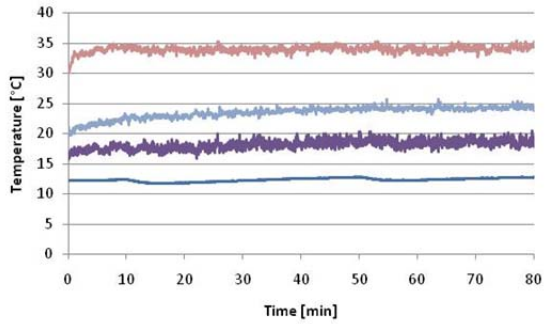

Figure 3–34: 4kW Rack, Sample Thermocouple Measurements

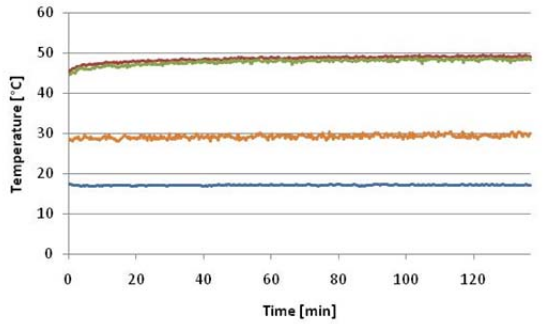

Figure 3–35: 10kW Rack, Sample Thermocouple Measurements

While all temperature and velocity measurements converged to a steady state solution, it was often found that the standard deviation around the average increased as measurements were taken higher up on the thermocouple poles. For points closest to the ceiling, one standard deviation could be as much as 0.8 °C (1.4 °F) without accounting for the accuracy of the thermocouples.

Repeatability experiments were also conducted to ensure that stable velocity and temperature measurements could be recreated and were not simply divergent, though stable, solutions to minor variations in the initial conditions. As expected, velocity

[28] (Chen & Srebric, Simplified Diffuser Boundary Conditions for Numerical Room Airflow Models, ASHRAE RP-1009, 2001)

repeatability (Figure 3–36), while achieved, had slightly more variability than repeated temperature measurements (Figure 3–36). However, all measurements were found to be repeatable within the limits of the laboratory controls and instrumentation (For reference to pole locations, see Figure 3–28 and Figure 3–32).

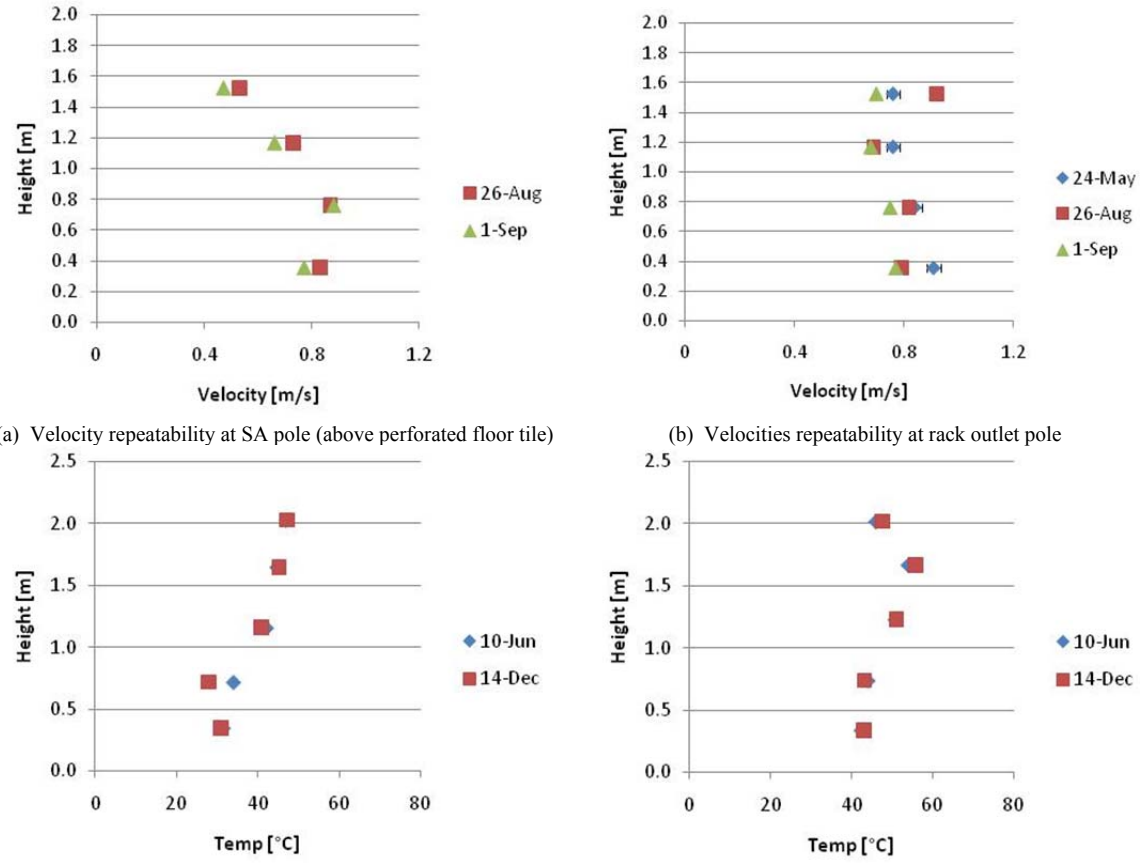

(a) Velocity repeatability at SA pole (above perforated floor tile)

(b) Velocities repeatability at rack outlet pole

(c) Temperature repeatability for 10kW Rack, Pole 2 (Rack Inlet)

(d) Temperature repeatability for 10kW Rack, Pole 4 (Under return air)

Figure 3–36: Temperature and Velocity Repeatability

3.3.3 Validating ΔT Across the Rack

Although the power drawn by each server simulator was verified at the panel by the Fluke 43B Power Quality Analyzer, it was also decided to look at the ΔT across the rack to see how closely the measured input and output temperatures agreed with the predicted ΔT that should come from a given power input.

In order to capture temperature variations horizontally, as well as vertically, thermocouple pole measurements were taken at the left edge, center and right edge of the rack inlet and exhaust faces. Each pole had seven thermocouples which were spaced evenly over the vertical distance, thus creating a three by seven point array across the front and back of the rack. While temperatures were not assumed to be uniform across the face of the rack, it was assumed that there were no extreme local temperature spikes and therefore this array of points could capture the trend of temperature variations closely enough for these calculations.

The temperatures were linearly interpolated in between the thermocouple measurements. For each matching point (same Y and Z coordinates) on the back and front faces of the rack a temperature difference was calculated and from that a differential amount of power was calculated using Equation (3–3). The differential mass flow rate for each differential cell was calculated from the fan speed settings for each server simulator.

$$\dot{q} = c_p \cdot \dot{m} \cdot (T_{ex} - T_{in}) \tag{3–3}$$

c_p – specific heat capacity of air $\left[\frac{J}{kg \cdot °K}\right]$

$\dot{m}$ – mass flow rate across the cell $\left[\frac{kg}{s}\right]$

T_{ex} – exhaust temperature for differential element on rack exhaust [°C]

T_{in} – temperature for differential element on rack inlet plate for same y,z [°C]

$\dot{q}$ – heat added across each differential element [W]

The results were added over the Y-Z domain to produce a total power across the rack face. Using the results for a rack that was set at 4kW (1kW per server simulator) with a fan speed of 0.56 m/s (800 cfm), this calculation method produced a measurement of 3.9 kW of heat (2.5 % error).

Potential sources of error here include temperature variations not captured in between the thermocouple points and air exiting through the spaces between the rack and the floor. Both of these appear to be fairly small and given that this is more of a quick check on calculations, it seems reasonable to assume that the power being put out by the rack is being reasonably captured by the array of temperature sensors at the front and back of the rack.

3.3.4 Validation Experiments

3.3.4.1 Variables

The following variables were decided upon for variation in the validation experiments:

- Rack load [kW]
- Server simulator fan speed [m/s]
- Load distribution within the rack

Rack load is the sum of the loads imposed by each individual server simulator. Server simulator fan speed is the velocity imposed by the fan units within each server simulator. While the manufacturer specified this as a volumetric flow rate, it was found to be easier to specify this as a velocity within the CFD model. Figure 3–37 compares velocities set by measuring the fan speed at the front of the server simulator versus interpreting the velocities from the flow rates given in the manufacturer's data[29]. As one can see there is fairly good agreement in data until one gets to the higher fan speeds. In all cases, the laboratory measured velocities were used.

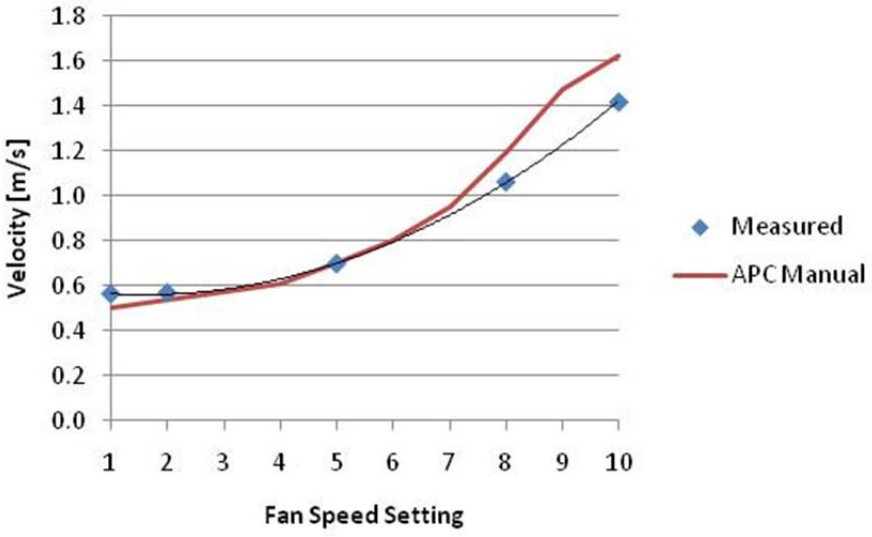

Figure 3–37: Fan speed settings by velocity

Load distribution within the rack was varied for one of the validation cases to test the effect of having different powered servers within the same rack. In this case a distribution was chosen to match, as closely as possible, (Shrivastava, Iyengar, Sammakia, Schmidt, & VanGilder, Experimental-Numerical Comparison For A High-Density Data Center: Hot Spot Heat Fluxes In Excess Of 500 W/Ft^2, 2006) which listed a enthalpy distribution of: middle 85%, bottom 6%, top 9%.

[29] (APC, 2005)

3.3.4.2 Validation Experiments

lists the validation experiments that were conducted. "Even" power distribution meant that every server simulator had the same load. For example, in the 4kW rack, each server simulator had a load of 1kW. For the 5.5kW "uneven" power distribution the server simulator power distribution, from bottom to top was 1kW, 2kW, 2kW, 0.5kW (Figure 3–38).

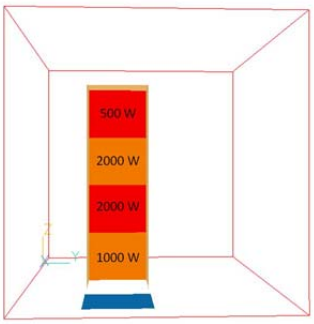

Figure 3–38: Power distribution for the 5.5kW Uneven case

The fan volumetric flow rate is given for the entire rack in both cubic feet per minute, cfm, and in meters per second. The speed was given in cubic feet per minute as a reference since this is a standard that most people working in data center design and management are familiar with. This column was taken from the APC manufacturer's data for the server simulators[30]. The fan speed given in meters per second was taken from experimental data and is the value that was used in the CFD model. The fan speed was kept constant over all four server simulators.

The supply air refers to the air delivered through the perforated floor tile. The temperature was a time average taken from thermocouple data. The volumetric flow rate was an average taken from the balometer at the beginning and end of the experiment.

Table 3–9a: Validation Experiments (IP)

Exp. No.	Rack Load [kW]	Distribution	fan speed		Supply Air		Purpose
			[cfm]	[fpm]	T [°F]	$\dot{V}$ [cfm]	
1	4	even	800	110	61.2	790	medium powered rack with even heat distribution
2	5.5	uneven	800	110	62.1	722	medium powered rack with typical uneven heating distribution
3	8	even	1920	209	53.6	804	higher powered rack with medium fan speed
4	10	even	800	110	60.1	790	higher powered rack with a low fan speed
5	10	even	2600	295	61.5	806	higher powered rack with a high fan speed

[30] (APC, 2005)

Table 3–10b: Validation Experiments (SI)

Exp. No.	Rack Load [kW]	Distribution	fan speed		Supply Air		Purpose
			[m^3/s]	[m/s]	T [°C]	$\dot{V}$ [m^3/s]	
1	4	even	0.38	0.56	16.2	0.37	medium powered rack with even heat distribution
2	5.5	uneven	0.38	0.56	16.7	0.34	medium powered rack with typical uneven heating distribution
3	8	even	0.91	1.06	12.0	0.38	higher powered rack with medium fan speed
4	10	even	0.38	0.56	15.6	0.37	higher powered rack with a low fan speed
5	10	even	2600	1.5	16.4	806	higher powered rack with a high fan speed

After the initial validation experiments were conducted and the models results validated, five additional experiments were conducted. The purpose of these additional experiments was to test out the rack under a wider variety of conditions that might be encountered in a data center. Table 3–11 lists the additional experiments that were conducted.

Table 3–11a: Additional Experiments (IP)

Exp. No.	Rack Load [kW]	Distribution	fan speed		Supply Air		Purpose
			[cfm]	[fpm]	T [°F]	$\dot{V}$ [cfm]	
6	4	even	800	110	72.0	844	Test high supply air temp
7	4	uneven	800	110	62.6	463	Low SA flow rate
8	4	even	1120	138	59.7	843	Top portion of the rack off / with blanking panels
9	4	even	1120	138	60.8	840	Top Portion of the rack removed
10	5.5	uneven	800	110	59.4	842	Test distance from perforated floor tile

Table 3–12b: Additional Experiments (SI)

Exp. No.	Rack Load [kW]	Distribution	fan speed		Supply Air		Purpose
			[m^3/s]	[m/s]	T [°C]	$\dot{V}$ [m^3/s]	
6	4	even	0.38	0.56	22.2	0.40	Test high supply air temp
7	4	uneven	0.38	0.56	17.0	0.22	Low SA flow rate
8	4	even	0.53	0.7	15.4	0.40	Top portion of the rack off / with blanking panels
9	4	even	0.53	0.7	16	0.40	Top Portion of the rack removed
10	5.5	uneven	0.38	0.56	15.2	0.40	Test distance from perforated floor tile

The first two additional experiments dealt with variations in the supply conditions from the perforated floor tile. Many of supply air temperatures used in the original validation experiments were at the lower end of the allowable window (60 °F / 15.5 °C). This experiment tested the model with a supply air temperature of 71.6 °F (22 °C) closer to the upper end of the recommended window (80.6 °F / 27 °C). The second experiment dealt with a lower volumetric flow rate through the perforated floor tile relative to that used in the original validation experiments. Most of the validation experiments had flow rates close to 800 cfm (0.378 m^3/s). For this second experiment a flow rate of about 67% (570 cfm / 0.269 m^3/s) was chosen.

The next two experiments looked at rack cabinets that were not completely full of servers. In both cases the bottom two server simulators were set at 2 kW each for a total rack load of 4 kW. The fan speed setting was put at 0.7 m/s (138 fpm) (corresponding to a fan speed setting of 5 on the front panel). In one experiment the top half of the rack was blocked with a panel (Figure 3–39). In the other experiment the two server simulators were removed and the top half of the cabinet was left open (Figure 3–40). It is important to note that Figure 3–40 is simply to show what the inside of the rack looked like. The side panel was replaced and the doors were closed during the experiment.

Figure 3–39: Top Servers Off and Blocked

Figure 3–40: Top Server Removed

For the final additional experiment, the rack was placed against the edge of the perforated floor tile (Figure 3–41). Figure 3–42 and Figure 3–43 show the altered thermocouple and anemometer pole placements. One anemometer pole was eliminated from this experiment.

Figure 3–41: Rack Against the Perforated Floor Tile

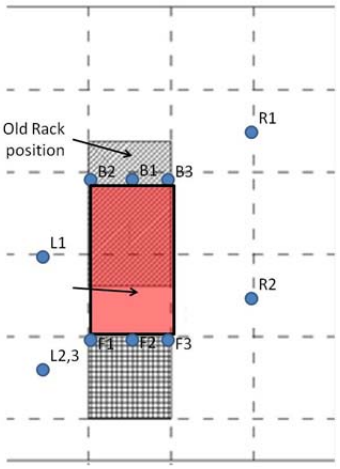

Figure 3–42: Thermocouple Poles for Rack Placed against Perforated Floor Tile

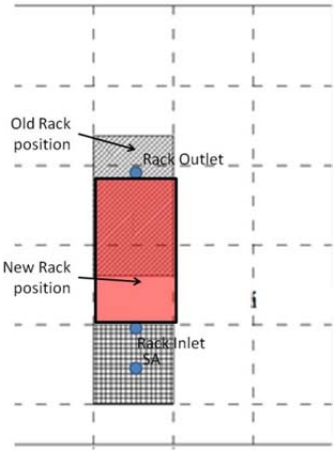

Figure 3–43: Anemometer Poles for Rack Placed against Perforated Floor Tile

4 CFD MODELING

4.1 Model Development Approach

This team developed two distinct CFD models, an open box model (OBM) and a black box model (BBM). Both models were designed to be simple and require minimal user inputs. All models were developed using the Flair module of Phoenix 2009. The OBM was developed first and its purpose was to be an interim step to inform the development of the BBM. While it was very simplified in its detail it was still an approximation of the server simulator and allowed for air to flow through the rack model. The BBM, by contrast, was a solid box. It took inputs at the rack inlet and outputs modified values at the rack outlet. Its assumptions were compared against both the experimental data and the OBM.

4.2 Open Box Model (OBM)

4.2.1 Initial Layout

Figure 4–44 shows the layout of the open box model. All sides of the rack are plates. Initially these were given a surface temperature from the thermocouple data. The front and back plates are given a specified porosity. The side plates extend to the floor. The fan and heating plate were broken up by server simulator but together form a continuous plate. The fan plate was given an X velocity. It was not modeled with any swirl. The heating plate was given a specified heat flux from both sides with each side being half of the total specified heat flux. It should also be noted that the rack model's lower coordinates begin at the floor. Even though there is 3.5 inch (9 cm) gap between the bottom of the server simulators and floor, in parametric studies there was found to be no significant difference in temperature and velocity predictions by starting the rack model at the floor.

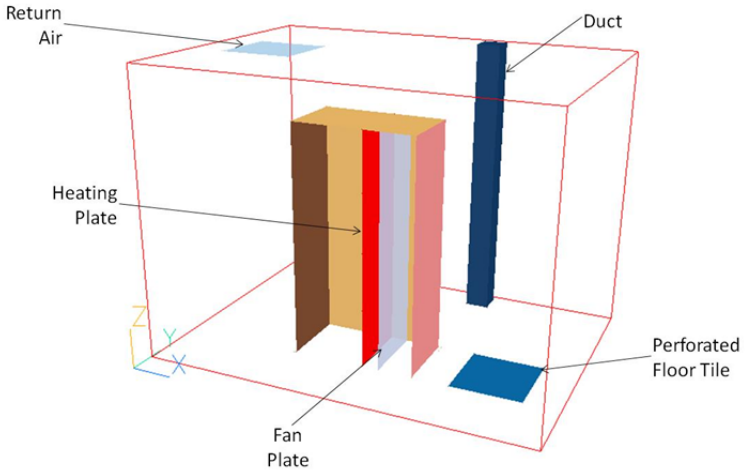

Figure 4–44: Open Box Model Layout

4.2.2 Porosity Treatment

In this study porosity refers to the percent open area at the rack inlet and exhaust. The velocity across the plate was calculated based on the porosity and the pressure differential. At the cell face the velocity is the upwind velocity times the specified porosity. The pressure drop across the plate was calculated using Equation (4–4).

$$\Delta p = \frac{1}{2} \cdot C \cdot \rho \cdot v^2 \qquad (4–4)$$

Where:

C – resistance coefficient. In these models this was specified as 1

ρ – density of air

v – velocity on the face of the plate (v_{upwind} * porosity)

Parametric runs were conducted to determine the sensitivity of the model to various levels of porosity. It was found that as long as rack inlet and exhaust porosities were within reasonable ranges (40-80%), porosity did not have any significant effect on temperature and velocity distributions (See Section 4.4.2). This was still a parameter in the models, but for general modeling guidelines, the selection is not critical as long as long as it is a reasonable value.

The Phoenix CFD software used in this project makes use of the momentum method and this has important implications for setting boundary conditions with regards to velocities at the rack inlet and outlet. The momentum method allows for different boundary conditions to be set for the mass and momentum equations. This allows for the velocity at the inlet or exhaust face to be set to one value and for the total flow rate to be set to an independent value.

For example, if the flow rate through the rack is 800 cfm (0.38 m^3/s), the rack face has an open area of 12.8 ft^2 (1.19 m^2) and a 60% open area then the actual velocity on the face would be 104 ft/min (0.53 m/s). However, Phoenix allows this to be specified as the velocity at the face (for the momentum equation) while still maintaining a volumetric flow rate of 800 cfm (0.378 m^3/s) over the surface (for the mass equation).

4.2.3 Perforated Floor Tile Modeling

This team first tried modeling the perforated floor tile as an inlet with an effective area, volumetric flow rate and temperature. The effective area, similar to porosity, adjusted the velocity based on the total area of the object and the volumetric flow rate. This method was found to significantly overestimate the velocity at the center of perforated floor tile and the edge of the perforated floor tile (Figure 4–45).

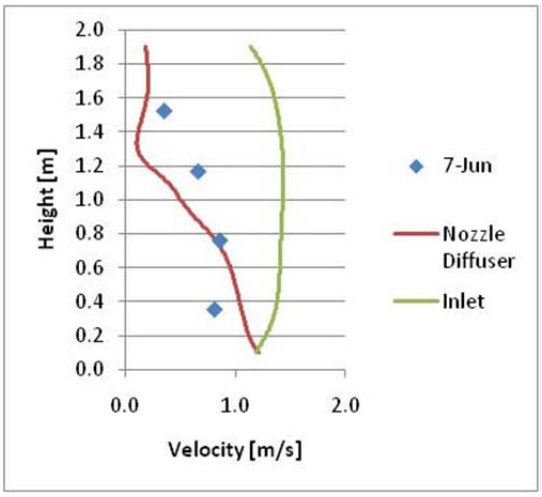

(a) Velocity profile over center of perforated floor tile

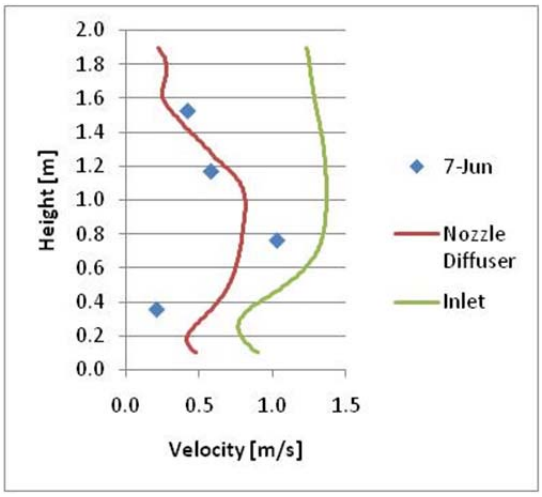

(b) Velocity profile at edge of perforated floor tile

Figure 4–45: Modeling the perforated floor tile as a simple inlet

The Flair module of Phoenix gives five different diffusers that can be modeled: round, vortex, 4-way regular/directional, Grille/Nozzle and Displacement[31]. While none of these diffuser models was designed to model a perforated floor tile for a data center, the grille/nozzle was found to most closely describe the observed airflow. A perforated floor tile can be thought of as a grid of nozzle diffusers with an angle of 0°.

All diffuser models in Phoenix are implemented based on the "Momentum method" described in ASHRAE Report RP-1009[32]. The momentum method specifies mass, momentum, energy and concentration at the supply diffuser surface – in this case, the perforated floor tile. It avoids a detailed representation of the perforated floor tile geometry. It is not as accurate for the area immediately outside of the perforated floor tile (as compared to the remainder of the space) and the authors of RP-1009 noted that it did have a tendency to overestimate velocities in the immediate vicinity of the nozzle diffuser[33]. However, once the jets from the individual "nozzles" have merged, this model gives a much better representation of the airflow and decay in the main section of the jet flow.

The perforated floor tile model was somewhat more difficult to calibrate in that it requires a series of measurements with anemometers (or preferable PIV) to determine the initial velocity and decay. The temperature was taken from a thermocouple reading over the course of the experiment. The volumetric airflow for each experiment and model was determined by using the flow hood. The variables of V_X and A_0 were kept constant for all models. The variables of K, V_0 and X were, calibrated for each experiment. While they were often close to the initial values as determined in Section 3.3.1, it was found that factors like rack volumetric flow rate and distance to the perforated floor tile had some minor effects on these variables. The exact relationships were not however explicitly determined in this study.

[31] (CHAM, 2010)
[32] (Chen & Srebric, Simplified Diffuser Boundary Conditions for Numerical Room Airflow Models, ASHRAE RP-1009, 2001)
[33] Ibid.

4.2.4 Effects of Radiative Heat Transfer on Thermocouple Points

All of the sources reviewed for this background research looked at forced convection as the dominant form of heat transfer. Only one paper, (Rambo 2005) even mentioned it – and only to then to say that it was ignored to simplify the model. For this primary reason, radiative heat transfer was largely ignored during the experimental phase of this project. One particular area that was not quantified experimentally was the effect of radiative heat transfer on the thermocouples. While the effect of radiation on the thermocouple measurements is arguably small, its effect on uncertainty was never tested experimentally and so it is an un-quantified possible source of error for some thermocouples that had a high view factor of the heat fins or other higher temperatures surfaces on the race.

While the effects of radiative heat transfer on the thermocouples was never measured, the effects of radiative heat transfer was examined during model development. The building module of the Phoenix CFD software (FLAIR) uses the IMMERSOL radiation model, where IMMERSOL stands for immersed solids. The IMMERSOL radiation model was developed as method for radiative transfer involving numerous sources where calculating the form factors was not economically feasible or realistic.

The equations for the IMMERSOL model involve adding a radiation temperature variable to the energy equation that is tracked for cells within the fluid medium (air, for this case). This variable, which the software calls T3, is defined as the fourth root of the radiosity divided by the Stefan-Boltzmann constant.

The local medium (air) temperature is set as equal to the solid's surface temperature at the solid-fluid boundary cells. The T3 variable is calculated at the air-surface boundary based on:

- the radiosity
- surface temperature
- the absorptivity of the air
- the scattering coefficient of the air
- the distance between adjacent walls

Then between cells within the fluid, the T3 variable is calculated based on the specific heat capacity and the thermal conductivity. A more complete treatment of the equations used to solve the IMMERSOL radiation model is available in the Phoenix documentation[34].

[34] CHAM. (n.d.). *The IMMERSOL model of Radiative Heat Transfer*. Retrieved July 5, 2011, from POLIS: http://www.cham.co.uk/phoenics/d_polis/d_enc/enc_rad3.htm

The software user defines emissivity values for all surfaces. For simplicity, these were assumed to be 0.95 for all surfaces. The effects of radiation were checked against Experiment 2 (5.5kW rack with uneven power distribution) and 4 (10kW rack with low fan speed) (See Table 3–10 for full details). Without accounting for radiation, the model predicted surface temperatures roughly within 1-2 °C of those observed for the rack, floor, wall and ceiling based on thermocouple measurements and spot checks done with a thermal imaging camera. For both the OBM and BBM, the difference between model that considered radiation versus the model that was run without the radiation was less than 0.8 °C for all thermocouple pole points. Figure 4–46 shows a representative pole comparison for the Back Pole Averages and L1 on the 10kW Rack. These poles were chosen because they should have been the most effected by radiative heat. The thermocouples on the back poles could "see" the surfaces of the heat fins and the Left-1 pole could see the side surface of the rack. Even for these poles, which should have been most affected by radiative heat transfer, on one of the highest powered racks used in this study, one can see that there are minimal effects on the thermocouple points.

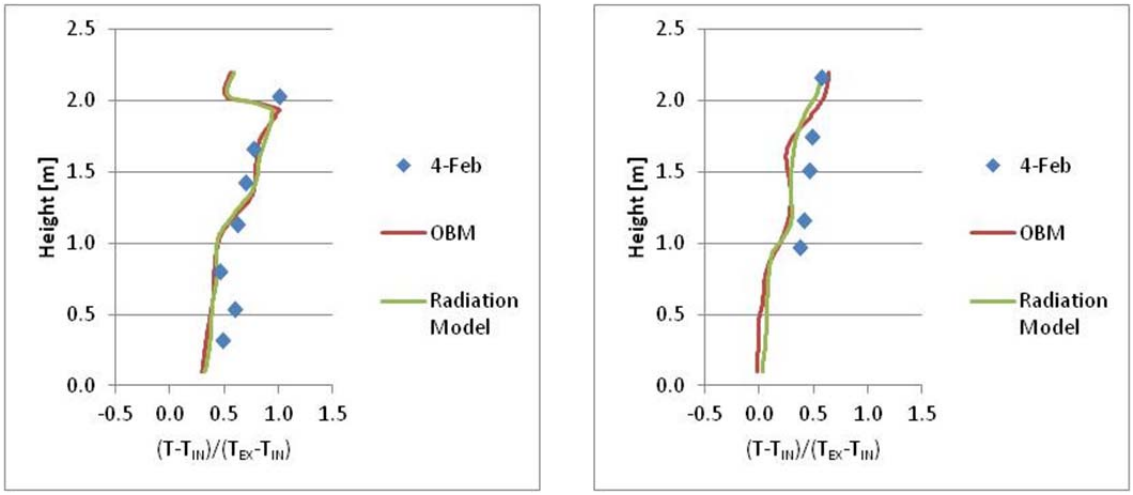

(a) Back Pole Avg. For Experiment 4 (b) Left-1 Pole for Experiment 4

Figure 4–46: Effects of Radiation on Thermocouple Readings

4.2.5 Mesh Generation

Phoenix uses a structured mesh. Objects within the domain can either be made to affect the grid, meaning that the edge of a cell will border the plate or one can chose to have them not affect the grid and instead use a partial solid treatment. With the exception of the vertical duct (Figure 4–44), heating plate and fan, the grid conformed to all surfaces.

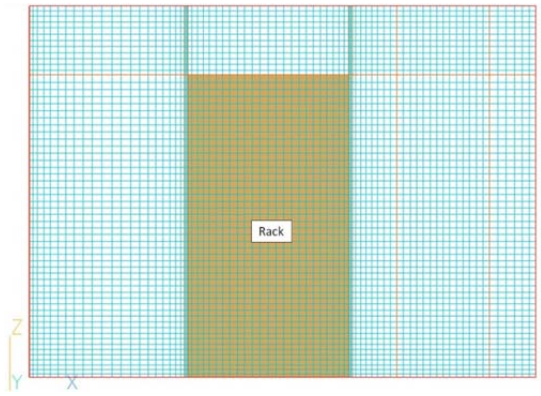

Figure 4–47: Y-Plane mesh at middle of rack

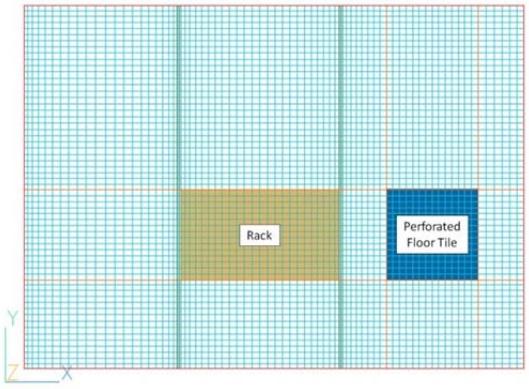

Figure 4–48: Z-Plane mesh at middle height of rack

While a power factor was used to refine the grid around the walls, this study did not focus on quantifying precisely what this refinement should be. Most practitioners of CFD related to data centers did not feel that this study should focus on impingement and other issues related to flow near walls.

The most significant issue came when trying to compare temperature and velocity data at the rack inlet and outlet. Both the thermocouple poles 2 and 3 (Figure 3–28) as well as the rack inlet and rack outlet anemometer poles (Figure 3–31) were within an inch of their respective grills on the rack. Since the inlet and outlet grills for the rack had a certain porosity there is a velocity increase that occurs as the air passes through the grill.

In reality this occurs within millimeters of the grill. In the CFD program this takes four to five cells to resolve. The initial unrefined grid had the inlet pole within the first cell of the rack inlet (Figure 4–49) causing it to over-predict velocities when compared with experimental results. The grid had to be refined at both the rack inlet and outlet so that the CFD pole readings were at least four cells away from the rack inlet and outlet.

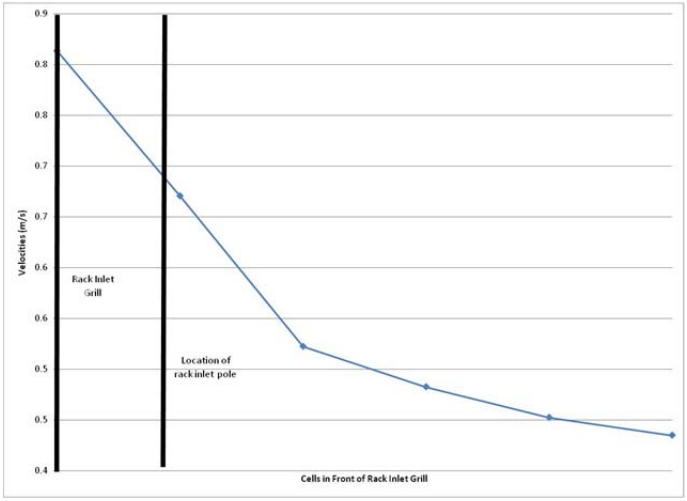

Figure 4–49: Velocity in Front of Rack Inlet Grill with Unrefined Mesh

4.2.6 Grid Independence

The normalized root mean squared error (NRMSE) was used to analyze the results of different levels of meshes to find the grid independent solution. Equation (4–5) shows how the NRMSE is reached.

$$NRMSE\left(\overline{P^c}, \overline{P^{2c}}\right) = \sqrt{\frac{\sum_{i=1}^{n}\left(p_i^c - p_i^{2c}\right)^2}{\sum_{i=1}^{n}\left(p_i^f\right)^2}} \quad (4–5)$$

$$p_i^c - \ prediction\ of\ coarse\ grid\ at\ point\ i$$
$$p_i^{2c} - \ prediction\ of\ doubled\ grid\ at\ point\ i$$
$$p_i^f - \ prediction\ of\ fine\ grid\ at\ point\ i$$

For grid independence analysis, the rack inlet anemometer pole was used. This was due to the fact that there is more variability in the velocity measurements versus the temperature measurements and if sufficient accuracy and resolution due to grid refinement is achieved here it can be reasonably assumed to apply for temperature as well. Grids were chosen by roughly doubling the number of total cells in the domain with the exception of the first interval (72,000 to 244,800). The cells were proportionally increased in each of the cardinal directions.

Table 4–13 shows the results of the grid independence analysis. The third column shows the percentage of computational time that the lower cell count had versus the highest (1,150,000 cells). Interestingly the NRMSE stayed around the same range and actually went up slightly. This team ultimately decided to go with 244,800 cells since this still had a reasonable computational time but would give greater resolution than 72,000 cells.

Table 4–13: Grid Independence Analysis

Cell Comparison	NRMSE	Computational Time
72,000 vs. 244,800	0.084	6%
244,800 vs. 576,000	0.109	25%
576,000 vs. 1,150,000	0.102	100%

4.2.7 Adiabatic Surfaces and Pole Refinement

As mentioned above, the team initially wanted to capture surface temperatures since this is a common boundary condition in CFD models. However, given the finite number of channels available in the DAQ, it was also thought that if specifying these surface temperatures did not significantly affect the results, that it would be more beneficial to experiment validation to use these thermocouples on additional poles and on additional resolution for the current poles.

The original models (with surface temperature boundary conditions) were compared to models with all surfaces (ceiling, walls, floor) being specified as adiabatic surfaces, to include rack surfaces. It was found from the results that less than 3% of heat flux in the system was through these surfaces. This was well within the error in the model itself and greatly simplified the model as well.

Figure 4–50 shows the pole locations that were used for the validation experiments. One item that is important to note is the three thermocouple poles in the front and the back of the rack. It was found that since the heat was not distributed evenly across the seven heating elements that it was necessary to get an average temperature across the rack outlet. For the rack inlet the average was necessary because warmer air recirculated from the sides and there was not a constant temperature profile across the front of the rack horizontally or vertically.

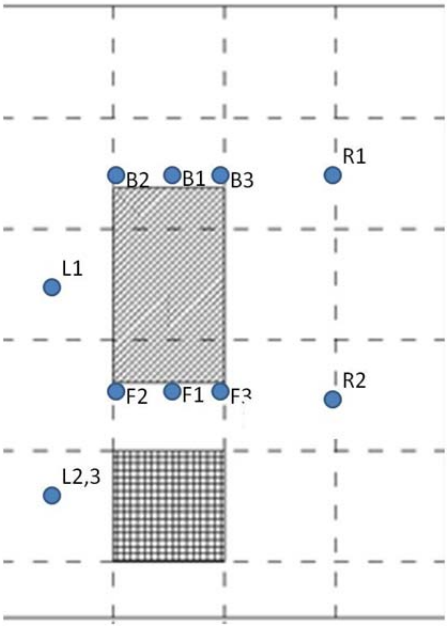

Figure 4–50: Final Thermocouple Pole Locations

4.2.8 Other modeling details

It was found that 1000 iterations produced sufficient resolution. Running an additional 1000-4000 iteration was found to produce around 0.1°C improvement in average temperature errors and negligible improvements in velocity errors. Given the increase in computational time, this was determined not to be worthwhile.

4.3 Black Box Model Development

While the author did not develop the programming for the black box model (BBM) the concepts behind this model are central to what was researched in this thesis and were developed jointly with Saleh Al-Saadi and Dr. John Zhai.

The boundary conditions outside of the rack were the same for both models. In other words, conditions for the perforated floor tile, room exhaust, surfaces (floor, walls and ceiling) and mesh structure were kept identical.

4.3.1 Boundary Conditions for the BBM

Setting up the boundary conditions for the BBM was codified as a series of steps. The intent was for these steps to be simple enough to be implemented in any numerical modeling program and require simple inputs for the rack.

Step 1: Set velocity boundary conditions at rack inlet. In order to do this one needs a reasonable estimate of the porosity of rack cabinet door (percent open area), a reasonable weighted average for the flow rate through the rack and the cross-sectional area of the rack. For the rack porosity any value (between 0.5 to 0.75) can be assumed. The velocity on the inlet face is set uniformly across the rack inlet face / rack cabinet door as shown in Equation (4–6).

$$V_{face} = \frac{\left(\frac{\dot{V}_{rack}}{A_{rack}}\right)}{plate\ porosity} \tag{4–6}$$

$$V_{face} - velocity\ normal\ to\ the\ rack\ inlet\ or\ exhaust\ at\ the\ plate$$
$$\dot{V}_{rack} - weighted\ average\ of\ volumetric\ flow\ rate\ through\ the\ rack$$
$$A_{rack} - \ Cross\ sectional\ area\ of\ the\ rack$$
$$plate\ porosity - expressed\ as\ a\ percent\ open\ area$$

As noted in Section 4.2.2, if the inlet / exhaust face velocity is specified as a boundary condition for the momentum equation one must be careful that this is separated from the volumetric flow rate as related to the boundary conditions for the mass equation. Alternatively, one could simply use the volumetric flow rate through the rack which as noted in Section 4.2.2 will produce acceptable results.

$$V_{face} = \left(\frac{\dot{V}_{rack}}{A_{rack}}\right) \tag{4–7}$$

$$V_{face} - velocity\ normal\ to\ the\ rack\ inlet\ or\ exhaust\ at\ the\ plate$$
$$\dot{V}_{rack} - weighted\ average\ of\ volumetric\ flow\ rate\ through\ the\ rack$$
$$A_{rack} - \ Cross\ sectional\ area\ of\ the\ rack$$

Step 2: Set the temperature boundary condition at the rack inlet. For any given location on the rack face, this is equal to the temperature of the upwind adjacent cell (as shown in Figure 4–51). In other words, the temperature at any location on the rack inlet is equal to the entering air temperature.

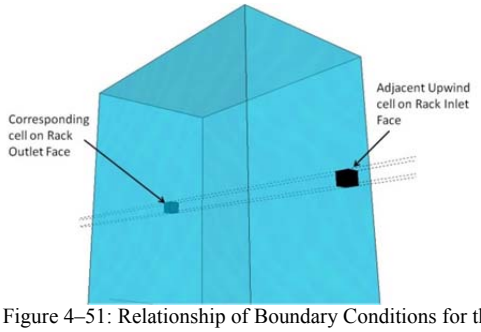

Figure 4–51: Relationship of Boundary Conditions for the BBM

Step 3: Set the velocity boundary condition at the rack exhaust face. In order to conserve momentum, this is set as the same uniform velocity which was specified for the rack inlet.

Step 4: Set the temperature boundary condition at the rack exhaust face. Here one takes the temperature from the rack inlet cell at the same vertical and horizontal position on the rack inlet face (Figure 4–51) and adds the appropriate amount of heat based on the thermal load of the server. One must know the mass flow rate which can be determined from the volumetric flow rate of the fans in the server and the density of air for the server location. Equation (4–8) shows the relationship used across each server (or across each cell).

$$T_{ex} = T_{in} + \frac{q_{server}}{c_p \cdot \dot{m}} \tag{4–8}$$

T_{ex} – exhaust temperature for cell on rack exhaust plate [°C]

T_{in} – temperature at rack inlet plate for same y,z [°C]

q_{server} – heat added by server [W] (for BBM total heat generated by each server is assumed to be evenly distributed over the server cross sectional area)

c_p – specific heat capacity of air $\left[\frac{J}{kg \cdot °K}\right]$

$\dot{m}$ – mass flow rate across the cell $\left[\frac{kg}{s}\right]$

4.4 Sources of uncertainty in modeling assumptions

In a model that attempts to make multiple levels of simplification assumptions there are bound to be numerous sources of uncertainty. Some of these are more related to the data gathering during the experiments, but all of these sources of uncertainty cross over into the model and many are primarily issues of uncertainty in the model itself.

4.4.1 Temperature

The thermocouples themselves are sufficiently accurate for the resolution required by this experiment (

Table 3–5) and were calibrated during the experimental phase. The major issue of uncertainty in temperature measurements relates to locations in the test chamber where there are high temperature gradients. In these locations, small discrepancies in recording the location of the thermocouples could potentially make significant differences in the measured temperature.

This is primarily an issue near the rack inlet and exhaust. Poles L1, L2, R1 and R2 have fairly gradual and linear temperature slopes. Figure 4–52 (a) shows a typical example of one of these poles. One can see that a linear regression has a very high R squared value and a very low slope for all points. If a thermocouple was off by, two to three centimeters, it would only result in an error of approximately 0.1°C. Therefore temperature uncertainty from experimentation is fairly low for these poles.

On the other hand, Figure 4–52 (b) shows a fairly typical profile for the rack exhaust from one of the experiments. Here an error of three centimeters would result in an error of approximately 0.36°C (0.6 °F). This is only about 4% of the typical ΔT across the rack (approximately 10°C / 18 °F) and is a relatively small, although not entirely insignificant, source of uncertainty.

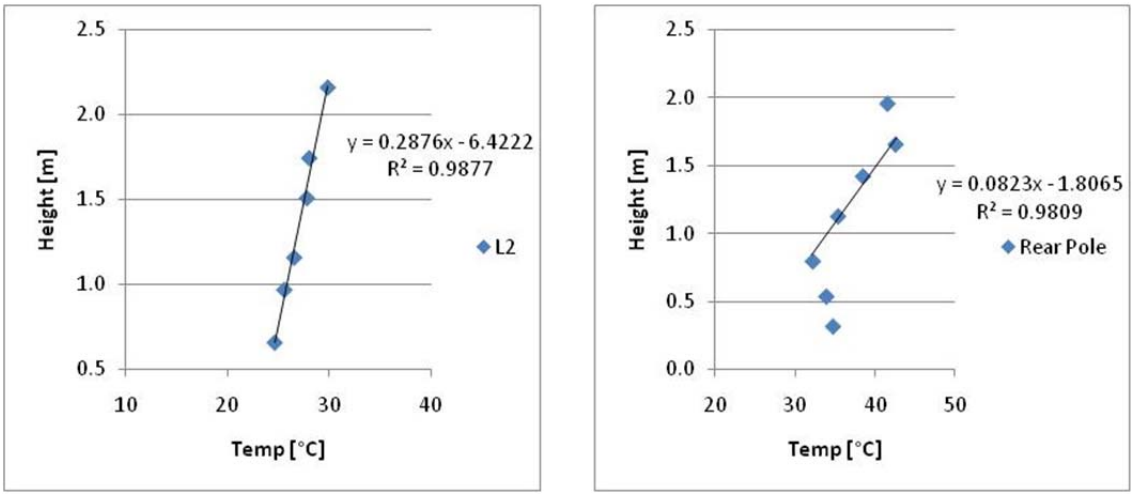

(a) Typical example of non-rack inlet / exhaust thermocouple pole

(b) Rack exhaust profile for Experiment 3

Figure 4–52: Temperature gradients for poles near and away from rack inlet / exhaust

4.4.2 Porous media

As mentioned in Section 4.2.2, it was found that reasonable ranges of porosity did not significantly affect temperature and velocity distributions across the room. This conclusion was arrived at by examining a range of rack inlet and exhaust porosities and comparing the results for velocity and temperature. Rack porosities of 20, 40, 60 and 80% open area were modeled and compared against the results from Experiment 5 (10kW Rack with a high fan speed). The velocity comparisons can be seen in Figure 4–53. Here one can see that the porosities farthest from realistic IT industry values (20 and 40%) produce the most divergent results. There is also very little difference between 60 and 80% porosities. The conclusions translate very well into the temperature comparisons that can be seen in Figure 4–54.

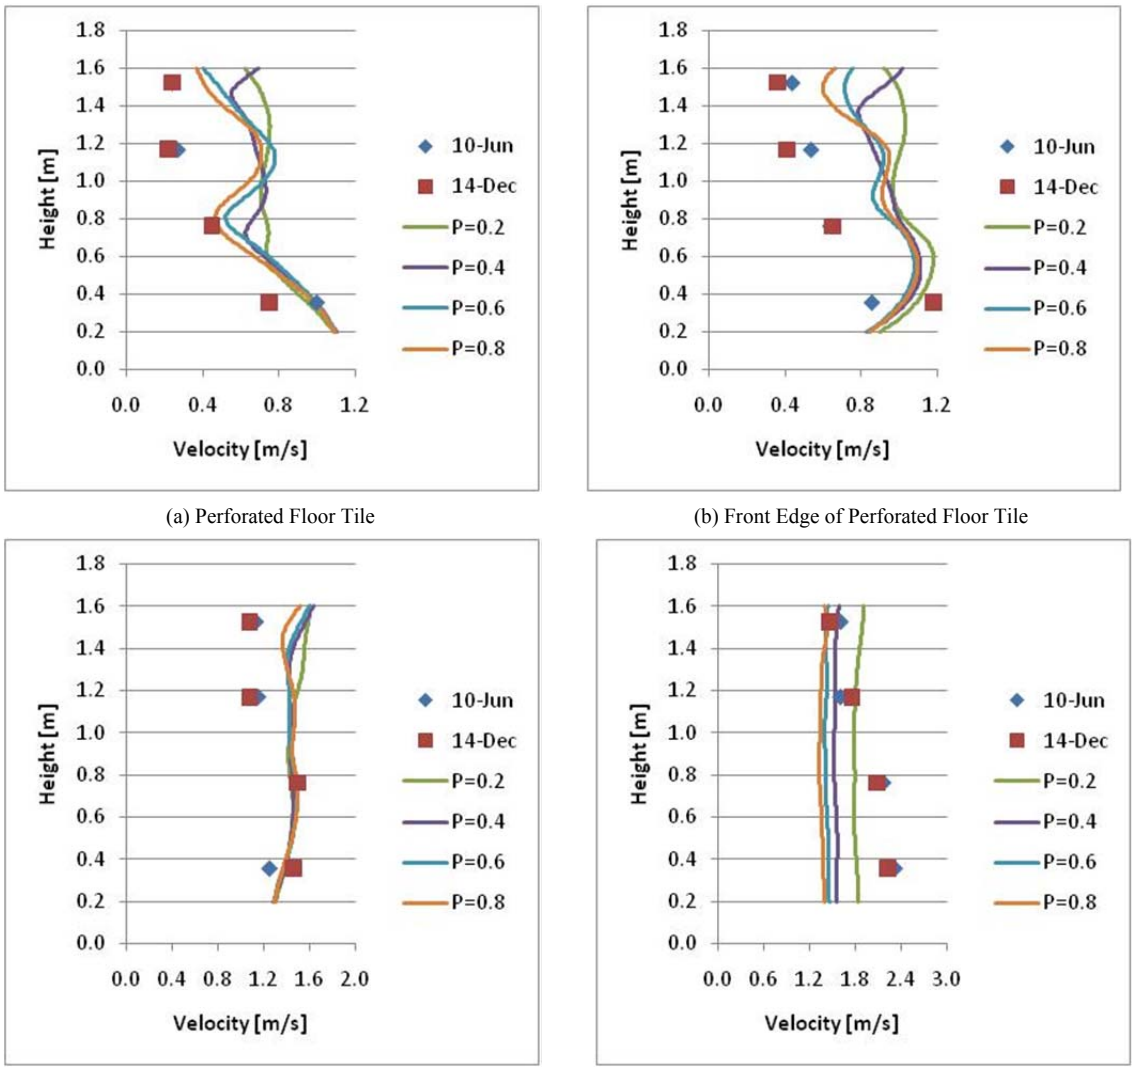

(a) Perforated Floor Tile

(b) Front Edge of Perforated Floor Tile

(c) Rack Inlet

(d) Back Average

Figure 4–53: Effects of Rack Porosity on Velocity Predictions for Experiment 5

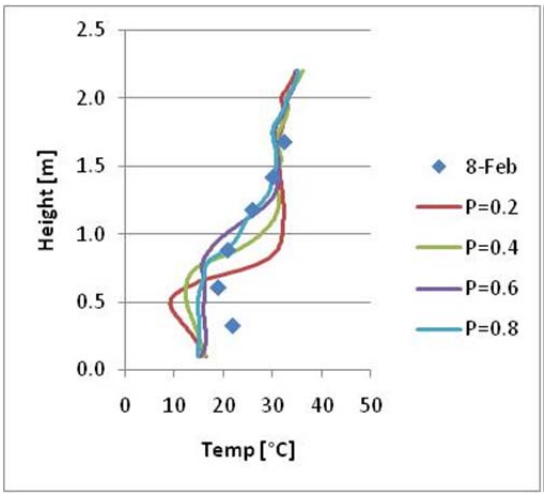

(a) Front Average

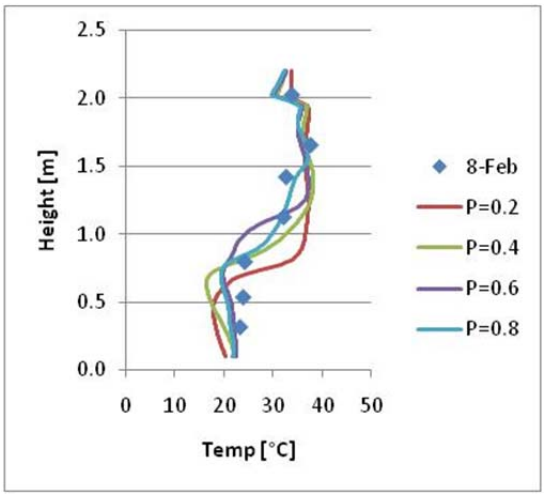

(b) Back Average

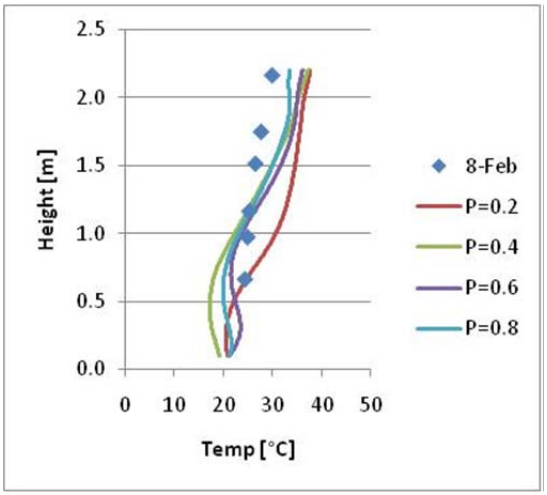

(c) Left 1

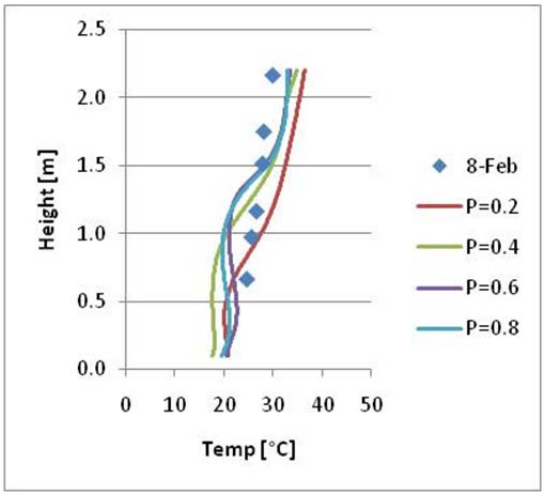

(d) Left 2

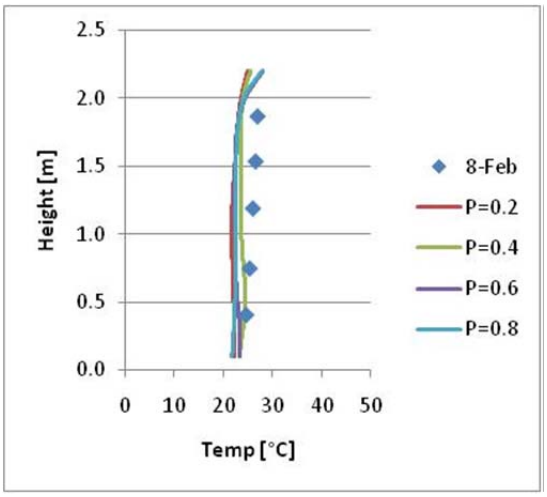

(e) Right 1

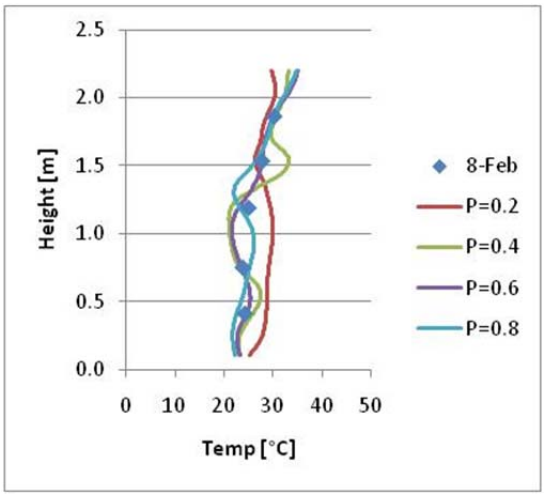

(f) Right 2

Figure 4–54: Effects of Rack Porosity on Temperature for Experiment 5

4.4.3 Fan Speed

Figure 3–37 shows that for both the manufacturer's data and experimental results, the fans follow the relationship predicted by the fan affinity laws. However, there is not precise agreement between these two curves. In particular, the two curves (manufacturer's and experimental) diverge by up to 0.2 m/s (39 fpm) at a fan speed setting of 10.

To address fan speed uncertainty and quantify its effects on the models, four experiments were chosen from across the range of fan speed settings. For lower fan speeds (Experiments 1 and 8) the effects of underestimating the fan speed by 10% were considered. The regular models for Experiments 1 and 8 used fan speeds of 0.56 (110 fpm) and 0.70 m/s (138 fpm) (respectively) based on the data used to make Figure 3–37.

For higher fan speeds, the models tested the effects of fan speeds that could be as high as the manufacturer's suggested fan speeds (up to 14% higher than the experimentally validated fan speeds). For this scenario, the OBM and BBM models for Experiments 3 and 5 were considered. Experiment 3 was a rack with a fan speed setting of 8. For the validation models, a fan speed of 1.06 m/s (209 fpm) was used. The manufacturer's literature suggested that this could be as high as 1.2 m/s (236 fpm). Experiment 5 used a fan speed setting of 10 which corresponded to a setting of 1.42 m/s (280 fpm) on the regular models and 1.62 m/s (319 fpm) in the models created based on the manufacturer's literature.

Table 4–14 shows the results of this parametric analysis. The last two columns on the right show the average differences in temperature and velocity (from experimental results) across all poles between the regular (validation) models and the alternate fan speeds. A negative number indicates that the regular (validation) model was closer to the experimental results than the alternate fan speed. One can see that in nearly all categories the validation model, with its fan speeds based on experimental measurements, was more accurate. The one case where it was not more accurate, it was the same (average temperature delta for Experiment 8). This validates the method of determining the fan speed used in these experiments.

Table 4–14a: Uncertainty Analysis for fan speeds (SI)

Exp. No.	Description	Validation Fan Speed [m/s]	Alternate Fan Speed [m/s]	Avg. Temp Delta [°C]	Avg. Velocity Delta [m/s]
1	4kW Fan1	0.56	0.5	-0.2	0.01
3	8kW Fan8	1.06	1.2	-0.1	-0.06
5	10kW Fan10	1.42	1.62	-1.2	-0.10
8	4kW Top Servers Blocked	0.7	0.63	0.0	-0.04

Table 4–15a: Uncertainty Analysis for fan speeds (IP)

Exp. No.	Description	Validation Fan Speed [fpm]	Alternate Fan Speed [fpm]	Avg. Temp Delta [°F]	Avg. Velocity Delta [fpm]
1	4kW Fan1	110	98	-0.4	2
3	8kW Fan8	209	236	-0.2	-12
5	10kW Fan10	279	319	-2.2	-20
8	4kW Top Servers Blocked	138	124	0.0	-8

This table also shows that for errors of less than 10%, the errors in temperature and velocity estimates from the models are fairly small. While the experimentally gathered fan speeds were more accurate, using the manufacturer's data does not introduce significant error. With the exception of extreme error (14%) at the highest fan speed setting (Experiment 5), the average variation in temperature and velocity readings were 0.1 °C (0.2 °F) and 0.03 m/s (6 fpm).

4.4.4 Gap between rack and the floor

In section 4.2.1 it was mentioned that the models assumed no gap between the rack and floor. This assumption was taken to make the model simpler as well as more general. In order to test the sensitivity of this assumption, models for three different experiments were parametrically tested with a 9 cm (3.5 in) gap added between the floor and the rack (in the model space) to simulate the actual space created by the wheels / supports on the physical rack. There was no significant change in the velocity measurements and only a 0.5 °C (1 °F) change in temperature measurements across all of the poles. Interestingly, the models generally decreased in accuracy when the gap was added. This could be due to a need for increased grid refinement near the gap. However, increasing the number of cells near this gap would increase computational time and unless there can be shown to be an economical return on accuracy for this increased price, the assumption of maintaining a simpler model by eliminating this gap would seem to be justified.

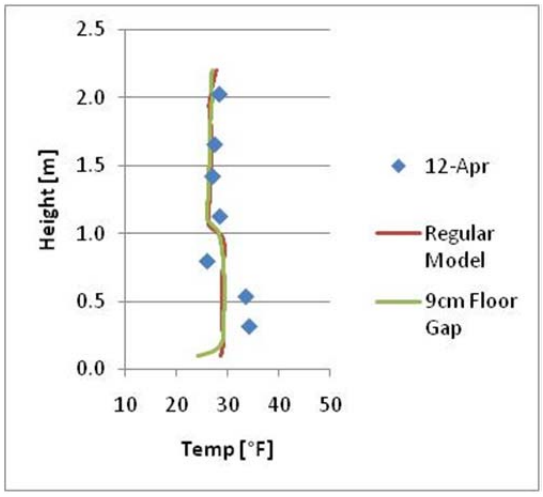

(a) Rack Outlet Temperature Pole, Experiment 8

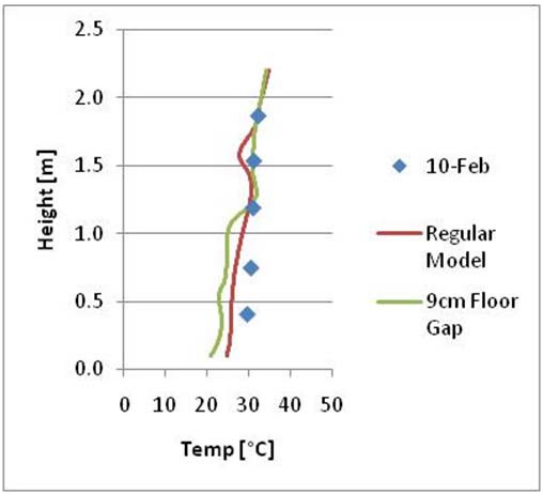

(b) R-2 Temperature Pole, Experiment 3

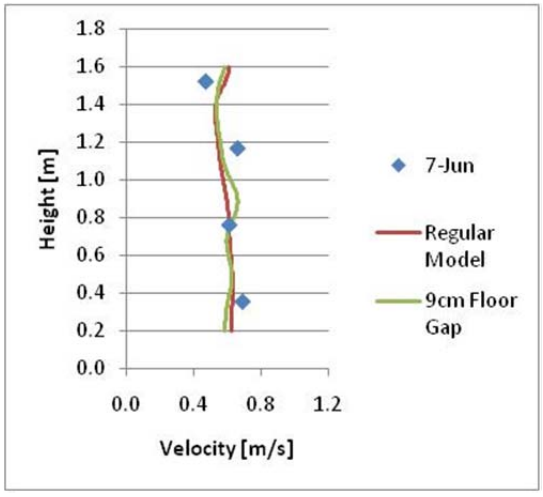

(c) Rack Outlet Velocity Pole, Experiment 1

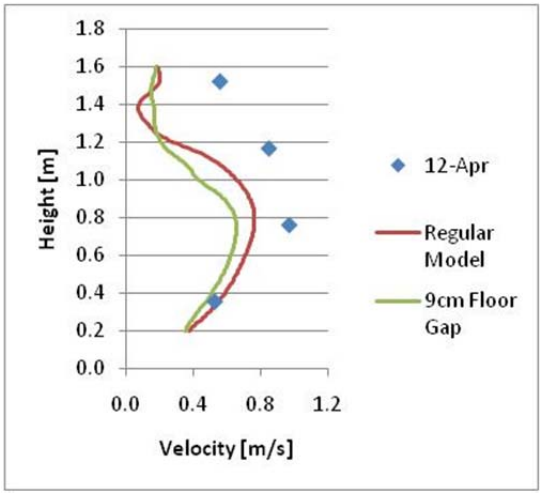

(d) SA-Rack Velocity Pole, Experiment 3

Figure 4–55: Effects of modeling a gap between the rack and the floor

4.4.5 Rack and other Power Sources

While power was measured at the panel for a sample of experiments, it was also deemed necessary to look at the possibility of not accounting for all possible power sources. It is also important to consider the sensitivity of the model to the precision of power used in order to know the effect of situations where the power might not be as precisely known as it was in these experiments. Three models were considered based on Experiments 1, 2 and 5. These experiments were chosen so as to include a high power rack (Experiment 5), a low-powered rack (Experiment 1) and a rack with uneven power distribution (Experiment 2). For all three models, 10% was added to the power put through the racks.

Across all three models, adding 10% power to the rack resulted in an average difference of 0.5 °C (1 °F) (Table 4–16). Interestingly though, for most poles, it tended to slightly increase the accuracy of the predicted results. This indicates that the current models may not be accounting for some minor power sources within the room.

Table 4–16a: Uncertainty Analysis for rack power by experiment (SI)

Exp. No.	Description	Avg. Temp Change [°C]	Avg. Velocity Change [m/s]
1	4kW Fan1	0.4	0.02
2	5.5kW Fan1 Uneven Heat	0.6	0.01
5	10kW Fan10	0.4	0.04

Table 4–17b: Uncertainty Analysis for rack power by experiment (IP)

Exp. No.	Description	Avg. Temp Change [°F]	Avg. Velocity Change [fpm]
1	4kW Fan1	0.7	3.9
2	5.5kW Fan1 Uneven Heat	1.1	2.0
5	10kW Fan10	0.7	7.9

Another interesting aspect is that when one breaks down the changes in temperature by pole, one can see that the poles closest to the rack are less accurate for having more power added.
Table 4–18 shows the change in the error between the model and the experimental results. The error for the model with 10% power added was subtracted from the validation model. Therefore if the model with 10% power was farther from the experimental results when compared with the validation model, then the difference was negative (and vice versa if the model with 10% power added was closer to experimental results).

Table 4–18: Uncertainty Analysis for rack power by temperature pole

Pole	Change in pole Avg.	
	[°C]	[°F]
Front Avg	-0.2	-0.4
Back Avg	-0.3	-0.5
L1	0.0	0.0

Pole	Change in pole Avg.	
	[°C]	[°F]
L2	0.7	1.3
R1	0.2	0.4
R2	0.6	1.1

Table 4–18 also shows that the poles closer to the front of the experimental chamber, poles L2 and R2 (Figure 4–50) showed the greatest improvement in accuracy. This could indicate that any unaccounted for heat sources may be more due to a heat source external to the rack towards the front of the room. Whether one has uncertainty in the rack energy output or possible additional energy sources within the room, if either of these uncertainties is within 10% of the total measured or estimated rack energy, the effect on the average temperature poles across the room still appears be fairly small (within 0.5 °C / 1 °F).

5 RESULTS

5.1 Analysis of Results

For each experiment the OBM and the BBM were compared to the temperature and velocity pole measurements. The deviation from the experimental results was measured for each point along the pole. For experiments with more than one set of temperature or velocity data points, the average of the two sets was used.

5.2 Temperature and Velocity Agreement

5.2.1 Temperature Agreement

What constitutes good temperature agreement is ultimately subjective, but the results can be quantified by either an absolute comparison or by a normalized temperature comparison. The first is easier to visualize, the second is better for comparison.

In the data center industry, temperature rise across the rack is a common metric and a good candidate for comparison. All temperature results were normalized against the lowest rack inlet temperature and highest rack exhaust temperature. This does occasionally produce graphs with results that fall outside of the zero to one scale since there can be lower temperatures in the room (i.e., at the perforated floor tile) and higher temperatures in the room (i.e., at the ceiling near the return air). It can also make it more nuanced to compare experiments with racks that have the same load but different supply air temperatures.

Across all experiments, the OBM and BBM were within 16% of the normalized results for temperatures. In the absolute sense, temperatures in the open box models and black box models were within 2.9 °C (5.2 °F) of experimental results, on average.

5.2.2 ΔT Across the Rack

One curious consequence of using the average of three poles at the rack inlet and rack exhaust is that there is not an even shift that one might expect in the graphs of the rack inlet and outlet temperatures (Figure 5–56). For example, in the black box model, a specified amount of enthalpy is added evenly across the rack exhaust plate, and therefore one might expect the graph of the rack exhaust temperatures to be the rack inlet temperature profile shifted by an amount equal to the energy added (divided by the mass flow rate and the specific heat).

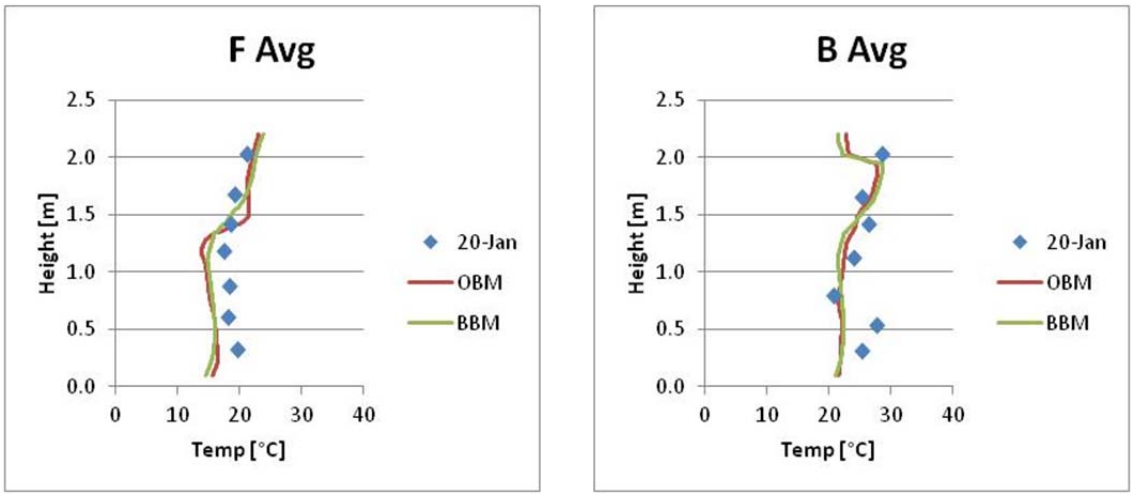

(a) Experiment 1 – 4kW Rack

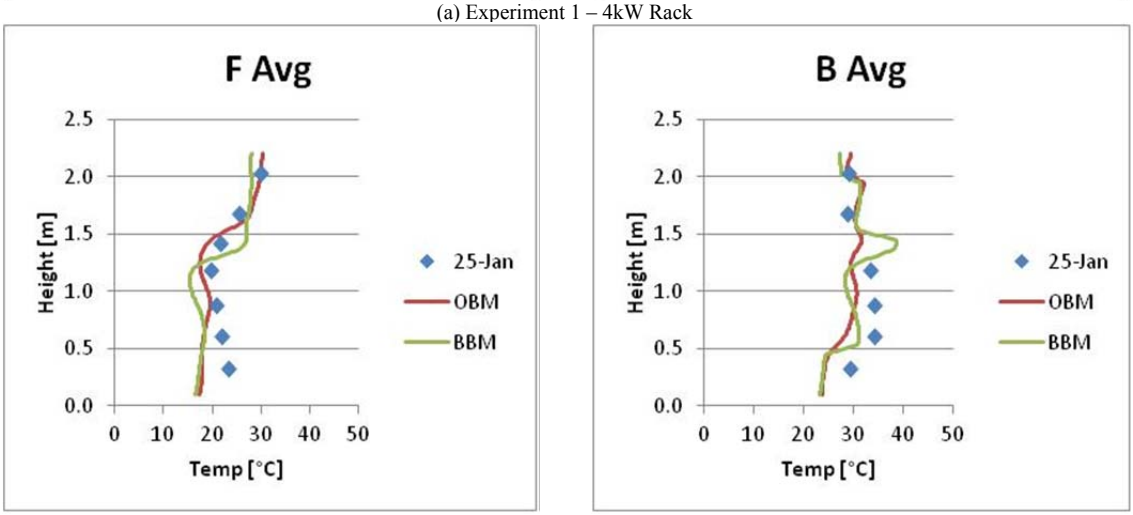

(b) Experiment 2 – 5.5kW Rack (Uneven server simulator distribution)

Figure 5–56: Temperature Shifts From Rack Inlet to Rack Outlet

The reason that the plots for these experimental and model results do not show an even shift is that they are the average of three poles (Figure 4–50) at the left, center and right of the rack inlet and exhaust. The inlet temperature profiles for each pole are slightly different.

Figure 5–57 shows the temperature difference across the rack broken down by individual thermocouple pole locations for the Black Box Models. For example, the "middle pole" indicates the temperature difference between the pole locations F1 and B1 (Figure 4–50). One can see that for each pair of thermocouple poles, the temperature difference, or the amount of enthalpy added is constant even though the temperature profiles for each pole may be different. Figure 5–57 (b) shows the jumps where the server simulator power changes in Experiment 2 (1 kW for the bottom server simulator, 2 kW for the middle two server simulators and 0.5 kW for the top server simulator), but across each server simulator the lines are constant. There are still minor fluctuations in the lines (they are not perfectly straight). This is due to the fact temperature values in Phoenix are stored in the center of the cell and not directly on the plate of the rack. Values for pole positions are extrapolated between the values stored in the center of these cells and so non-linear

jumps or high temperature gradients can cause minor fluctuations in what should be ideally straight lines. This program issue aside, one can see that the temperature shifts are essentially consistent and agree with the amount of heat added by the server simulators even if the overall average does not produce an even shift.

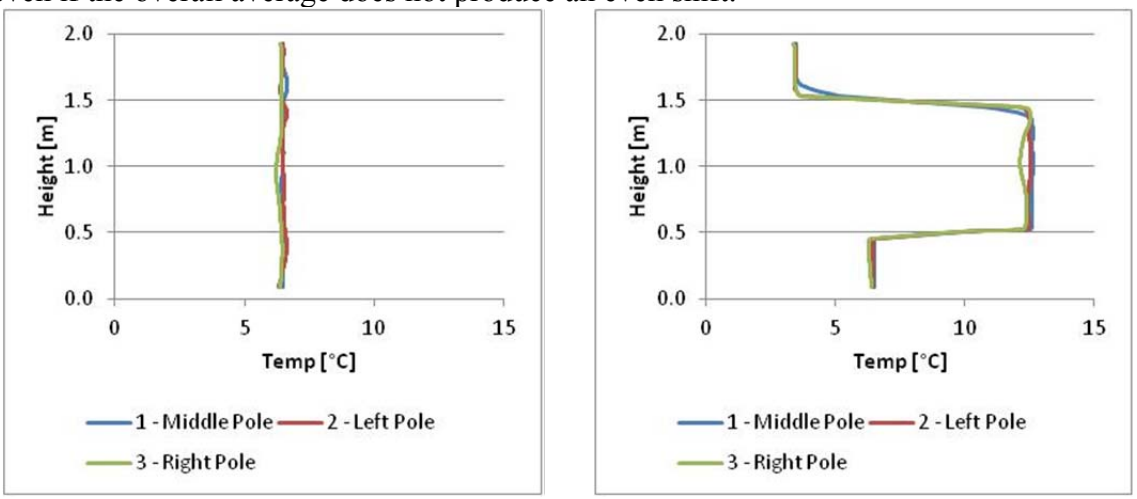

(a) Experiment 1 ΔT Across the Rack (b) Experiment 2 ΔT Across the Rack

Figure 5–57: ΔT Across Rack Inlet and Exhaust By Thermocouple Poles

This may also raise the question of the value of using an average value for the rack inlet and exhaust temperatures. Using an average rack exhaust value was done due to the fact that there is uneven heat distribution in the racks due to how the heating coils are set up, as explained in *Section 3.1.3.3 Heating Elements*. Displaying average temperature profiles makes the results more succinct.

5.2.3 Velocity Agreement

Across all of the models, values were within 0.2 m/s (39 fpm) for velocities. Approximately 82% of all temperature poles (in other words, the average error over the entire pole) and 75% of all data points were within 4°C (7.2 °F). Over 90% of all velocity poles and 82% of all anemometer readings were within 0.3 m/s (59 fpm).

Sections 5.3 – 5.12 discuss the results by experiment.

5.3 Experiment 1 - 4 kW Even Case

The 4kW Rack had an even power distribution and the lowest fan speed setting (0.56 m/s / 110 fpm). As Table 5–19 shows, all poles were within 3°C (5.4 °F) on average. Figure 5–58 shows the normalized temperature results. Absolute temperature results are shown in Appendix A.

Table 5–19: Average Error for Experiment 1 Temperature Poles

	OBM		BBM	
	[°C]	[°F]	[°C]	[°F]
F Avg	2.7	4.9	2.3	4.1
B Avg	2.4	4.3	2.5	4.5
L1	2.5	4.5	2.7	4.9
L2	1.0	1.8	1.4	2.5
R1	1.8	3.2	1.9	3.4
R2	1.5	2.7	2.3	4.1

Velocity measurements were within the acceptable range of 0.3 m/s (59 fpm) for all pole averages (Table 5–20). Figure 5–59 shows the pole measurements. There is some velocity under-prediction for the pole over the center of the perforated floor tile, but this could be due to some uncertainty in determining the throw value. This is a common issue across the experiments which often under-predict near the top. While the throw, terminal velocity and decay constant could be further refined, the values used were considered to be close enough to give good results.

Table 5–20: Average Error for Experiment 1 Velocity Poles

	OBM		BBM	
	[m/s]	**[fpm]**	[m/s]	**[fpm]**
Perforated Floor Tile	0.18	35	0.25	49
Front Edge of PFT	0.21	41	0.26	51
Rack Inlet	0.17	33	0.10	20
Rack Exhaust	0.07	14	0.09	18

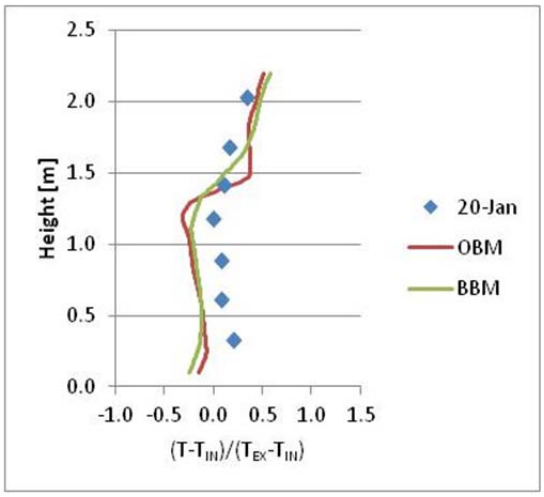

(a) Front Average

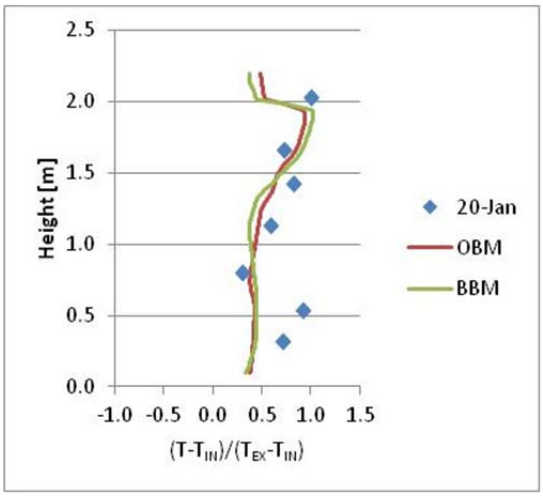

(b) Back Average

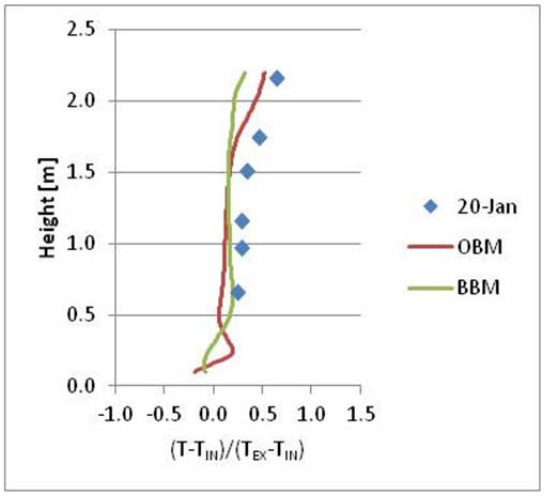

(c) Left 1

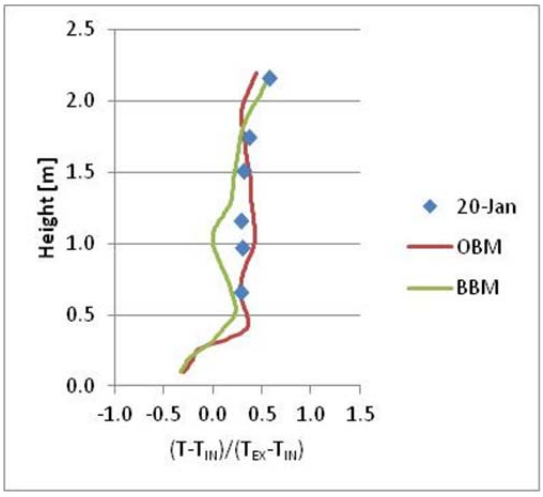

(d) Left 2

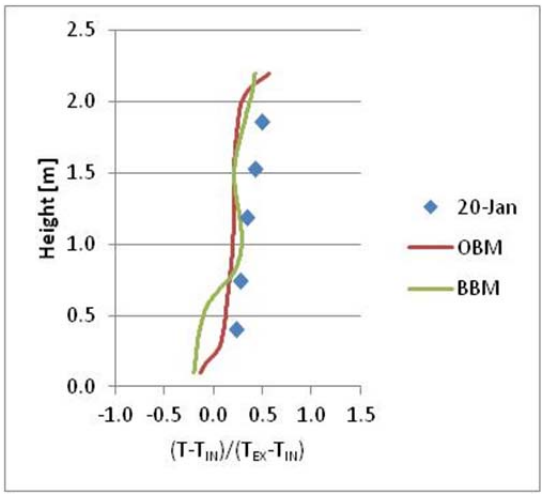

(e) Right 1

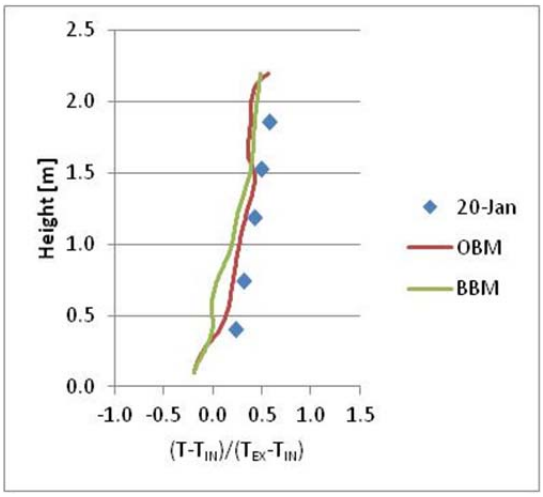

(f) Right 2

Figure 5–58: Normalized Temperature Profiles for Experiment 1

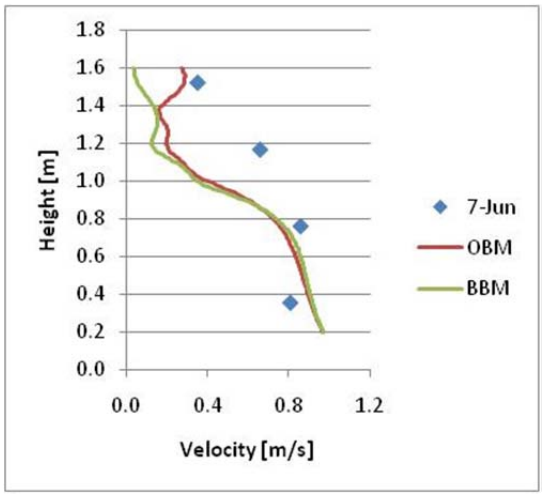

(a) Perforated Floor Tile

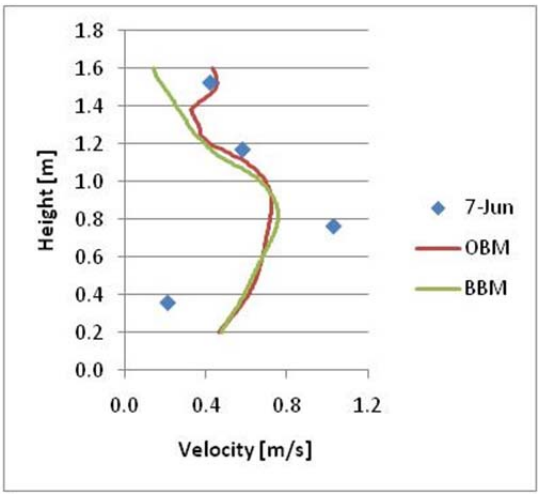

(b) Front Edge of Perforated Floor Tile

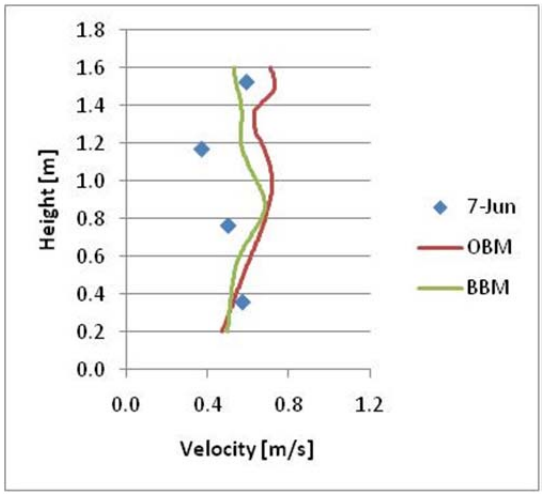

(c) Rack Inlet

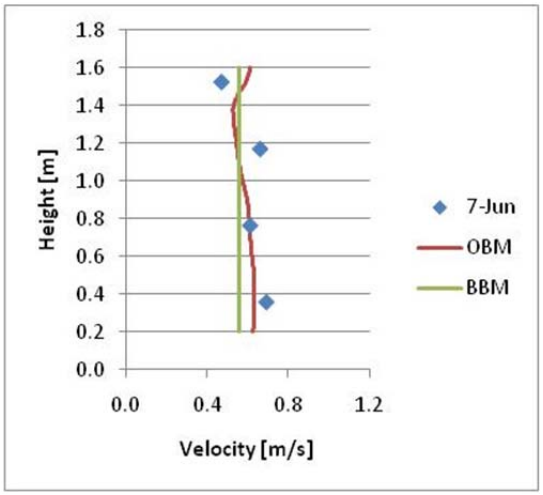

(d) Back Average

Figure 5–59: Velocity Profiles for Experiment 1

5.4 Experiment 2 - 5.5 kW Uneven Case

For the 5.5 kW uneven case, heat distribution from top of the rack to bottom was given as follows, based on (Shrivastava, 2006): 500 W, 2000 W, 2000 W, 1000 W(Figure 3–38). Table 5–21 shows the average error for each of the temperature poles. All of the poles show good agreement except for the Right-1 temperature pole for both models and the Left-2 temperature pole for the OBM. Figure 5–60 shows the normalized temperature profiles and one can see that these are all under-predictions around the mid-height of the room. Some recirculation of the exhaust air is not be captured by the models.

Table 5–21: Average Error for Experiment 2 Temperature Poles

	OBM		BBM	
	[°C]	[°F]	[°C]	[°F]
F Avg	2.8	5.0	3.9	7.0
B Avg	3.3	5.9	3.0	5.4
L1	1.7	3.1	3.1	5.6
L2	4.7	8.5	2.6	4.7
R1	4.6	8.3	4.9	8.8
R2	1.3	2.3	2.8	5.0

Both models show very good agreement with velocity measurements. Table 5–22 shows that all pole averages are within 0.2 m/s (40 fpm) or better. Figure 5–61 shows this velocity agreement broken down by individual anemometer data points. It should be noted that even though the BBM specifies a constant velocity at the exhaust plate (Figure 5–61 (d)), there is some variation seen in this and other experiments. This is due to the fact the pole in the CFD model is four cells away from the rack exhaust and so there is some variation that happens in the space of those four cells.

Table 5–22: Average Error for Experiment 2 Velocity Poles

	OBM		BBM	
	[m/s]	[fpm]	[m/s]	[fpm]
Perforated Floor Tile	0.15	30	0.07	14
Front Edge of PFT	0.18	35	0.17	33
Rack Inlet	0.14	28	0.13	26
Rack Exhaust	0.10	20	0.13	26

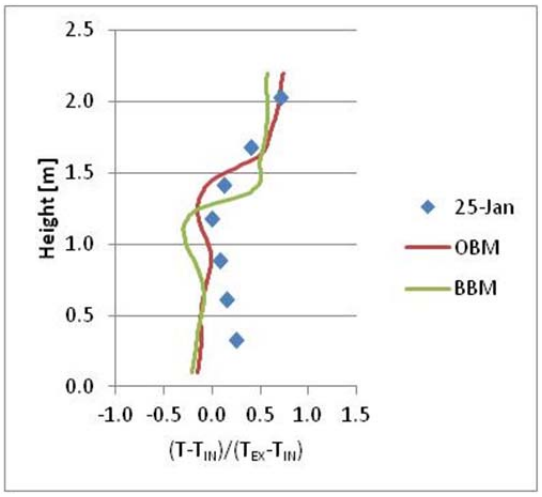

(a) Front Average

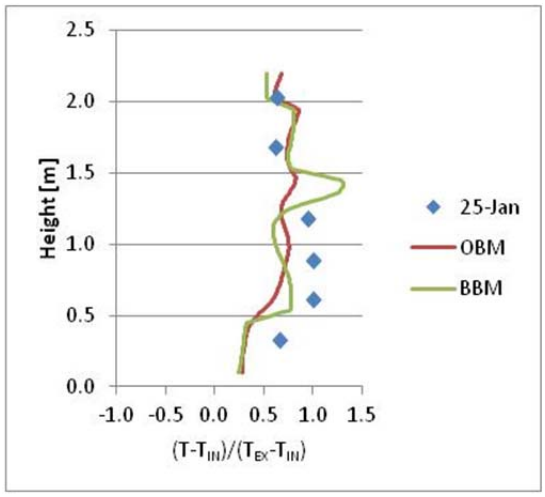

(b) Back Average

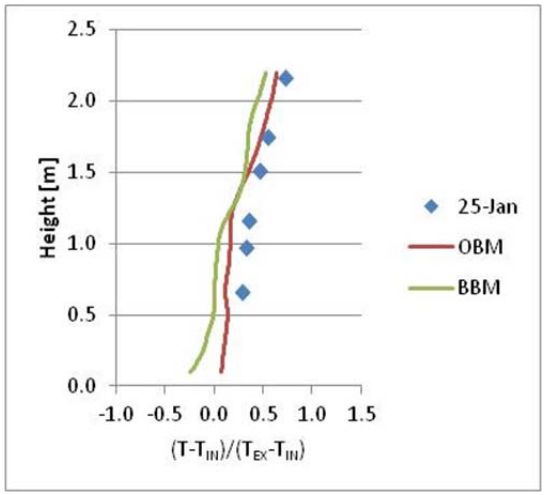

(c) Left 1

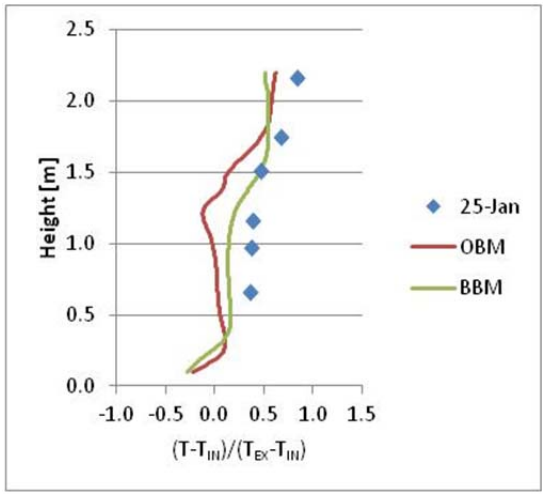

(d) Left 2

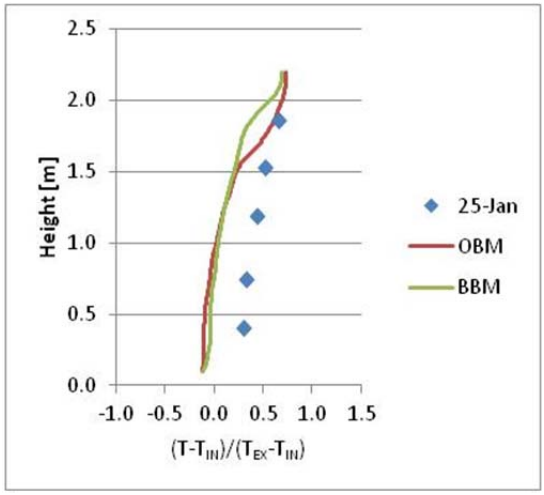

(e) Right 1

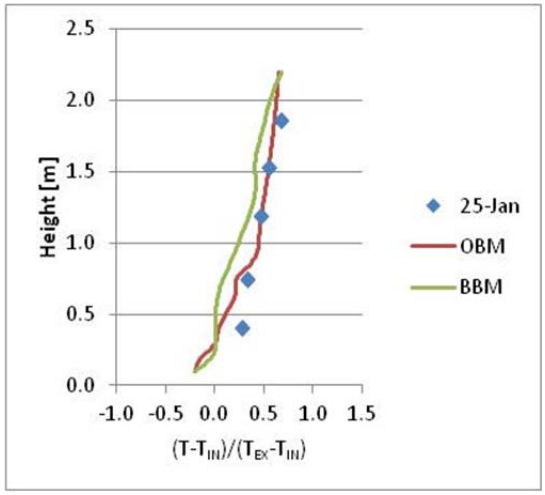

(f) Right 2

Figure 5–60: Normalized Temperature Profiles for Experiment 2

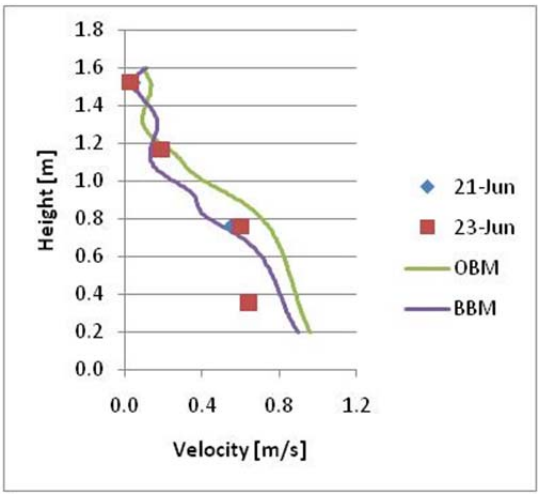

(a) At Perforated Floor Tile

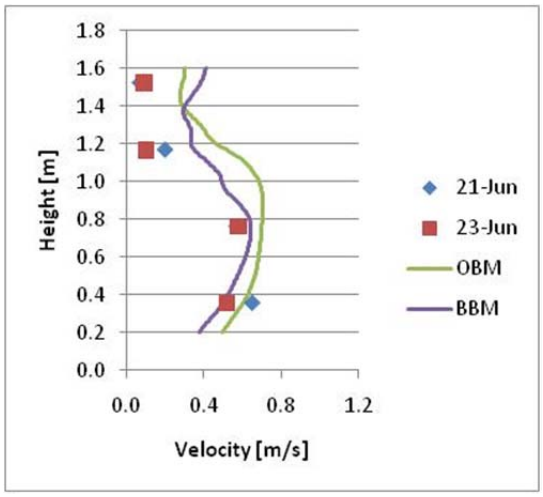

(b) Forward Edge of Perforated Floor Tile

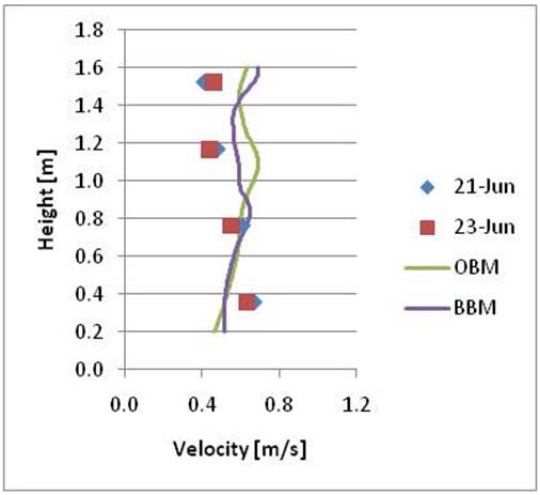

(c) Rack Inlet

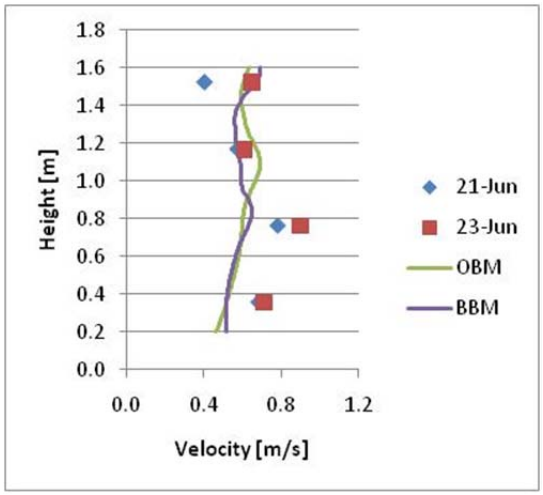

(d) Rack Outlet

Figure 5–61: Velocity Profiles for Experiment 2

5.5 Experiment 3 - 8 kW Case

Experiment 3 is an 8kW rack with a fan speed of 1.06 m/s (209 fpm) or a volumetric flow rate of 1920 cfm (0.91 m^3/s). In spite of a volumetric flow rate through the rack that is around twice the flow through the perforated floor tile, there is fairly good agreement in temperatures (Table 5–23). The pole averages for both models fall within 4°C (7.2 °F) of the experimental results except for the Right-1 pole. Figure 5–62 shows that this is due to a consistent under-prediction. While there is an under-prediction, it follows the trend very closely and appears to be linked to issues that these models have resolving recirculation in the Z-plane at higher rack velocities.

Table 5–23: Average Error for Experiment 3 Temperature Poles

	OBM		BBM	
	[°C]	[°F]	[°C]	[°F]
F Avg	3.0	5.4	3.0	5.4
B Avg	3.3	5.9	3.1	5.6
L1	3.7	6.7	3.6	6.5
L2	3.1	5.6	2.4	4.3
R1	4.5	8.1	4.5	8.1
R2	2.0	3.6	3.5	6.3

Table 5–24 shows very good velocity agreement for this model. The one discrepancy worth noting in this model is at the rack exhaust which can be seen in Figure 5–63 (d). There is an unusually high velocity coming out of the lowest server simulator. This could be due to fact that for this experiment a different set of 2kW heating coils was used which may have caused a hotter airstream to flow over the hotwire anemometer.

Table 5–24: Average Error for Experiment 3 Velocity Poles

	OBM		BBM	
	[m/s]	**[fpm]**	[m/s]	**[fpm]**
Perforated Floor Tile	0.15	30	0.12	24
Front Edge of PFT	0.14	28	0.13	26
Rack Inlet	0.14	28	0.17	33
Rack Exhaust	0.17	33	0.19	37

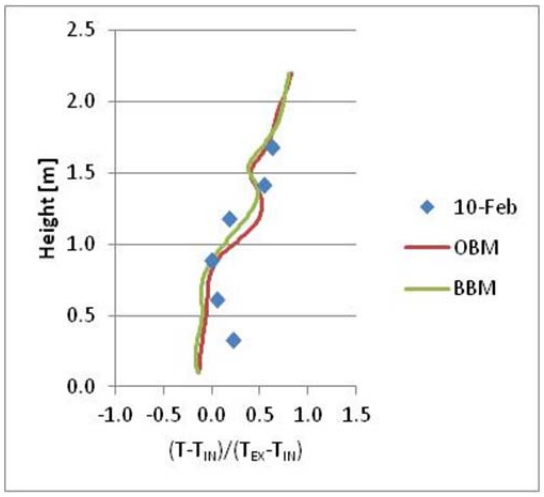

(a) Front Average

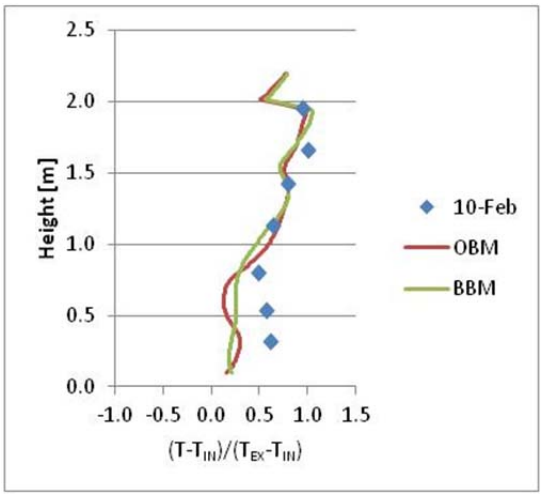

(b) Back Average

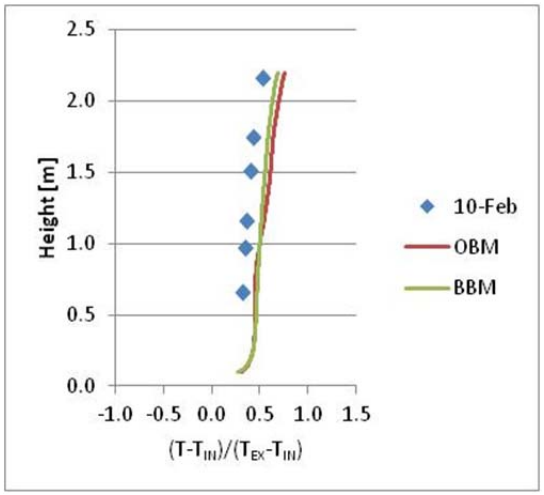

(c) Left 1

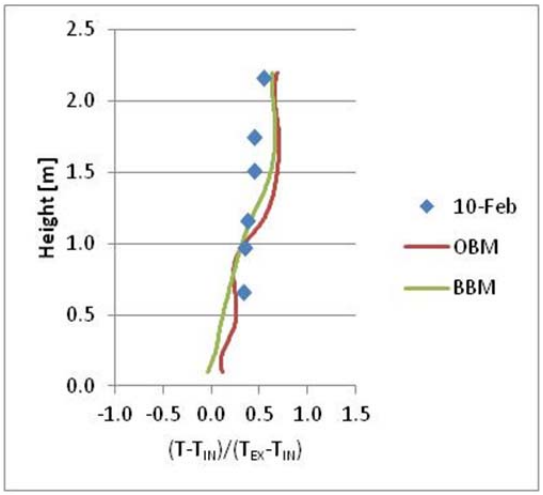

(d) Left 2

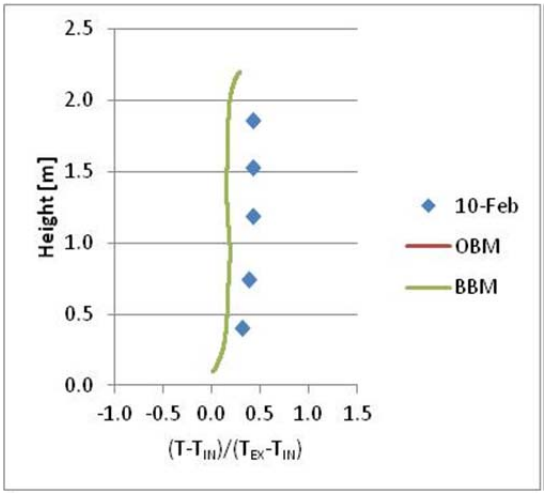

(e) Right 1

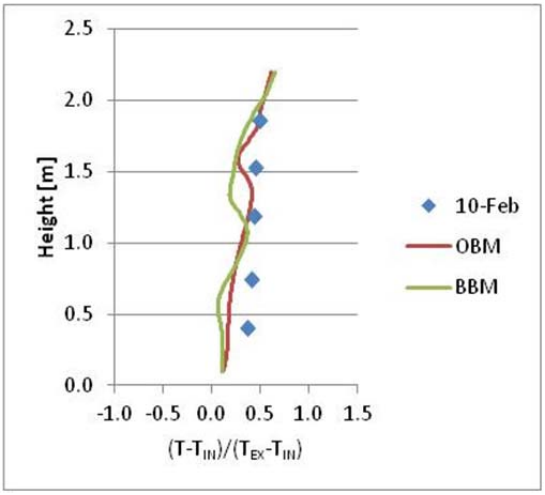

(f) Right 2

Figure 5–62: Normalized Temperature Profiles for Experiment 3

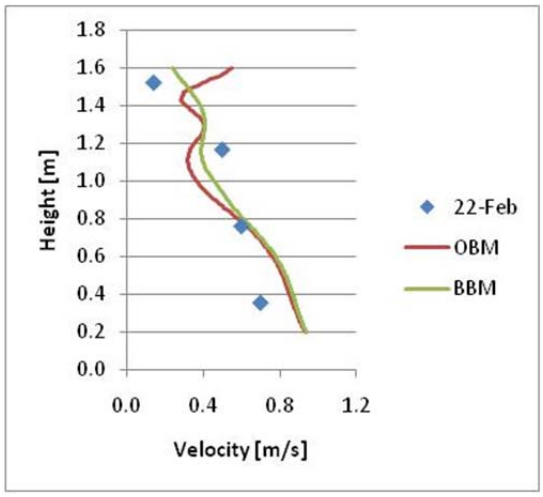

(a) At Perforated Floor Tile

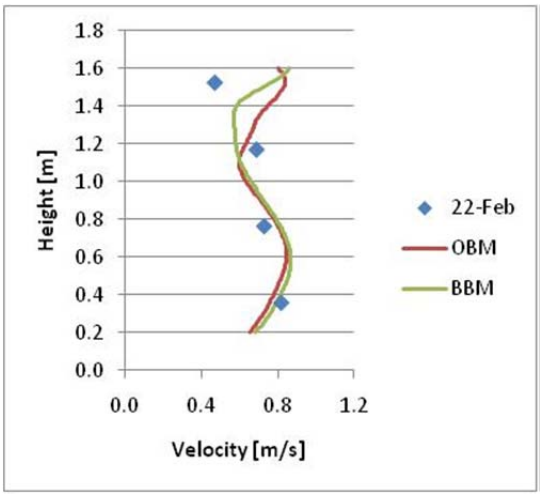

(b) Forward Edge of Perforated Floor Tile

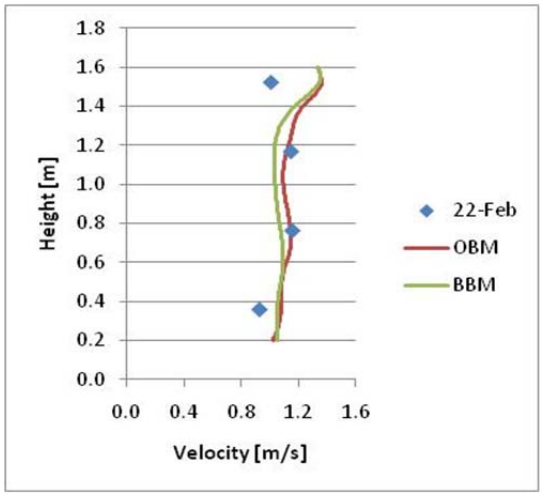

(c) Rack Inlet

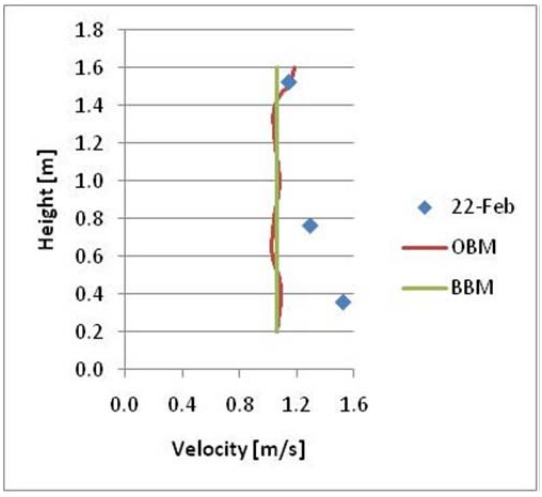

(d) Rack Outlet

Figure 5–63: Velocity Profiles for Experiment 3

5.6 Experiment 4 - 10 kW, Low Fan Speed (0.56 m/s / 110 fpm)

The next three experiments all deal with higher powered racks in the 10kW range. While temperature agreement is good overall, the under-prediction issues towards the back of the room that were observed in experiment 3 show up again for the next two 10kW experiments. Table 5–25 shows that both Left-1 and Right-1 pole results for both models have under-predictions which are over 6 °C (10.8 °F) off (on average) from experimental results. For Left-1 the under-prediction occurs in the lower half of the pole indicating that air recirculation is not being captured very well in this narrower part of the room but that this problem is resolved as one gets above the height of the rack (Figure 5–64 (c)). On Right-1, the temperature profiles follow the trend, but are consistently low indicating issues with recirculation prediction (Figure 5–64 (e)).

Table 5–25: Average Error for Experiment 4 Temperature Poles

	OBM		BBM	
	[°C]	[°F]	[°C]	[°F]
F Avg	4.8	8.6	3.1	5.6
B Avg	2.8	5.0	3.0	5.4
L1	6.7	12.1	7.9	14.2
L2	2.1	3.8	2.7	4.9
R1	6.7	12.1	6.7	12.1
R2	2.3	4.1	2.6	4.7

Problems with temperature prediction toward the back of the room clearly are not due to velocity predictions, as all velocity pole averages are with 0.2 m/s (40 fpm) of experimental data (Table 5–26). Figure 5–65 shows this agreement for experimental data points versus the respective OBM and BBM.

Table 5–26: Average Error for Experiment 4 Velocity Poles

	OBM		BBM	
	[m/s]	[fpm]	[m/s]	[fpm]
Perforated Floor Tile	0.17	33	0.20	39
Front Edge of PFT	0.13	26	0.13	26
Rack Inlet	0.13	26	0.09	18
Rack Exhaust	0.13	26	0.10	20

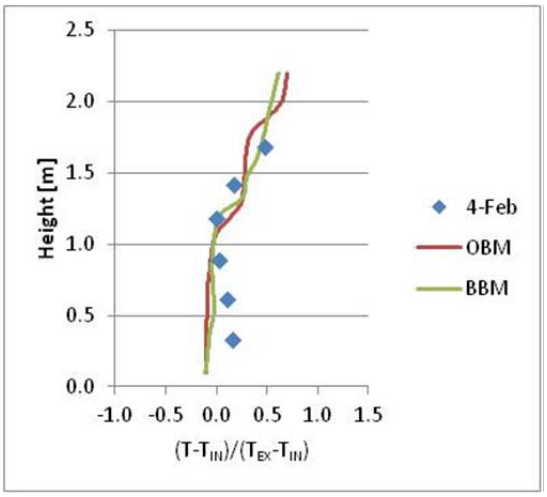

(a) Front Average

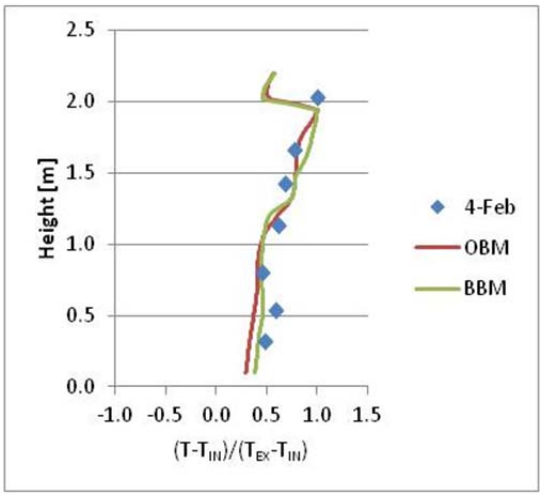

(b) Back Average

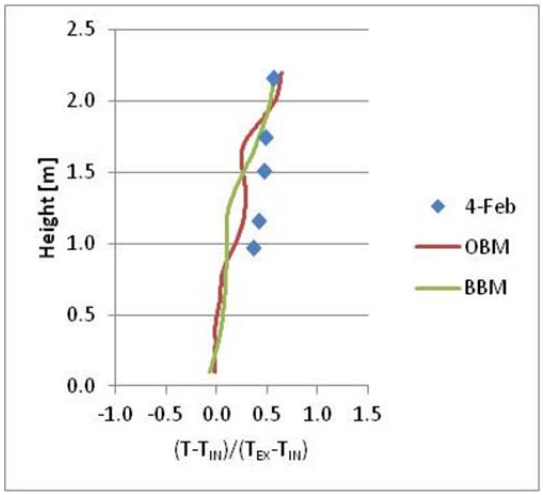

(c) Left 1

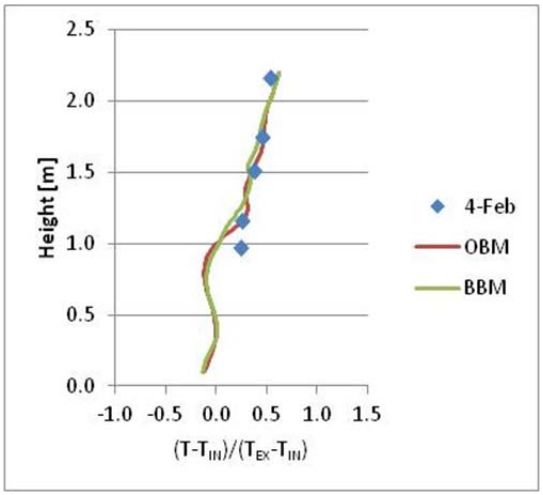

(d) Left 2

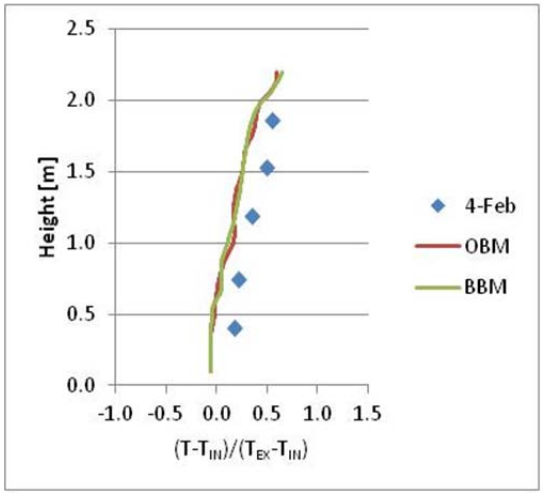

(e) Right 1

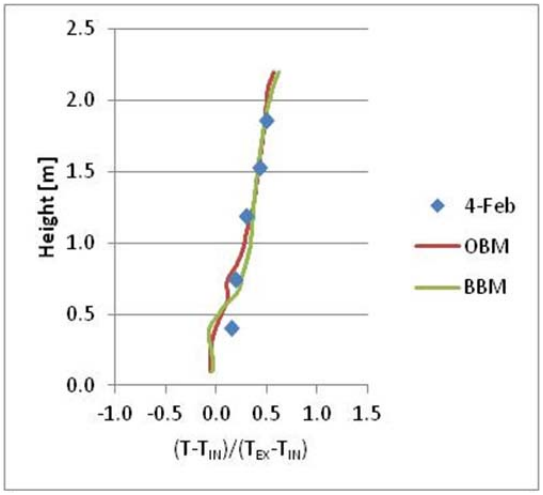

(f) Right 2

Figure 5–64: Normalized Temperature Profiles for Experiment 4

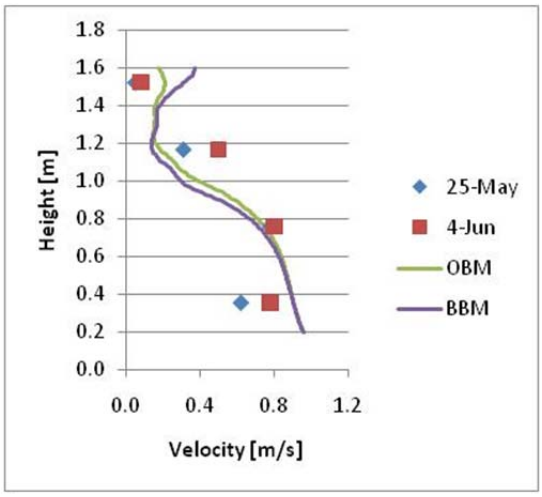

(a) At Perforated Floor Tile

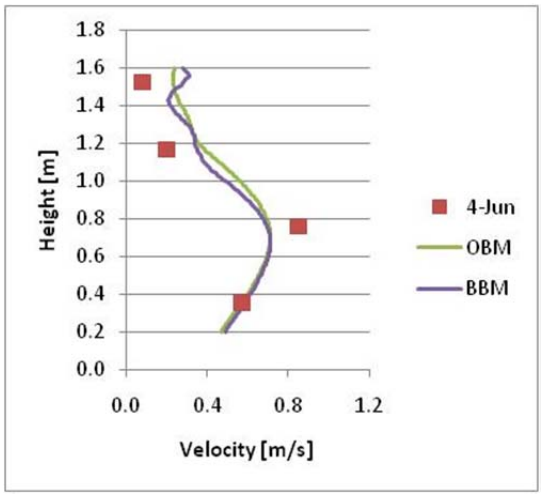

(b) Forward Edge of Perforated Floor Tile

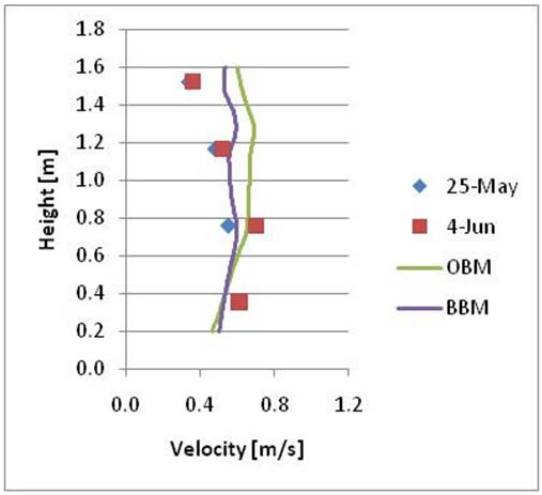

(c) Rack Inlet

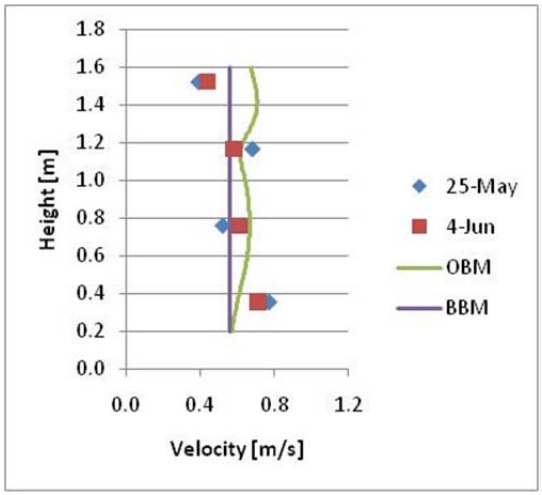

(d) Rack Exhaust

Figure 5–65: Velocity Profiles for Experiment 4

5.7 Experiment 5 - 10 kW, High Fan Speed (1.42 m/s / 280 fpm)

Table 5–27 shows decent agreement across all of the pole averages with the small exception of the front pole average in the BBM model. This is due to an under-prediction near the bottom of the pole (Figure 5–66 (a)). This appears to be tied to an under-prediction of the velocity coming out of the exhaust from the lower server simulator (Figure 5–67 (d)). Since this issue has shown up on several experiments with higher powered racks, there may be some issue with the fan for this server simulator when it is operated at these power settings.

Table 5–27: Average Error for Experiment 5 Temperature Poles

	OBM		BBM	
	[°C]	[°F]	[°C]	[°F]
F Avg	2.7	4.9	4.3	7.7
B Avg	2.1	3.8	2.1	3.8
L1	2.8	5.0	0.8	1.4
L2	3.5	6.3	3.0	5.4
R1	3.6	6.5	2.4	4.3
R2	2.4	4.3	2.9	5.2

Both models demonstrated good agreement in the velocity poles (Table 5–28) with the exception of the exhaust pole. This was largely due to the issues noted in the discussion above.

Table 5–28: Average Error for Experiment 5 Velocity Poles

	OBM		BBM	
	[m/s]	**[fpm]**	[m/s]	**[fpm]**
Perforated Floor Tile	0.14	28	0.25	49
Front Edge of PFT	0.30	59	0.26	51
Rack Inlet	0.19	37	0.24	47
Rack Exhaust	0.48	94	0.49	96

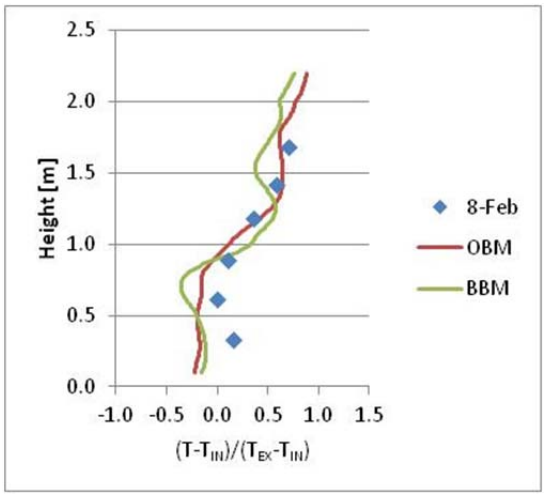

(a) Front Average

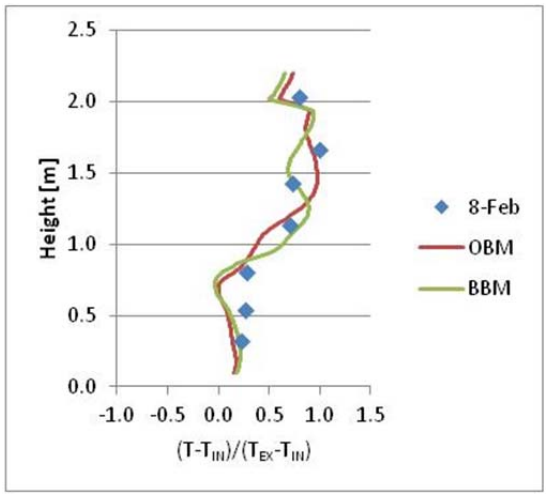

(b) Back Average

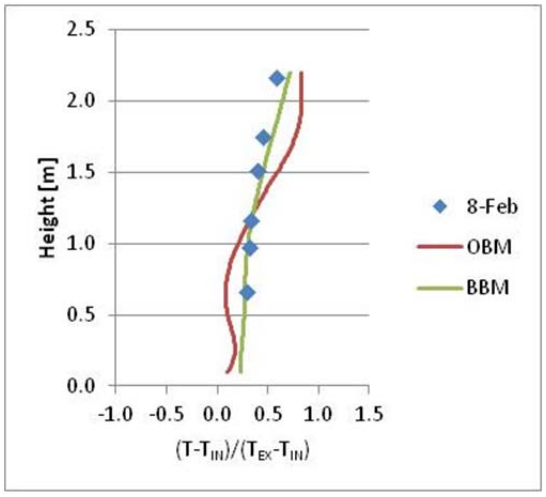

(c) Left 1

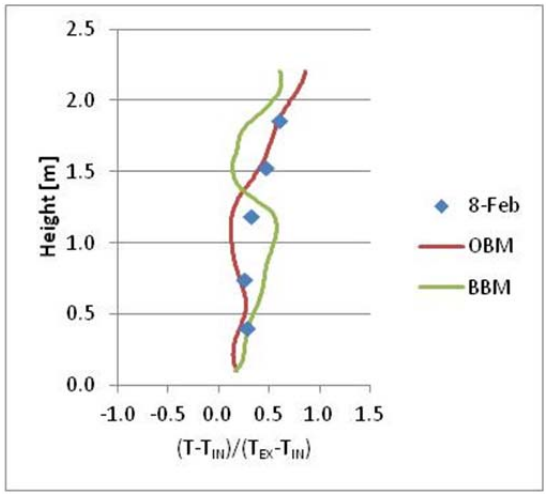

(d) Left 2

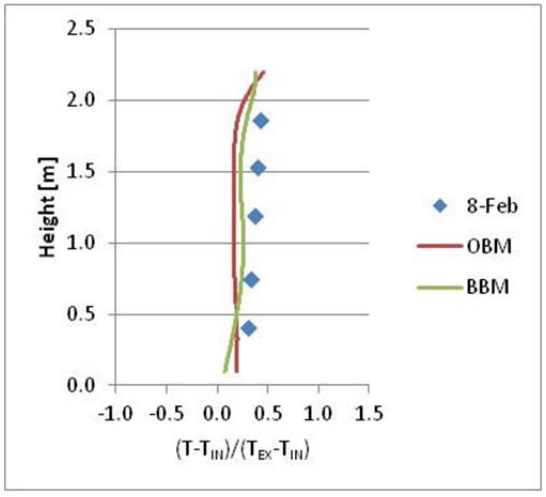

(e) Right 1

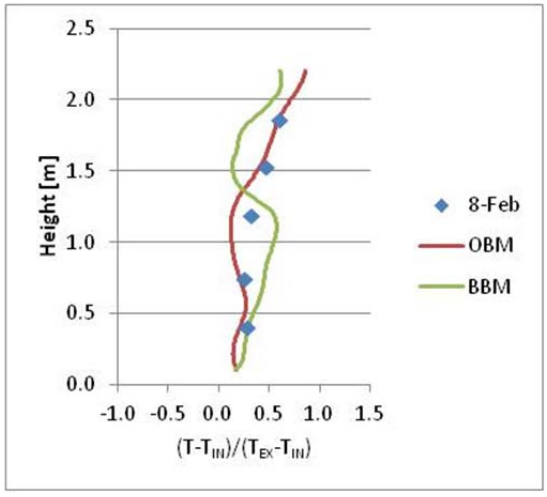

(f) Right 2

Figure 5–66: Normalized Temperature Profiles for Experiment 5

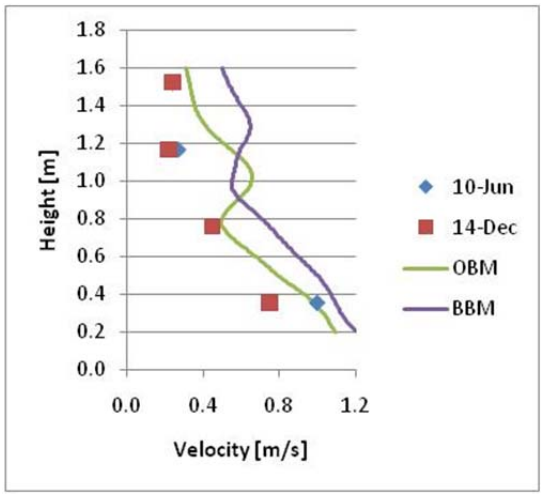

(a) At Perforated Floor Tile

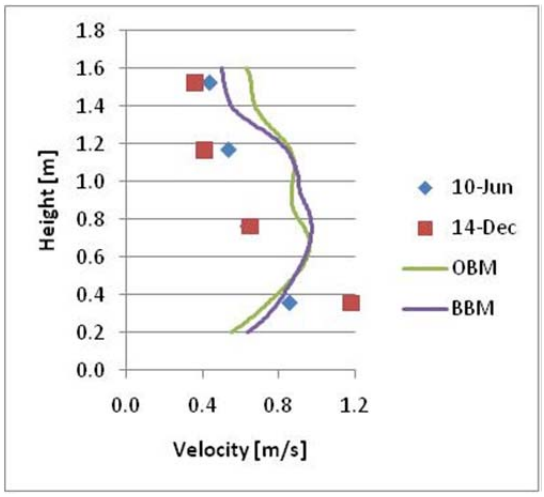

(b) Forward Edge of Perforated Floor Tile

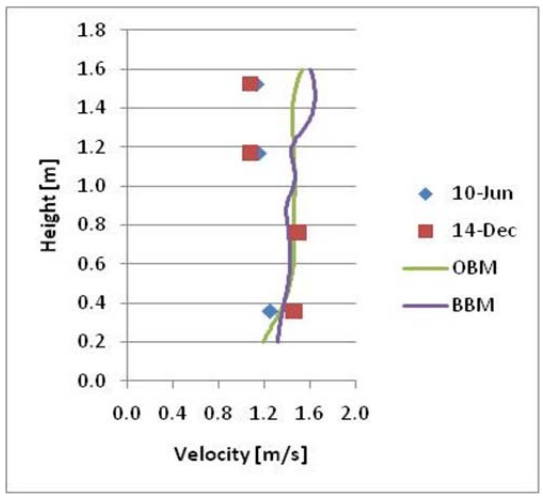

(c) Rack Inlet

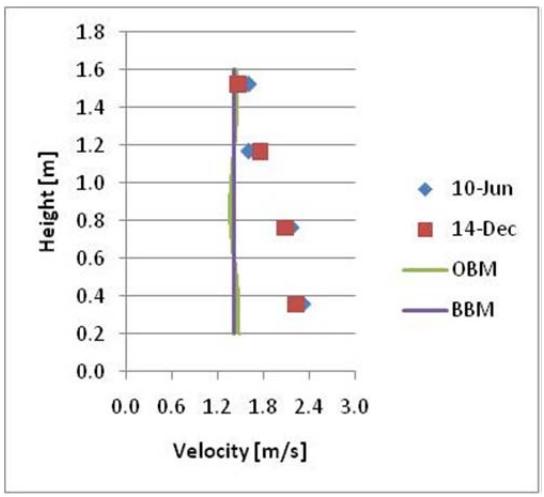

(d) Rack Outlet

Figure 5–67: Velocity Profiles for Experiment 5

5.8 Experiment 6 - 4kW Rack, Higher Supply Air Temperature (22.2 °C / 72.0 °F)

This was the first of five experiments conducted after the initial validation experiments (1-5) that were designed to look at some of the variables that had been kept largely the same for the initial experiments (supply air flow rate, temperature, rack location, etc.). Experiments 7 and 8 are similar to experiment 1 except that the supply conditions are altered. As with Experiment 1, however, the temperature agreement turns out to be very good across all temperature poles (Table 5–29). Figure 5–68 shows the agreement of the models with the normalized thermocouple data.

Table 5–29: Average Error for Experiment 6 Temperature Poles

	OBM		BBM	
	[°C]	[°F]	[°C]	[°F]
F Avg	2.6	4.7	3.1	5.6
B Avg	3.3	5.9	2.7	4.9
L1	1.4	2.5	1.7	3.1
L2	1.4	2.5	1.8	3.2
R1	0.6	1.1	0.4	0.7
R2	1.3	2.3	0.6	1.1

Velocity agreement was also good with the small exception of the pole in the center of the perforated floor tile for the BBM (Table 5–30). The discrepancy for both the OBM and BBM is due to an under-prediction near the top of the pole (Figure 5–69 (a)). This is most likely due to some uncertainty in calculating the throw for the perforated floor tile.

Table 5–30: Average Error for Experiment 6 Velocity Poles

	OBM		BBM	
	[m/s]	[fpm]	[m/s]	[fpm]
Perforated Floor Tile	0.29	57	0.35	69
Front Edge of PFT	0.22	43	0.26	51
Rack Inlet	0.15	30	0.14	28
Rack Exhaust	0.07	14	0.05	10

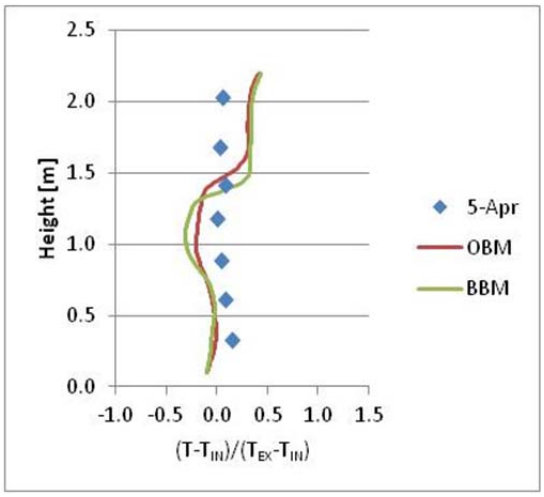

(a) Front Average

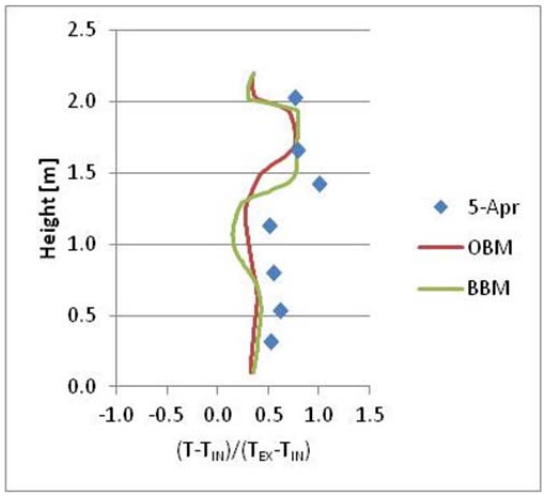

(b) Back Average

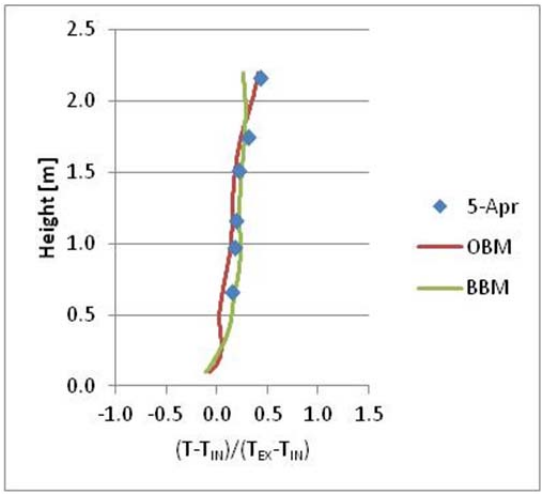

(c) Left 1

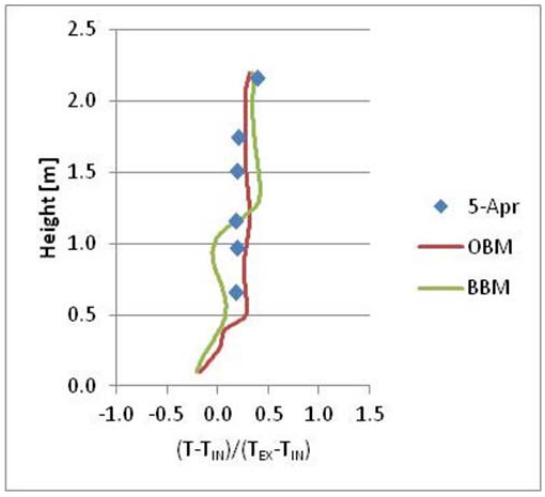

(d) Left 2

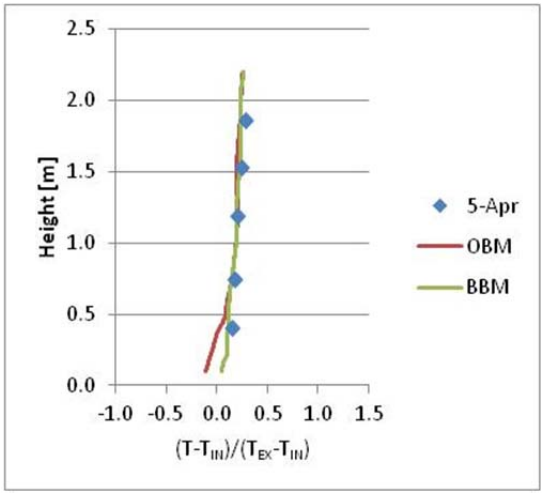

(e) Right 1

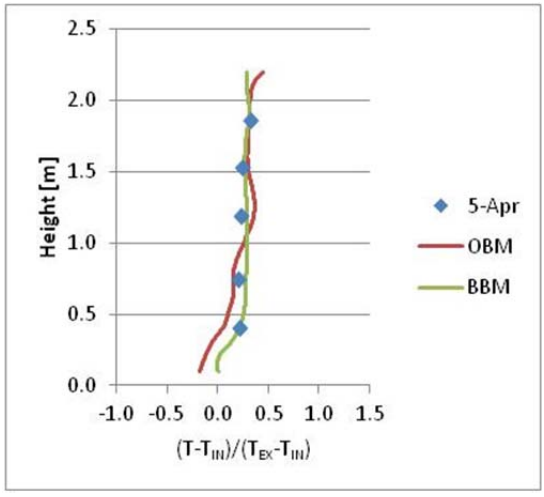

(f) Right 2

Figure 5–68: Normalized Temperature Profiles for Experiment 6

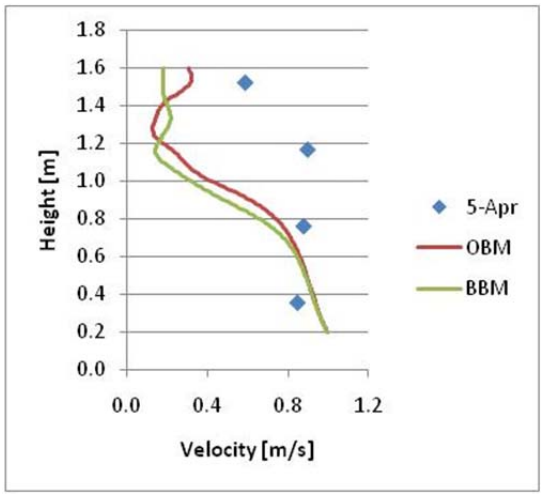

(a) At Perforated Floor Tile

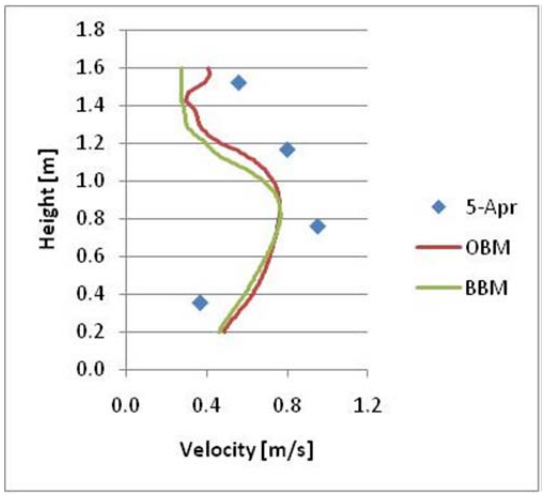

(b) Forward Edge of Perforated Floor Tile

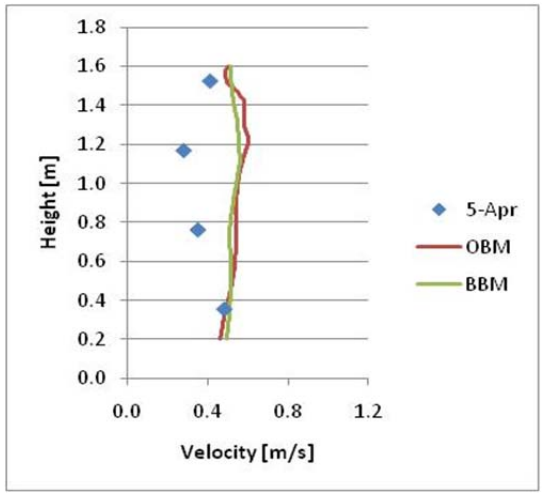

(c) Rack Inlet

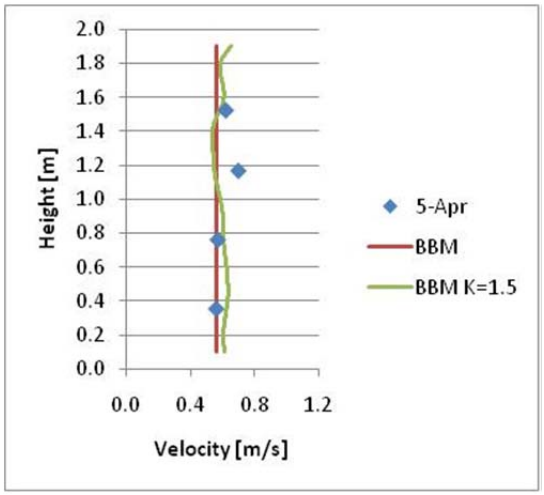

(d) Rack Outlet

Figure 5–69: Velocity Profiles for Experiment 6

5.9 Experiment 7 - 4 kW Rack, Lower Supply Air Flow Rate (460 cfm / 0.22 m^3/s)

Temperature agreements for experiment 7 were good, but not quite as good as experiments 1 and 6. Table 5–31 shows that while majority of the poles were within 3 - 4°C (5.4 – 7.2 °F), there were a few poles that are slightly over 5°C (9 °F) off of experimental results. Similar to the 10 kW cases, these were mostly in the back of the room (Left-1 and Right-1), closer to the rack exhaust. Given the disparity between the perforated floor tile and rack volumetric flow rates, there was most likely some recirculation around the sides which was not fully captured. Figure 5–70 shows that most of this error occurred around the middle height of the poles.

Table 5–31: Average Error for Experiment 7 Temperature Poles

	OBM		BBM	
	[°C]	[°F]	[°C]	[°F]
F Avg	2.5	4.5	2.3	4.1
B Avg	2.4	4.3	2.0	3.6
L1	4.6	8.3	5.2	9.4
L2	3.4	6.1	2.2	4.0
R1	4.3	7.7	5.1	9.2
R2	2.5	4.5	3.3	5.9

Table 5–32 and Figure 5–71 show that lower volumetric airflows did not cause any issues with model agreement.

Table 5–32: Average Error for Experiment 7 Velocity Poles

	OBM		BBM	
	[m/s]	[fpm]	[m/s]	[fpm]
Perforated Floor Tile	0.12	24	0.03	6
Front Edge of PFT	0.17	33	0.19	37
Rack Inlet	0.09	18	0.08	16
Rack Exhaust	0.10	20	0.09	18

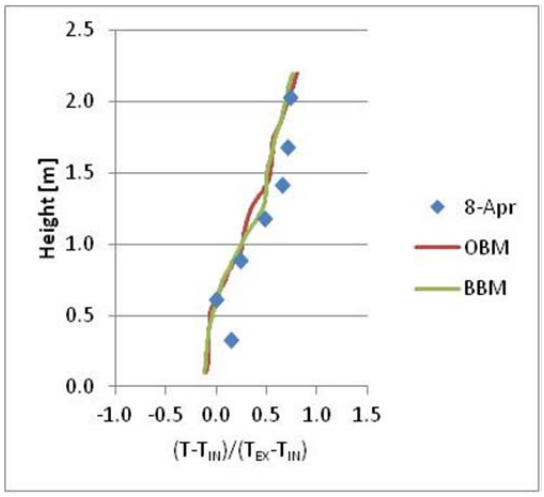

(a) Front Average

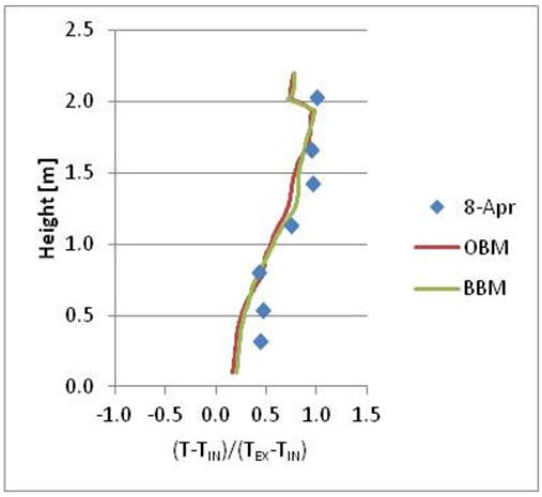

(b) Back Average

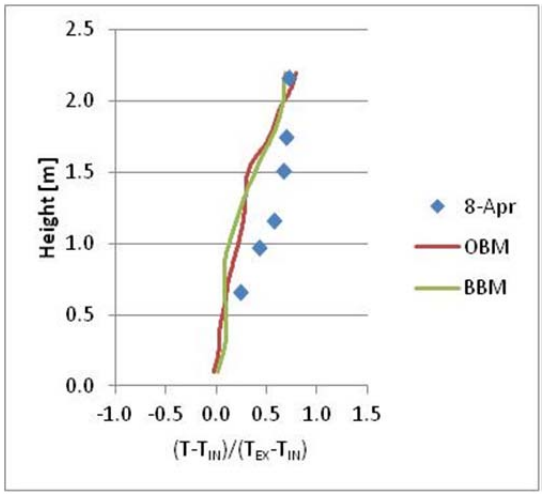

(c) Left 1

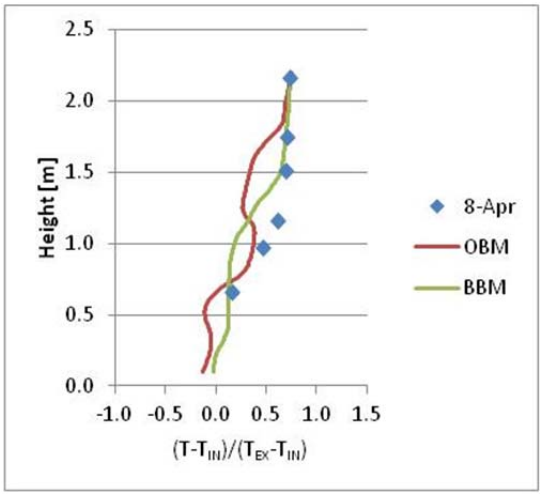

(d) Left 2

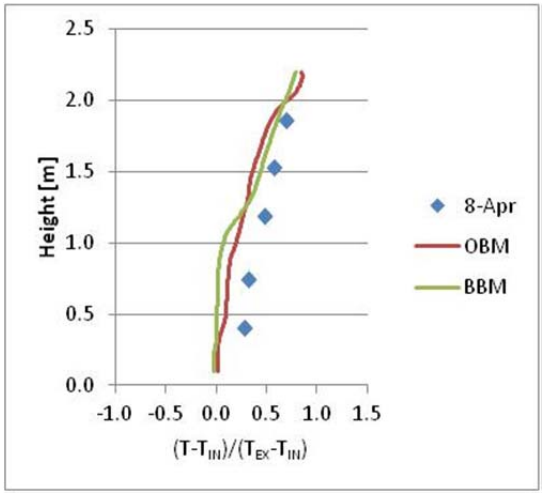

(e) Right 1

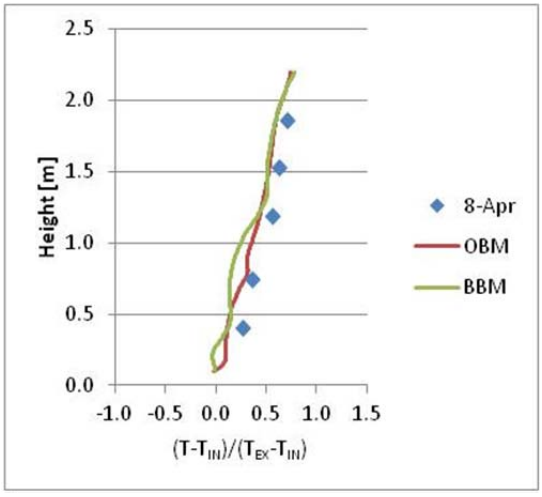

(f) Right 2

Figure 5–70: Normalized Temperature Profiles for Experiment 7

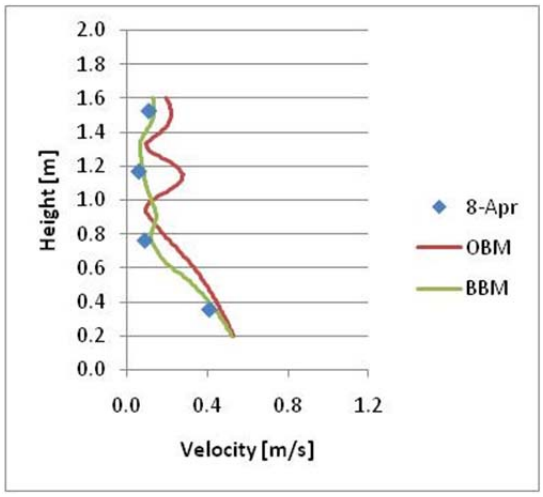

(a) At Perforated Floor Tile

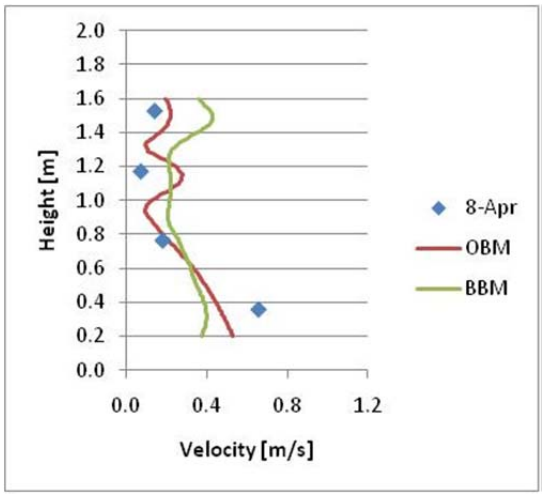

(b) Forward Edge of Perforated Floor Tile

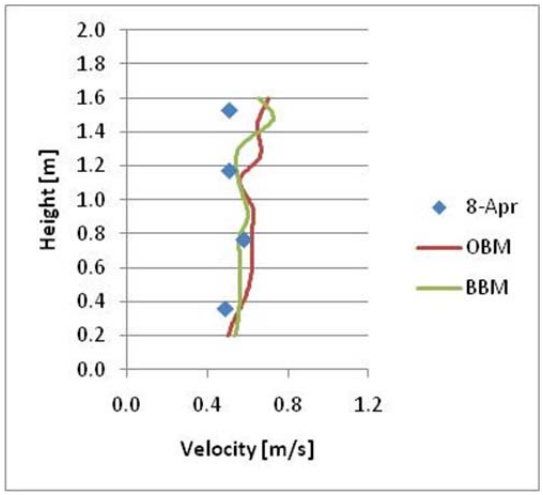

(c) Rack Inlet

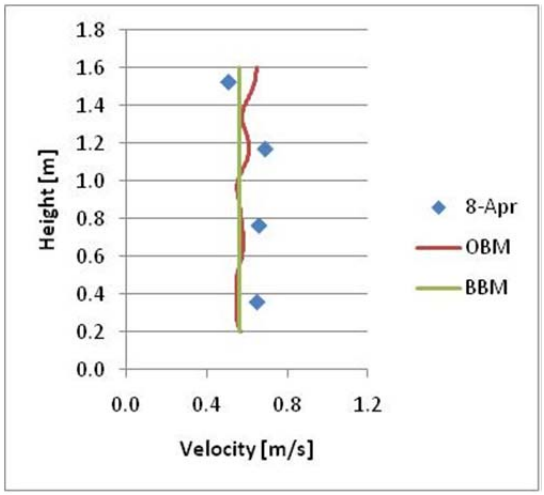

(d) Rack Outlet

Figure 5–71: Velocity Profiles for Experiment 7

5.10 Experiment 8 - 4 kW, Top Servers Off and Blocked

Experiment 8 and 9 look at how well the model handles racks that are not fully populated by servers. Sometimes these racks are left open, other times there is a blanking panel placed in the empty space. The intent of the blanking panel is to prevent hot air from the rack exhaust from recirculating back to the front of the rack.

The model handled the blanking panel well. The greatest discrepancies were the front average and the Left-2 poles in the BBM (Table 5–33). The front poles average has higher over-predictions near the top of the pole (Figure 5–72 (a)). It is possible that without the effects of the fans in the upper levels of the rack that the model is under-predicting mixing near the front of the room due to recirculation outside of the rack. It is also possible that both models under-predict the amount of air from the upper part of the space that is induced in from the lower server simulators.

Table 5–33: Average Error for Experiment 8 Temperature Poles

	OBM		BBM	
	[°C]	[°F]	[°C]	[°F]
F Avg	3.7	6.7	4.1	7.4
B Avg	2.6	4.7	2.1	3.8
L1	1.5	2.7	0.3	0.5
L2	3.1	5.6	4.9	8.8
R1	1.4	2.5	1.8	3.2
R2	2.2	4.0	2.1	3.8

From Table 5–34, one can see that the average velocity predictions are all within accepted boundaries with the slight exception of the perforated floor tile. One can see from Figure 5–73 (a) that there is an under-prediction near the top of the pole. This is due to uncertainty in calculating the K value from the data available.

Table 5–34: Average Error for Experiment 8 Velocity Poles

	OBM		BBM	
	[m/s]	[fpm]	[m/s]	[fpm]
Perforated Floor Tile	0.31	61	0.32	63
Front Edge of PFT	0.27	53	0.29	57
Rack Inlet	0.05	10	0.04	8
Rack Exhaust	0.14	28	0.16	31

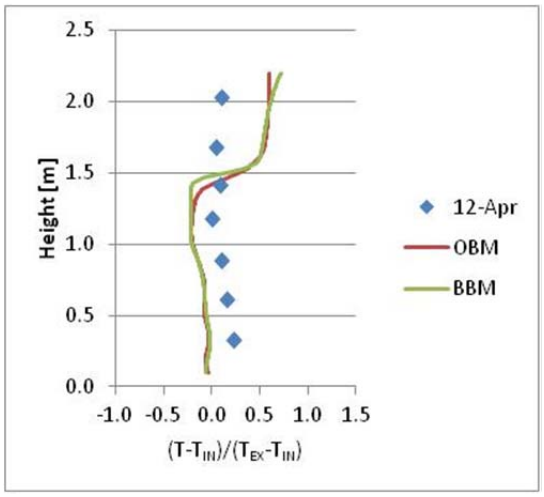

(a) Front Average

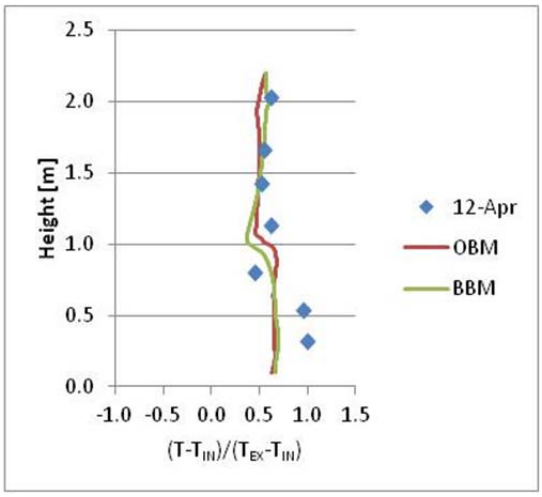

(b) Back Average

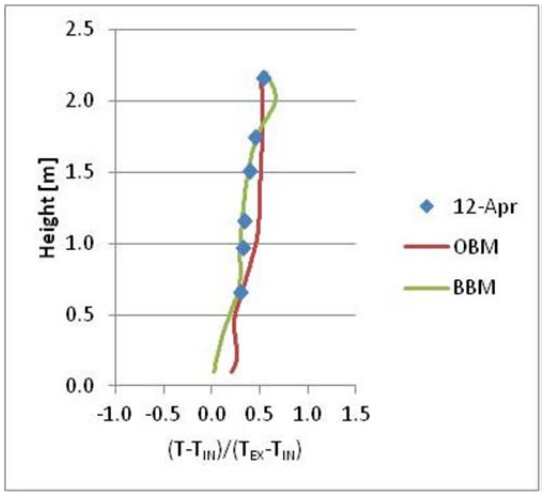

(c) Left 1

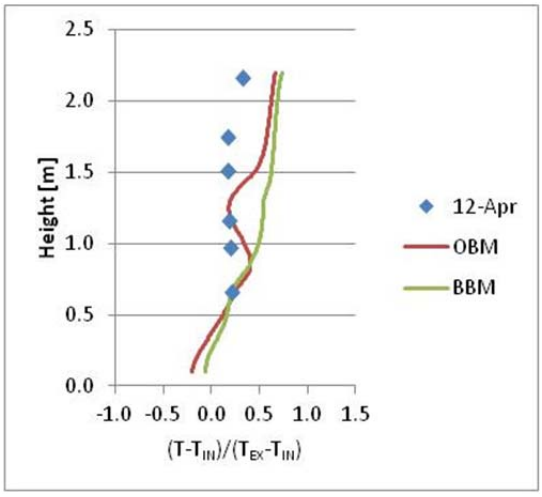

(d) Left 2

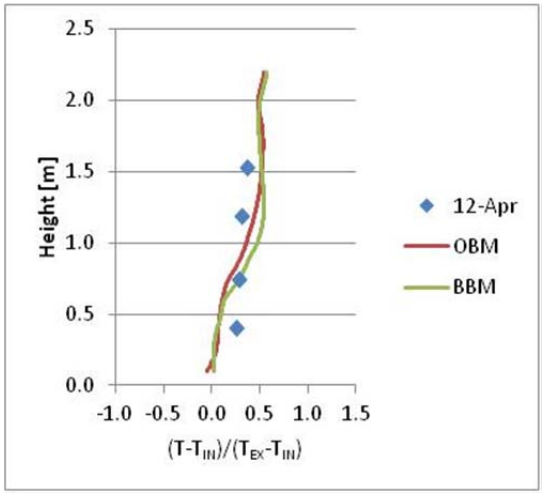

(e) Right 1

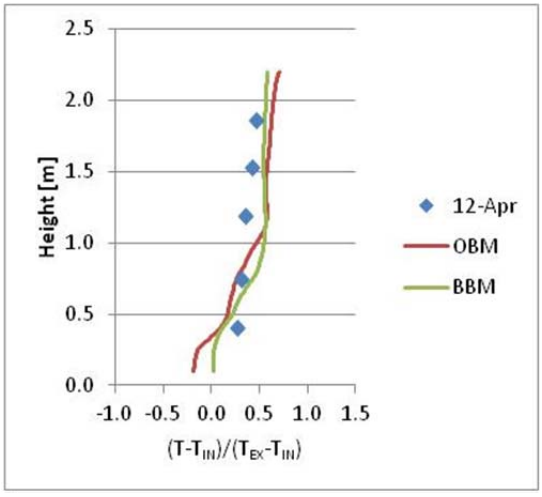

(f) Right 2

Figure 5–72: Normalized Temperature Profiles for Experiment 8

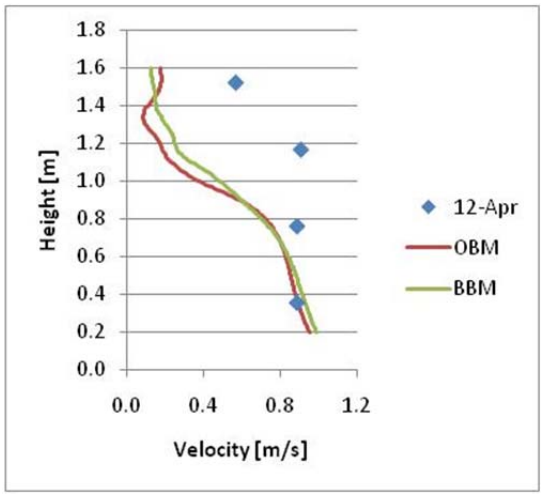

(a) At Perforated Floor Tile

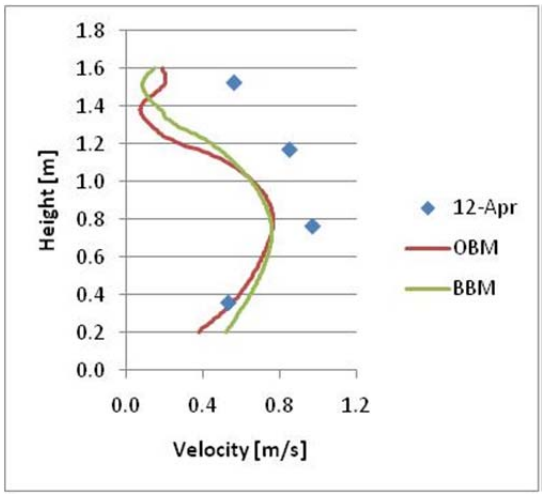

(b) Forward Edge of Perforated Floor Tile

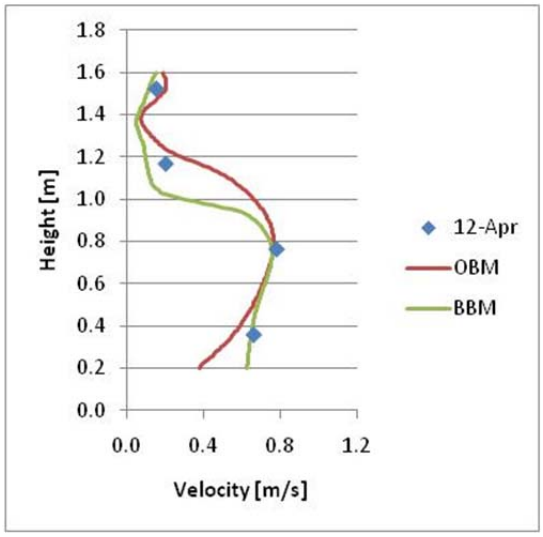

(c) Rack Inlet

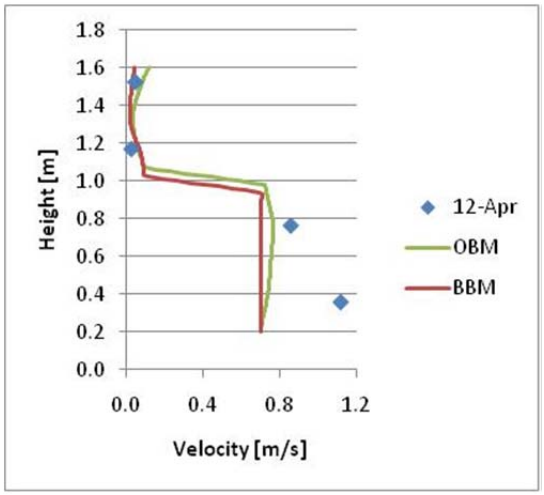

(d) Rack Outlet

Figure 5–73: Velocity Profiles for Experiment 8

5.11 Experiment 9 - 4 kW, Top Servers Removed

This experiment was done to determine the magnitude of recirculated airflow through rack cabinet when it was not full and not blocked with a blanking panel as in Experiment 8. However, it was observed during the experiment that there did not seem to be any significant recirculation through the rack (Figure 5–74). With these observations in mind, the same model that was set up for the previous experiment was used and results were generally good.

Figure 5–74: Flow visualization of airflow behind rack with top server simulators removed

Temperature agreements were good across all of the poles with the exceptions of the Front and Left-2 poles for the BBM (Table 5–35). The issue with front pole in the BBM is due to slightly more pronounced over-predictions at the top (Figure 5–75 (a)). A similar phenomena was observed in Experiment 8 and probably has similar causes.

Table 5–35: Average Error for Experiment 9 Temperature Poles

	OBM		BBM	
	[°C]	[°F]	[°C]	[°F]
F Avg	3.7	6.7	4.1	7.4
B Avg	2.6	4.7	2.1	3.8
L1	1.5	2.7	0.3	0.5
L2	3.1	5.6	4.9	8.8
R1	1.4	2.5	1.8	3.2
R2	2.2	4.0	2.1	3.8

Velocities were also within good agreement (Table 5–36). There were under-predictions near the top of the perforated floor tile pole, but this appears to be a slight mis-calibration of the throw and decay values.

Table 5–36: Average Error for Experiment 9 Velocities Poles

	OBM		BBM	
	[m/s]	[fpm]	[m/s]	[fpm]
Perforated Floor Tile	0.31	61	0.32	63
Front Edge of PFT	0.27	53	0.29	57
Rack Inlet	0.05	10	0.04	8
Rack Exhaust	0.14	28	0.16	31

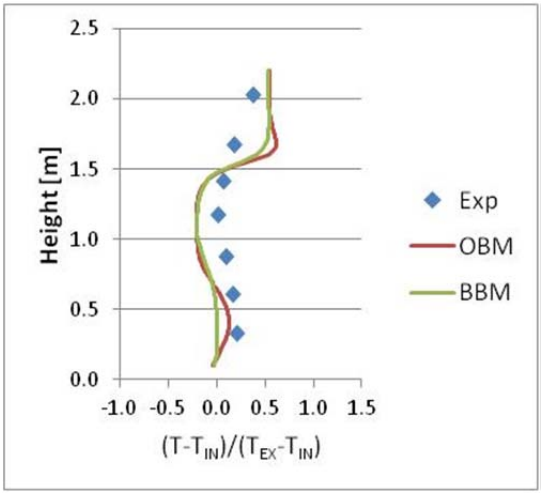

(a) Front Average

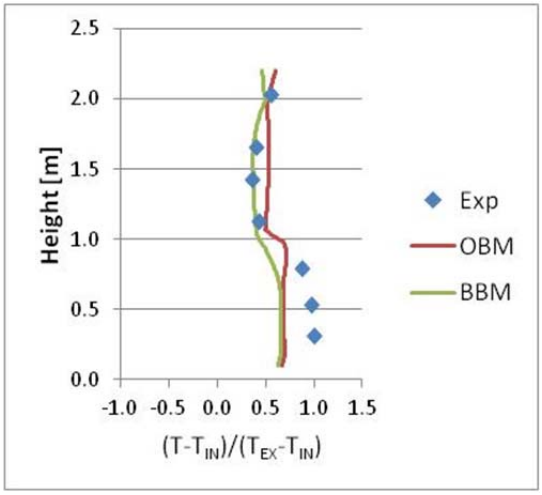

(b) Back Average

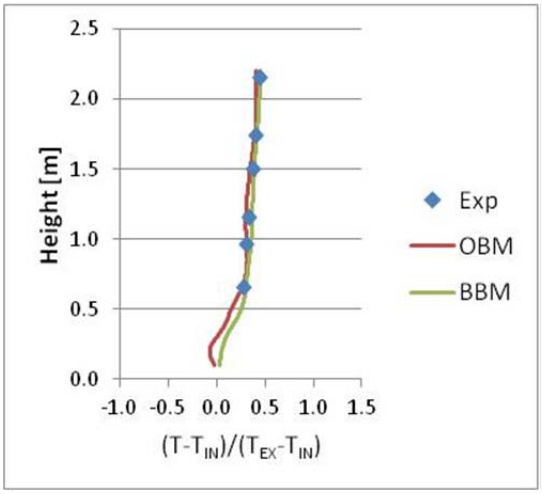

(c) Left 1

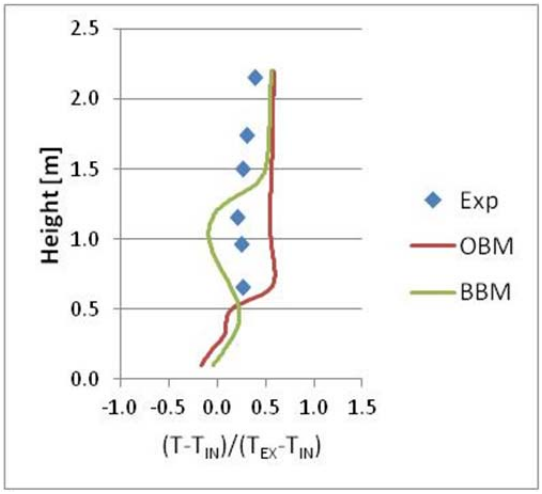

(d) Left 2

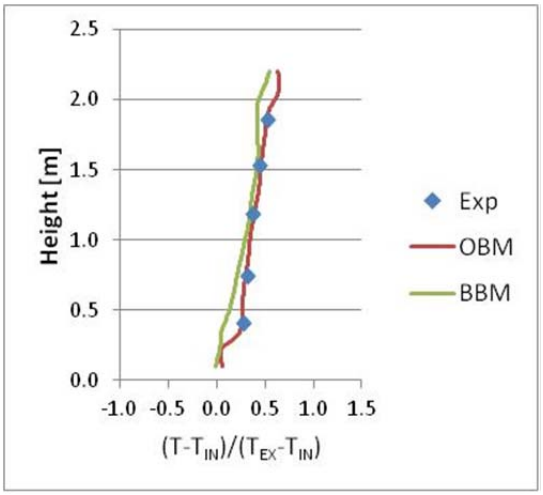

(e) Right 1

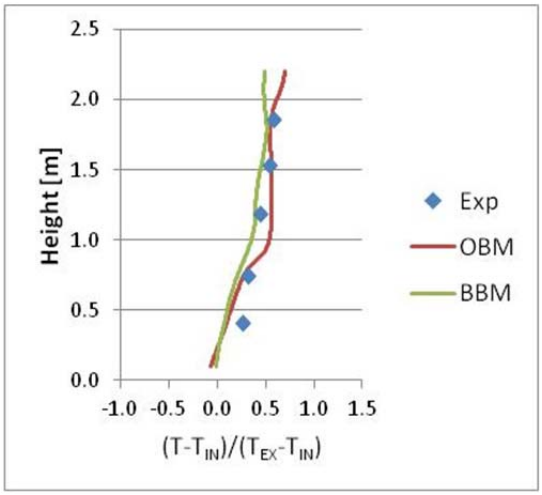

(f) Right 2

Figure 5–75: Normalized Temperature Profiles for Experiment 9

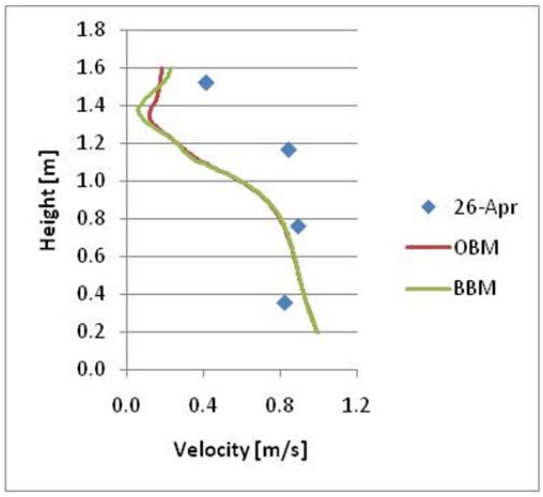

(a) At Perforated Floor Tile

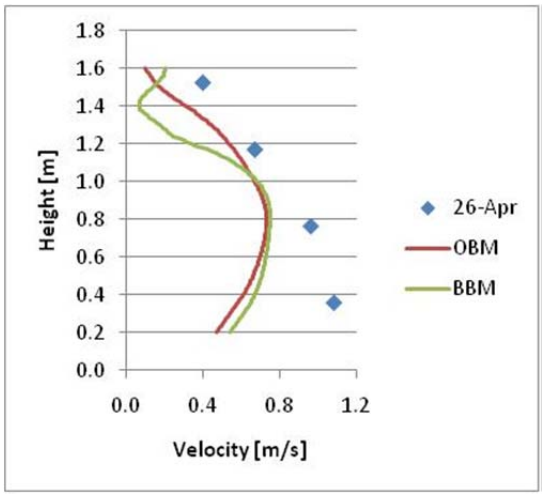

(b) Forward Edge of Perforated Floor Tile

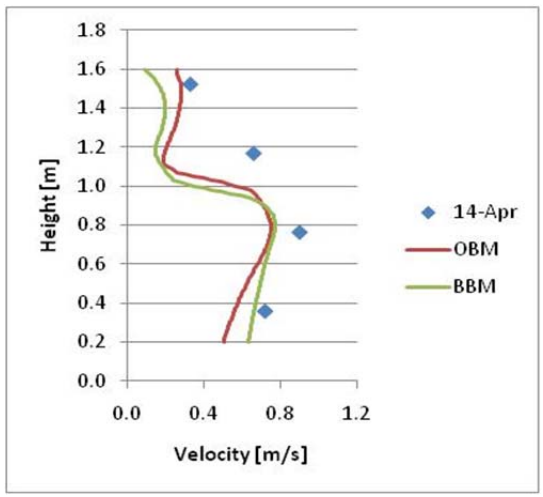

(c) Rack Inlet

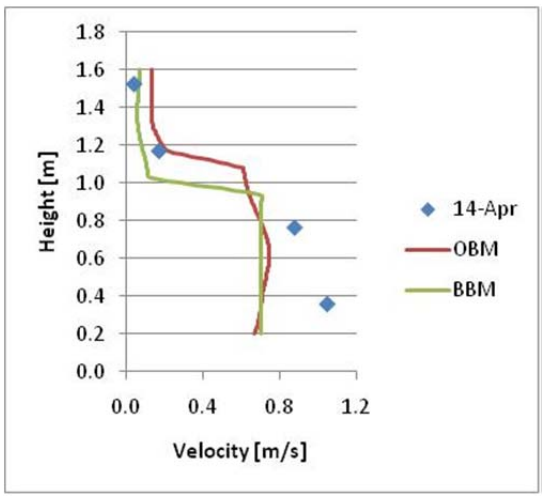

(d) Rack Outlet

Figure 5–76: Velocity Profiles for Experiment 9

5.12 Experiment 10 - 5.5 kW, Located at Edge of Perforated Floor Tile

The final experiment considers a rack that is up against the edge of a perforated floor tile. This is a fairly common condition in data centers. It was not used in the original validation experiments because it was felt that it would be better to keep the rack somewhat equidistant between the perforated floor tile and the return air (the return air could not be moved significantly due to how the ceiling partitions were arranged).

From Table 5–37, one can see decent temperature agreement except for the Right-1 pole. Due to the rack being forward, this pole was in the middle of a fairly open space. As observed in past experiments, the models often have a difficult time capturing the recirculated hot air here. Figure 5–77 shows that most of the poles that are off a under-predictions, but ones that follow the trends of the experimental data points.

Table 5–37: Average Error for Experiment 10 Temperature Poles

	OBM		BBM	
	[°C]	[°F]	[°C]	[°F]
F Avg	2.5	4.5	2.9	5.2
B Avg	3.1	5.6	2.7	4.9
L1	3.4	6.1	4.0	7.2
L2	3.8	6.8	4.3	7.7
R1	5.4	9.7	6.2	11.2
R2	3.6	6.5	3.5	6.3

Table 5–38 shows that both the inlet and exhaust have good temperature agreements but that there are minor issues with the airflow from the perforated floor tile. This is due to an under-prediction towards the upper part of the supply air column (Figure 5–78).

Table 5–38: Average Error for Experiment 10 Velocities Poles

	OBM		BBM	
	[m/s]	[fpm]	[m/s]	[fpm]
Perforated Floor Tile	0.39	77	0.38	75
Rack Inlet	0.14	28	0.15	30
Rack Exhaust	0.14	28	0.15	30

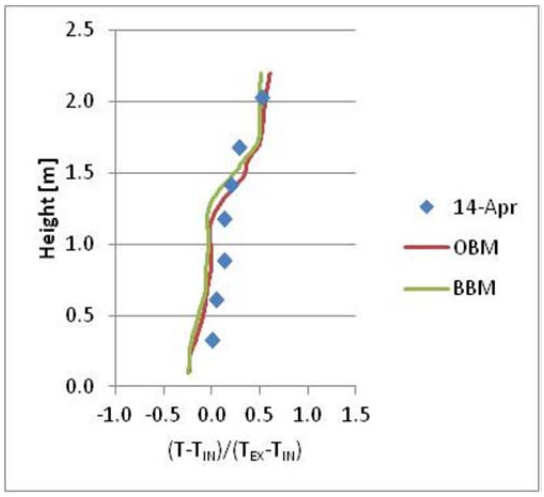

(a) Front Average

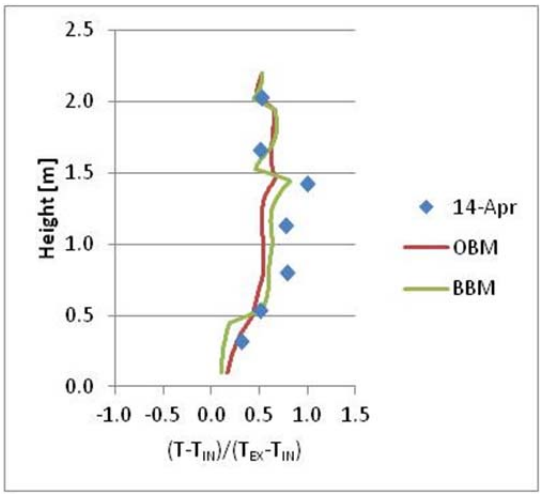

(b) Back Average

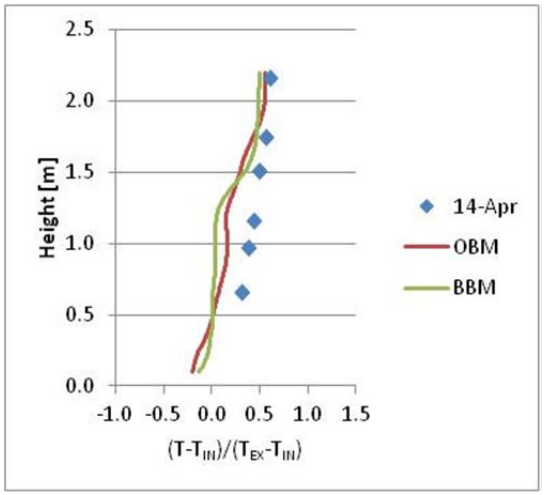

(c) Left 1

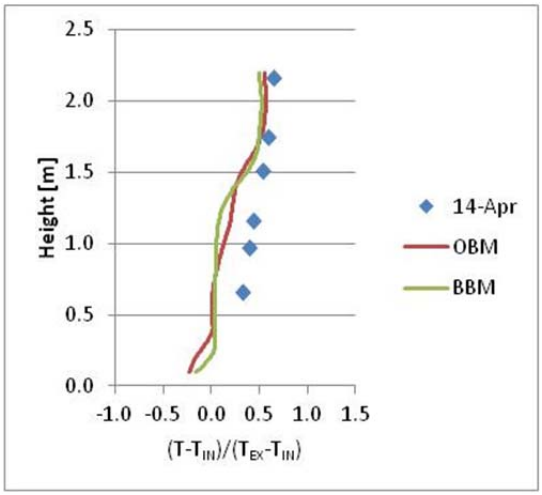

(d) Left 2

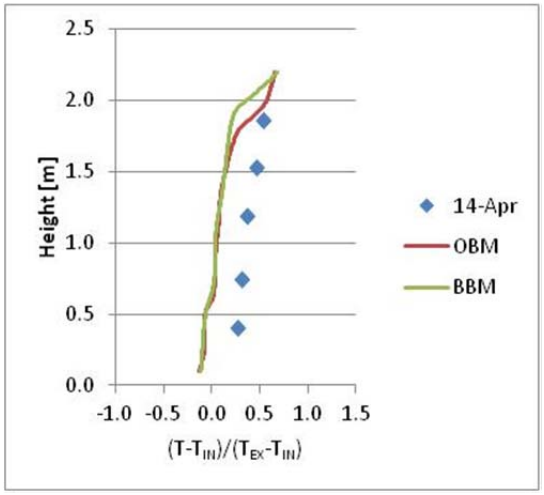

(e) Right 1

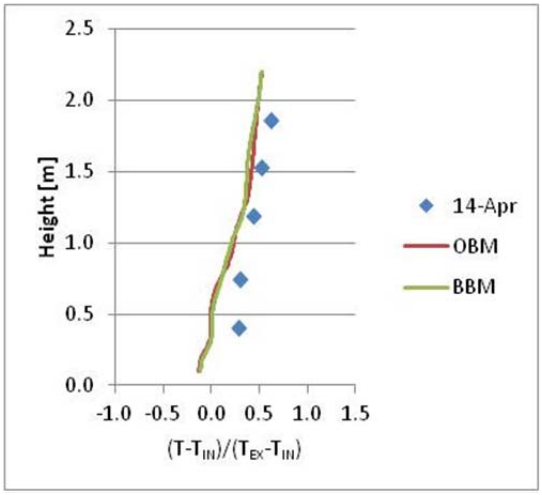

(f) Right 2

Figure 5–77: Normalized Temperature Profiles for Experiment 10

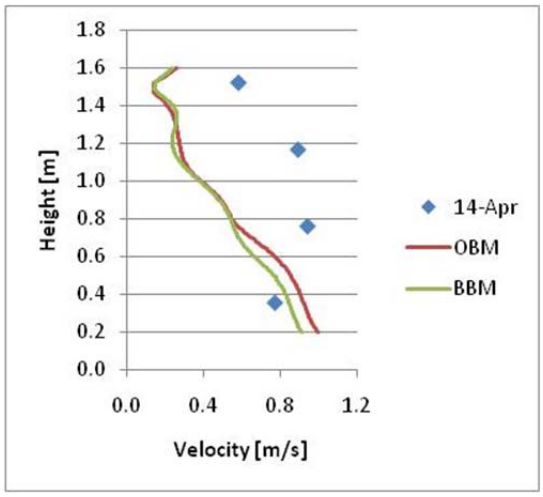

(a) At Perforated Floor Tile

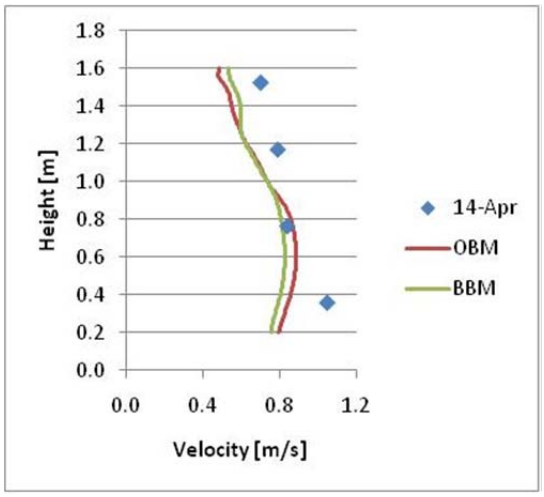

(b) Rack Inlet

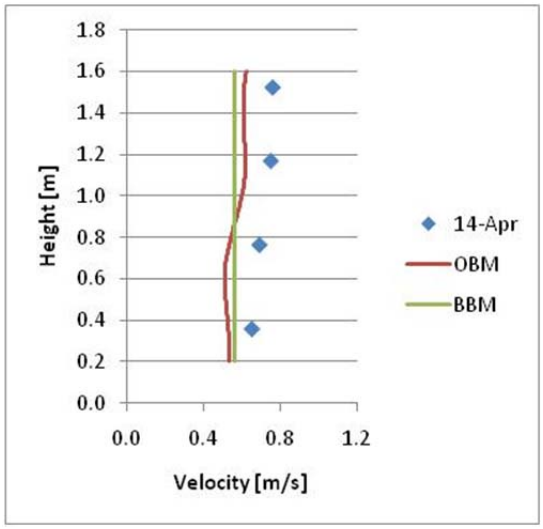

(c) Rack Outlet

Figure 5–78: Velocity Profiles for Experiment 10

6 CONCLUSIONS

6.1 Discussion

This study examined the development of simplified boundary conditions for a rack model designed for use in a computer data center. It was found that either an open or closed model performed at a generally equivalent level. Over 80% of all poles were found to be within 3°C (5.4 °F) or 0.3 m/s (59 fpm), for temperature and velocity poles, of experimental values. Given the simplified nature of these models this can be considered good agreement.

6.1.1 Perforated Floor Tile

It was found that correctly modeling the boundary conditions through the perforated floor tile was more important that had been initially suspected. Simple inlet boundary conditions were not sufficient and tended to significantly over-predict velocities. Instead, modeling the perforated floor tile as a grid of nozzle diffusers, following the momentum method outlined in ASHRAE RP-1009, was found to produce acceptable agreement. In order to use this method it was necessary to get a reasonable approximation of the terminal velocity, throw and decay constant. However, it was also found that these do not require a high degree of accuracy to produce acceptable temperature results. Methods for determining these constants are also outline in ASHRAE RP-1009. Most errors in velocity measurements appear due to issues with calculating the throw value, K. In general these lead to under-predictions of velocities near the upper anemometers in the poles. Interestingly though, temperatures towards the front of the room (closer to the supply air) were generally more accurate.

6.1.2 Temperature predictions

Both models proved better at predicting temperatures towards the front of the rack. Towards the back of the rack, both models tended to under-predict. There appears to be an under-prediction of rack exhaust recirculation within the back of the room. This is could possibly be due to exfiltration, but this is not as likely since the room was fairly well sealed and any problem points would have most likely been around the door in the front of the room. At the same time it is possible that the room exhaust was not properly modeled. Of all of the boundary conditions in the room, this was one of the least examined in this study. It is possible that air in the room, particularly towards the back, may have had a longer residence time that was predicted.

6.1.3 Open Box and Black Box Model Boundary Conditions

Assuming a constant pressure (velocity) across an individual server simulator based on the volumetric flow rate was validated as a reasonable assumption. The Open Box Model was found to be easier to implement, but it does give results within the rack which must be neglected. For the Black Box Model, extracting the temperatures from the upwind

cells at the rack inlet plate and adding the appropriate amount of enthalpy based on the mass flow rate and portion of the rack load was also found to be a valid assumption. The black box model is attractive in that it is fairly common in data center modeling and its validation by experimental results for both temperature and velocity add credibility to its use as a design tool.

6.2 Future Work

Future work could be done on quantifying the return air boundary condition and its effect on temperatures towards the rear of the room. It would also be useful to try this experiment in a larger room that could accommodate PIV and flow visualization measurements to better quantify the recirculation patterns around the return air. Finally, it would be interesting to examine the rack model for a rack with actual servers with smaller thicknesses of varying loads and volumetric flow rates.

7 REFERENCES

42U. (2011). *42U*. Retrieved June 2009, from 42U: http://www.42u.com/cooling/computer-room-cooling.htm

Akella, C., & Ortega, A. (2005). Experimental characterization of pressure loss through EMI shields and 1RU card passages in dense electronic enclosures. *Twenty First Annual IEEE Semiconductor Thermal Measurement and Management Symposium (IEEE Cat. No.05CH37651)* (pp. 86-93). Piscataway, NJ: IEEE.

APC. (2005). *Server Thermal Simulator Users Manual.* Billerica, MA: American Power Conversion Corporation.

ASHRAE. (2005). *Fundamentals.* Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

ASHRAE. (2006). *Method of testing for rating the performance of air outlet and inlets, ANSI/ASHRAE Standard 70-2006.* Atlanta: American Society of Heating, Refrigeration and Air Conditioning Engineers.

ASHRAE. (2008). *Environmental Guidelines for Datacom Equipment.* Atlanta, GA: Amer. Soc. Heating Refrig. and Air Cond. Eng. Inc.

ASHRAE. (2008). *Thermal Guidelines for Data Processing Environments.* Atlanta: American Society of Heating, Refrigeration and Air-Conditioning Engineergs.

ASM Modular Products. (2010). *http://www.asmproducts.com.* Retrieved February 2, 2010, from ASM Modular Products.

ASM Modular Products. (2010, February 2). *http://www.asmproducts.com.* Retrieved February 2, 2010, from ASM Modular Products.

Beaty, D. (2005). Cooling data centers with raised-floor plenums. *HPAC Heating, Piping, AirConditioning Engineering*, 58-65.

Beaty, D., & Davidsor, T. (2003). New Guidelines for Data Center Cooling. *ASHRAE Journal*, 28-36.

Beaty, D., Chauhan, N., & Dyer, D. (2005). High density cooling of data centers and telecom facilities - Part 1. *ASHRAE Transactions* (pp. 921-931). Orlando, FL: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Beaty, D., Chauhan, N., Dyer, D., & Davis, B. (2005). High density cooling of data centers and telecom facilities - Part 2. *ASHRAE Transactions* (pp. 932-944). Orlando, FL: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Bedekar, V., Karajgikar, S., Agonafer, D., Iyyengar, M., & Schmidt, R. (2006). Effect of CRAC location on fixed rack layout. *Thermomechanical Phenomena in Electronic Systems -Proceedings of the Intersociety Conference* (pp. 421-425). San Diego, CA: Institute of Electrical and Electronics Engineers Inc.

Beitelmal, A., & Patel, C. (2007). Thermo-fluids provisioning of a high performance high density data center. *Distributed and Parallel Databases*, 227-38.

Bhopte, S., Agonafer, D., Schmidt, R., & Sammakia, B. (2006). Optimization of data center room layout to minimize rack inlet air temperature. *Transactions of the ASME. Journal of Electronic Packaging*, 380-7.

Bhopte, S., Sammakia, B., Schmidt, R., Iyengar, M., & Agonafer, D. (2006). Effect of under floor blockages on data center performance. *Thermomechanical Phenomena in Electronic Systems -Proceedings of the Intersociety Conference* (pp. 426-433). San Diego, CA: Institute of Electrical and Electronics Engineers Inc.

CHAM. (2010, February 22). *FLAIR User's Guide - CHAM Technical Report TR 313*. Retrieved May 2, 2011, from Phoenix: http://www.cham.co.uk/phoenics/d_polis/d_docs/tr313/tr313.htm

CHAM Ltd. (2010, January 17). *CHAM Case Study – Data Centre Simulation for Thomas Reuters, UK.* Retrieved May 2010, from Phoenics: http://www.cham.co.uk

CHAM. (n.d.). *The IMMERSOL model of Radiative Heat Transfer*. Retrieved July 5, 2011, from POLIS: http://www.cham.co.uk/phoenics/d_polis/d_enc/enc_rad3.htm

Chen, Q., & Srebric, J. (2001). *How to Verity, Validate and Report Indoor Environmental Modeling CFD Anlaysis.* Atlanta: American Society of Heating, Refrigeration and Air Conditioning Engineers, Inc.

Chen, Q., & Srebric, J. (2001). *Simplified Diffuser Boundary Conditions for Numerical Room Airflow Models, ASHRAE RP-1009.* Altanta: American Society of Heating, Refrigeration and Air Conditioning Engineers, Inc.

Chen, Q., & Zhai, Z. (2004). The Use of CFD Tools for Indoor Environmental Design. *Advanced Building Simulation*, 119-140.

Choi, J., Kim, Y., Sivasubramaniam, A., Srebric, J., Wang, Q., & Lee, J. (2007). Modeling and managing thermal profiles of rack-mounted servers with ThermoStat. *Proceedings - International Symposium on High-Performance Computer Architecture* (pp. 205-215). Scottsdale, AZ: Inst. of Elec. and Elec. Eng. Computer Society.

EPA. (2007). *Report to Congress on Server and Data Center Energy Efficiency: Public Law 109-431.* Washington, D.C.: U.S. Environmental Protection Agency: Energy Star Program.

EPA. (2009, November 12). *Energy Star Data Center Infrastructure Rating Development Update.* Retrieved June 2010, from EPA: energystar.gov

Fanara, A. (2008). *Global Trends: Government Initiatives to Reduce Energy Use in Data Centers.* Retrieved June 2010, from EPA: energystar.gov

Furihata, Y., Hayama, H., Enai, M., & Mori, T. (2003). Efficient cooling system for IT equipment in a data center. *INTELEC '03. Proceedings. The 25th International*

Telecommunications Energy - INTELEC '03 (IEEE Cat. No.03CH7481) (pp. 152-9). Tokyo, Japan: Inst. of Electron., Inf. and Commun. Eng.

Furihata, Y., Hayama, H., Enai, M., Mori, T., & Kishita, M. (2004). Improving the efficiency of cooling systems in data centers considering equipment characteristics. *INTELEC 26th Annual International Telecommunications Energy Conference (IEEE Cat. No.04CH37562)* (pp. 32-7). Piscataway, NJ: IEEE.

Gondipalli, S., Bhopte, S., Sammakia, B., Iyengar, M., & Schmidt, R. (2008). Effect of isolating cold aisles on rack inlet temperature. *2008 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM '08)* (pp. 1247-54). Piscataway, NJ: IEEE.

Hamann, H., Iyengar, M., & O'Boyle, M. (2008). The impact of air flow leakage on server inlet air temperature in a raised floor data center. *2008 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM '08)* (pp. 1153-60). Piscataway, NJ: IEEE.

Hamann, H., Lacey, J., O'Boyle, M., Schmidt, R., & Iyengar, M. (2008). Rapid three-dimensional thermal characterization of large-scale computing facilities. *IEEE Transactions on Components and Packaging Technology*, 444-8.

Herrlin, M. (2005). Rack cooling effectiveness in data centers and telecom central offices: The Rack Cooling Index (RCI). *ASHRAE Transactions* (pp. 725-731). Denver, CO: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Herrlin, M. (2006). A new tool for evaluating and designing the thermal environment in telecom central offices. *INTELEC, International Telecommunications Energy Conference (Proceedings).* Providence, RI: IEEE, Power Electronics Society.

Herrlin, M., & Belady, C. (2006). Gravity-assisted air mixing in data centers and how it affects the Rack cooling effectiveness. *Thermomechanical Phenomena in Electronic Systems -Proceedings of the Intersociety Conference* (pp. 434-438). San Diego, CA: Institute of Electrical and Electronics Engineers Inc.

Herrlin, M., & Khankari, K. (2008). Method for optimizing equipment cooling effectiveness and HVAC cooling costs in telecom and data centers. *ASHRAE Transactions* (pp. 17-21). New York, NY: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Hwang, Y., Radermacher, R., Spinazzola, S., & Menachery, Z. (2004). Performance measurements of a forced convection air-cooled rack. *The Ninth Intersociety Conference on Thermal and Thermomechanical Phenomena In Electronic Systems (IEEE Cat. No.04CH37543)* (pp. 194-8). Piscataway, NJ: IEEE.

Ishimine, J., Ohba, Y., Ikeda, S., & Suzuki, M. (2009). Improving IDC cooling and air conditioning efficiency. *Fujitsu Scientific and Technical Journal*, 123-133.

Iyengar, M., & Schmidt, R. (2009). Analytical modeling for thermodynamic characterization of data center cooling systems. *Journal of Electronic Packaging*, 021009 (9 pp.).

Jeonghwan, C., Youngjae, K., Sivasubramaniam, A., Srebric, J., Qian, W., & Joonwon, L. (2008). A CFD-based tool for studying temperature in rack-mounted servers. *IEEE Transactions on Computers*, 1129-42.

Karki, K., Patankar, S., & Radmehr, A. (2003). Techniques for controlling airflow distribution in raised-floor data centers. *Advances in Electronic Packaging* (pp. 621-628). Maui, HI: American Society of Mechanical Engineers.

Karki, K., Radmehr, A., & Patankar, S. (2003). Use of computational fluid dynamics for calculating flow rates through perforated tiles in raised-floor data centers. *HVAC and R Research*, 153-166.

Karki, K., Radmehr, A., & Patankar, S. (2007). Prediction of distributed air leakage in raised-floor data centers. *ASHRAE Transactions* (pp. 219-226). Dallas, TX: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Karlsson, J., & Moshfegh, B. (2005). Investigation of indoor climate and power usage in a data center. *Energy and Buildings*, 1075-1083.

Khattar, M. (2010). Data Center Retrofit: Heat Containment and Airflow Managment. *ASHRAE Journal*, 40-49.

Kurkjian, C., & Glass, J. (2005). Air-conditioning design for data centers - Accommodating current loads and planning for the future. *ASHRAE Transactions* (pp. 715-724). Denver, CO: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Li, G., Li, M., Azarm, S., Rambo, J., & Joshi, Y. (2007). Optimizing thermal design of data center cabinets with a new multi-objective genetic algorithm. *Distributed and Parallel Databases*, 167-92.

Malone, C., Vinson, W., & Bash, C. (2008). Data center TCO benefits of reduced system airflow. *2008 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM '08)* (pp. 1199-202). Piscataway, NJ: IEEE.

Moore, J., Chase, J., & Ranganathan, P. (2006). Weatherman: automated, online, and predictive thermal mapping and management for data centers. *Proceedings. 3rd International Conference on Autonomic Computing (IEEE Cat. No. 06EX1303)* (pp. 155-64). Piscataway, NJ: IEEE.

Mukherjee, T., Qinghui, T., Ziesman, C., Gupta, S., & Cayton, P. (2007). Software architecture for dynamic thermal management in datacenters. *Proceedings of the 2007 2nd International Conference on Communication System Software, Middleware and Workshops* (pp. 580-90). Piscataway, NJ: IEEE.

National Instruments. (2010). *NI Developer Zone*. Retrieved 02 10, 2011, from http://zone.ni.com/devzone/cda/tut/p/id/4237: http://zone.ni.com/devzone/cda/tut/p/id/4237

National Instruments. (2011, 02 10). *NI Developer Zone*. Retrieved 02 10, 2011, from http://zone.ni.com/devzone/cda/tut/p/id/4237: http://zone.ni.com/devzone/cda/tut/p/id/4237

Nelson, G. (2007). *Master's Thesis: Development of an Experimentally-Validated Compact Model of a Server Rack.* Georgia Institute of Technology.

Patankar, S., Karki, K., VanGlder, J., & Kearney, M. (2004). Distribution of cooling airflow in a raised-floor data center. *ASHRAE Transactions* (pp. 629-634). Nashville, TN: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Patel, C., Bash, C., Belady, C., Stahl, L., & Sullivan, D. (2001). Computational Fluid Dynamics Modeling of High Computer Density Data Centers to Assure System Inlet Air Specifications. *Advances in Electronic Packaging* (pp. 821-829). Kauai, HI: American Society of Mechanical Engineers.

Patel, C., Sharma, R., Bash, C., & Beitelmal, A. (2002). Thermal considerations in cooling large scale high computer density data centers. *ITherm 2002. Eighth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (Cat. No.02CH37258)* (pp. 767-76). Piscataway, NJ: IEEE.

Patterson, M. K. (2009). Energy-Efficiency Through The Integration Of Information And Communications Technology Management And Facilities Controls. *Proceedings of IPACK2009.* San Francisco, CA: American Society of Mechanical Engineers.

PG&E. (2006). *High Performance Data Centers: A Design Guidelines Sourcebook.* Berkley, CA: Pacific Gas & Electric.

Prisco, J. F. (2007). Characterization of a high-density data center. *ASHRAE Transactions* (pp. 137-148). Dallas, TX: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Prisco, J., Lembke, P., & Misgen, M. (2008). Electrical and heat load planning - Keep your data processing environment running. *ASHRAE Transactions* (pp. 22-27). New York, NY: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Qinghui, T. (2007). Sensor-based fast thermal evaluation model for energy efficient high-performance datacenters. *Fourth International Conference on Intelligent Sensing and Information Processing* (pp. 122-7). Piscataway, NJ: IEEE.

Radmehr, A., Schmidt, R., Karki, K., & Patankar, S. (2005). Distributed leakage flow in raised-floor data centers. *Proceedings of the ASME/Pacific Rim Technical Conference and Exhibition on Integration and Packaging of MEMS, NEMS, and Electronic Systems: Advances in Electronic Packaging 2005* (pp. 401-408). San Francisco, CA: American Society of Mechanical Engineers.

Rambo, J., & Joshi, Y. (2005). Thermal performance metrics for arranging forced air cooled servers in a data processing cabinet. *Transactions of the ASME. Journal of Electronic Packaging*, 452-9.

Rambo, J., & Joshi, Y. (2006). Convective transport processes in data centers. *Numerical Heat Transfer; Part A: Applications*, 923-945.

Rambo, J., & Joshi, Y. (2007). Modeling of Data Center Airflow and Heat Transfer: State of the Art and Future Trends. *Distrib Parallel Databases*, 21:193–225.

Rambo, J., & Joshi, Y. (2007). Reduced-order modeling of turbulent forced convection with parametric conditions. *International Journal of Heat and Mass Transfer*, 539-551.

Rambo, J., Nelson, G., & Joshi, Y. (2007). Airflow distribution through perforated tiles in close proximity to computer room air-conditioning units. *ASHRAE Transactions* (pp. 124-135). Long Beach, CA: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Rolander, N., Rambo, J., Joshi, Y., Allen, J., & Mistree, F. (2006). An approach to robust design of turbulent convective systems. *Journal of Mechanical Design, Transactions of the ASME*, 844-855.

Samadiani, E., Joshi, Y., & Mistree, F. (2008). The thermal design of a next generation data center: a conceptual exposition. *Journal of Electronic Packaging*, 041104 (8 pp.).

Schmidt, R. (2001). Effect of Data Center Characteristics on Data Processing Equipment Inlet Temperatures. *Advances in Electronic Packaging* (pp. 1097-1106). Kauai, HI: American Society of Mechanical Engineers.

Schmidt, R., & Cruz, E. (2002). Raised floor computer data center: effect on rack inlet temperatures of chilled air exiting both the hot and cold aisles. *ITherm 2002. Eighth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (Cat. No.02CH37258)* (pp. 580-94). Piscataway, NJ: IEEE.

Schmidt, R., & Cruz, E. (2004). Cluster of high-powered racks within a raised-floor computer data center: Effect of perforated tile flow distribution on rack inlet air temperatures. *Journal of Electronic Packaging, Transactions of the ASME*, 510-518.

Schmidt, R., & Iyengar, M. (2007). Best practices for data center thermal and energy management - Review of literature. *ASHRAE Transactions* (pp. 206-218). Dallas, TX: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Schmidt, R., & Iyengar, M. (2007). Comparison between underfloor supply and overhead supply ventilation designs for data center high-density clusters. *ASHRAE Transactions* (pp. 115-125). Dallas, TX: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Schmidt, R., Belady, C., Classen, A., Davidson, T., Herrlin, M., Novotny, S., et al. (2004). Evolution of data center environmental guidelines. *ASHRAE Transactions* (pp. 559-566). Anaheim, CA: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Schmidt, R., Cruz, E., & Iyengar, M. (2005). Challenges of data center thermal management. *IBM Journal of Research and Development*, 709-23.

Schmidt, R., Iyengar, M., Mayhugh, S., & Chandra, N. (2006). Thermal profile of world's third fastest supercomputer. *ASHRAE Transactions* (pp. 209-219). Quebec City, QC: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Schmidt, R., Karki, K., & Patankar, S. (2004). Raised-floor data center: Perforated tile flow rates for various tile layouts. *Thermomechanical Phenomena in Electronic Systems -Proceedings of the Intersociety Conference* (pp. 571-578). Las Vegas, NV: Institute of Electrical and Electronics Engineers Inc.

Schmidt, R., Karki, K., Kelkar, K., Radmehr, A., & Patankar, S. (2001). Measurements and predictions of the flow distribution through perforated tiles in raised-floor data centers. *Advances in Electronic Packaging* (pp. 905-914). Kauai, HI: American Society of Mechanical Engineers.

Schmidt, R., Kurkjia, C., Phillips, E., & Scofield, C. (2004). Thermal profile of a high-density data center - Methodology to thermally characterize a data center. *ASHRAE Transactions* (pp. 635-642). Nashville, TN: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Schmidt, R., Lyengar, M., Beaty, D., & Shrivastava, S. (2005). Thermal profile of a high-density data center: Hot spot heat fluxes of 512 W/ft2. *ASHRAE Transactions* (pp. 765-777). Denver, CO: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Sharma, R., Bash, C., Patel, C., Friedrich, R., & Chase, J. (2005). Balance of power: dynamic thermal management for Internet data centers. *IEEE Internet Computing*, 42-9.

Shrivastava, S. (2006). Experimental-numerical comparison for a high-density data center: Hot spot heat fluxes in excess of 500 W/FT2. *Thermomechanical Phenomena in Electronic Systems -Proceedings of the Intersociety Conference* (pp. 402-411). San Diego, CA: IEEE.

Shrivastava, S., Iyengar, M., Sammakia, B., Schmidt, R., & VanGilder, J. (2006). Experimental-Numerical Comparison For A High-Density Data Center: Hot Spot Heat Fluxes In Excess Of 500 W/Ft^2. *IEEE.* San Diego, CA: IEEE.

Shrivastava, S., Iyengar, M., Sammakia, B., Schmidt, R., & VanGilder, J. (2009). Experimental-numerical comparison for a high-density data center: Hot spot heat fluxes in excess of 500 W/ft2. *IEEE Transactions on Components and Packaging Technologies*, 166-172.

Shrivastava, S., VanGilder, J., & Sammakia, B. (2006). A statistical prediction of cold aisle end airflow boundary conditions. *2006 Proceedings. 10th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronics Systems (IEEE Cat. No. 06CH37733C)* (pp. 412-420). Piscataway, NJ: IEEE.

Shrivastava, S., VanGilder, J., & Sammakia, B. (2007). Prediction of cold aisle end airflow boundary conditions using regression modeling. *IEEE Transactions on Components and Packaging Technologies*, 866-874.

Shrivastava, S., VanGilder, J., & Sammakia, B. (2008). Optimization of cluster cooling performance for data centers. *2008 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM '08)* (pp. 1161-6). Piscataway, NJ: IEEE.

Sorell, V. (2008). Current best practices in high-density cooling applications. *ASHRAE Transactions* (pp. 12-16). New York, NY: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Sorell, V., Abougabal, Y., Khankari, K., Gandhi, V., & Watve, A. (2006). An analysis of the effects of ceiling height on air distribution in data centers. *ASHRAE*

Transactions (pp. 623-631). Chicago, IL: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Sorell, V., Escalante, S., & Yang, J. (2005). Comparison of overhead and underfloor air delivery systems in a data center environment using CFD modeling. *ASHRAE Transactions* (pp. 756-764). Denver, CO: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Srebric, J. (2000). *Simplified methodology for indoor environment design.* Cabridge, MA: Department of Architecture, Massachusetts Institute of Technology.

Stahl, L. (2004). Cooling of high heat density rooms today and in the future. *ASHRAE Transactions* (pp. 574-579). Anaheim, CA: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Sullivan, A. (2010, February 4). *Energy Star for Data Centers.* Retrieved June 2010, from EPA: energystar.gov

Tampere University of Technology. (2006, May 11). *Tampere University of Technology - Energy and Process Engineering*. Retrieved February 23, 2011, from Particle Image Velocimetry: http://www.tut.fi/units/me/ener/laitteistot/EFD/PIV.html

Tang, Q., Gupta, S., & Varsamopoulos, G. (2007). Thermal-aware task scheduling for data centers through minimizing heat recirculation. *2007 IEEE International Conference on Cluster Computing (CLUSTER)* (pp. 129-38). Piscataway, NJ: IEEE.

Tang, Q., Gupta, S., & Varsamopoulos, G. (2008). Energy-efficient thermal-aware task scheduling for homogeneous high-performance computing data centers: a cyber-physical approach. *IEEE Transactions on Parallel and Distributed Systems*, 1458-72.

Tang, Q., Gupta, S., Stanzione, D., & Cayton, P. (2006). Thermal-aware task scheduling to minimize energy usage of blade server based datacenters. *2006 2nd IEEE International Symposium on Dependable, Autonomic and Secure Computing.* Los Alamitos, CA: IEEE.

Tang, Q., Mukherjee, T., Gupta, S., & Cayton, P. (2006). Sensor-based fast thermal evaluation model for energy efficient high-performance datacenters. *Proceedings - 4th International Conference on Intelligent Sensing and Information Processing, ICISIP 2006* (pp. 203-208). Bangalore, India: Inst. of Elec. and Elec. Eng. Computer Society.

Tschudi, W., & Fok, S. (2007). Best practices for energy-efficient data centers identified through case studies and demonstration projects. *ASHRAE Transactions* (pp. 450-456). Dallas, TX: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Udakeri, R., Mulay, V., & Agonafer, D. (2008). Comparison of overhead supply and underfloor supply with rear heat exchanger in high density data center clusters. *SEMI-THERM '08. 2008 24th Annual IEEE Semiconductor Thermal Measurement and Management Symposium* (pp. 165-72). Piscataway, NJ: IEEE.

Uptime Institute. (2011). *Heat Density Trends in Data Processing, Computer Systems and Telecommunications Equipment.* Retrieved March 2009, from Uptime Institute: http://uptimeinstitute.org/

VanGilder, J., & Shrivastava, S. (2007). Capture index: An airflow-based rack cooling performance metric. *ASHRAE Transactions* (pp. 126-136). Dallas, TX: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

VanGilder, J., Shrivastava, S., Rahai, H., & Kmankari, K. (2006). Real-time prediction of rack-cooling performance. *ASHRAE Transactions* (pp. 151-162). Quebec City, QC, Canada: Amer. Soc. Heating, Ref. Air-Conditoning Eng. Inc.

Zhai, Z. (2006). Applications of CFD in Building Design: Aspects and Trends. *Indoor and Built Environment*, 15(4): 305-313.

Zhang, X., VanGilder, J., Iyengar, M., & Schmidt, R. (2008). Effect of rack modeling detail on the numerical results of a data center test cell. *2008 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM '08)* (pp. 1183-90). Piscataway, NJ: Institution of Engineering and Technology.

APPENDIX A – ABSOLUTE TEMPERATURE RESULTS

This appendix lists the temperature pole plots for the validation experiments in °C instead of the normalized results.

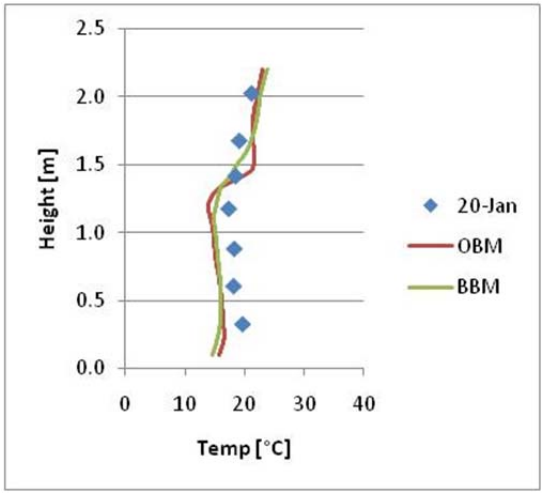

(a) Front Average

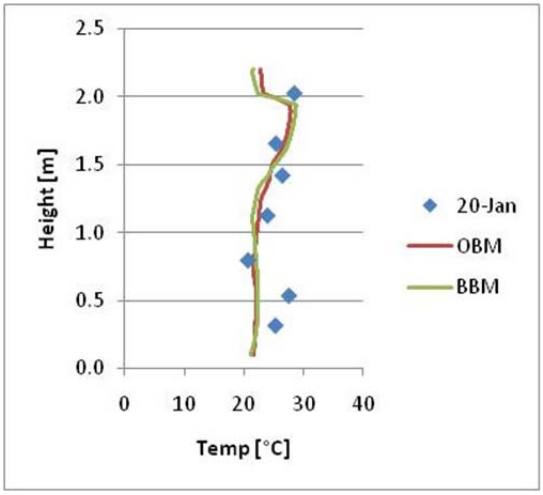

(b) Back Average

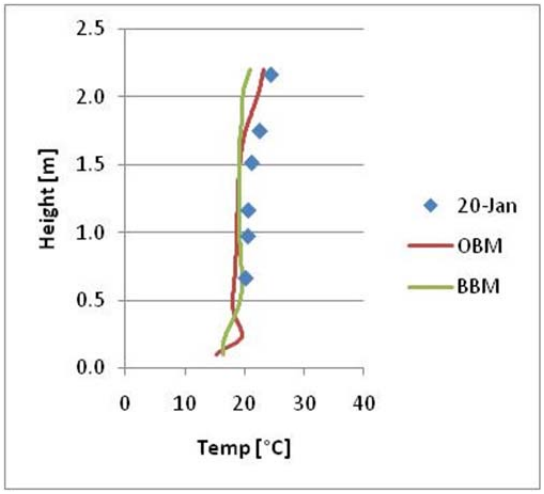

(c) Left 1

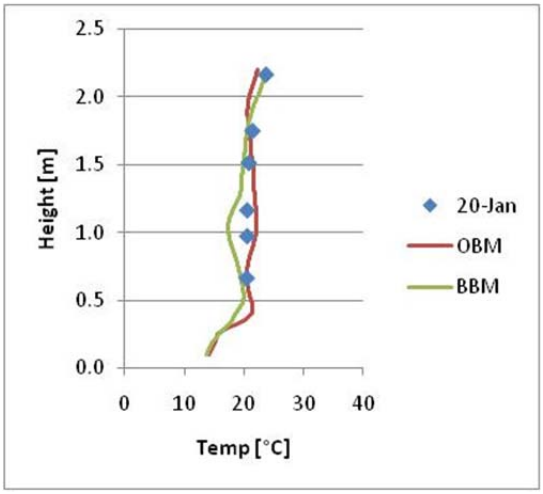

(d) Left 2

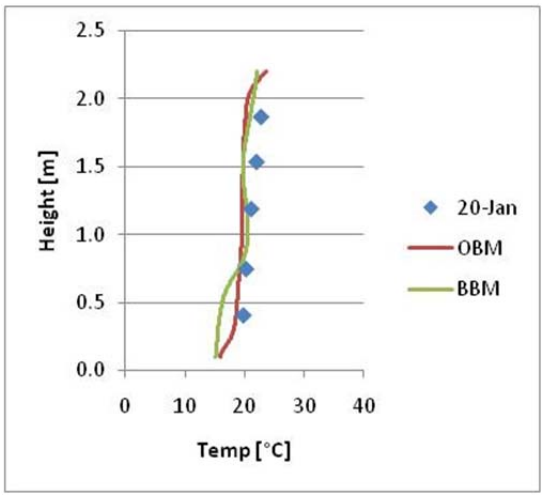

(e) Right 1

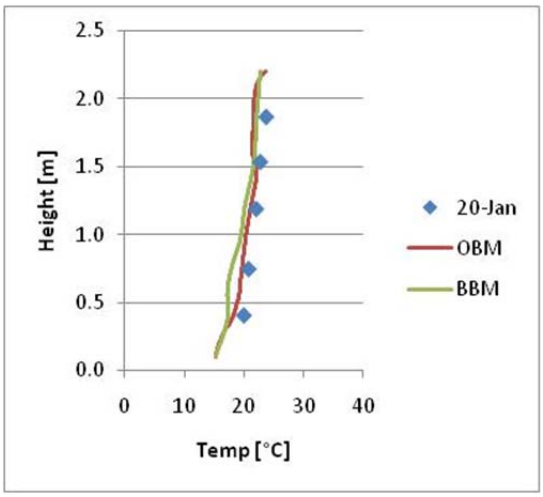

(f) Right 2

Figure 0–79: Temperature Profiles for Experiment 1

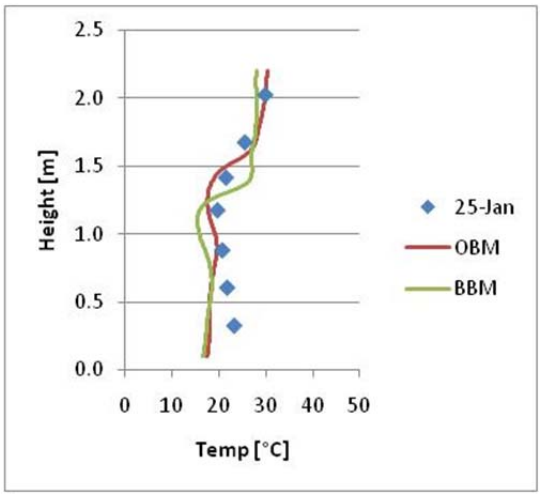

(a) Front Average

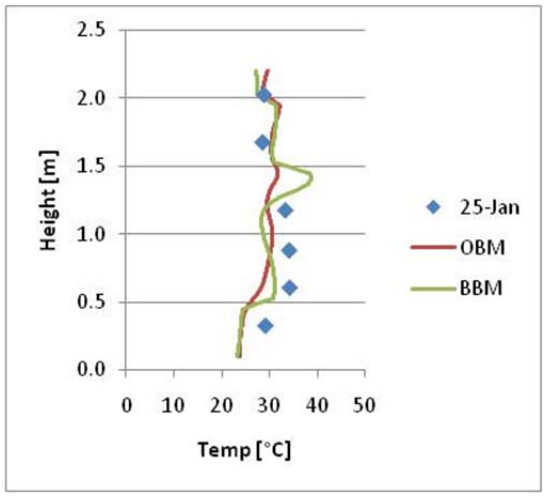

(b) Back Average

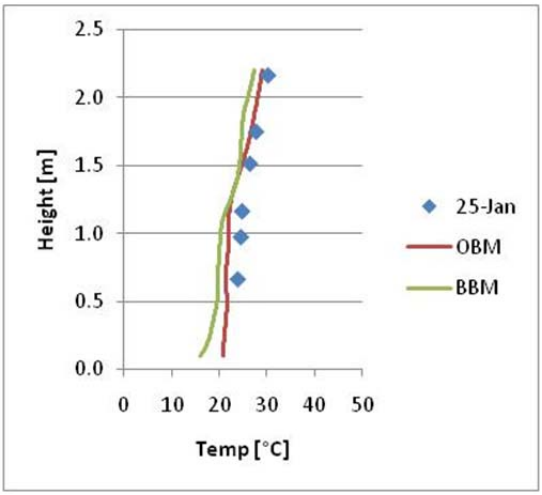

(c) Left 1

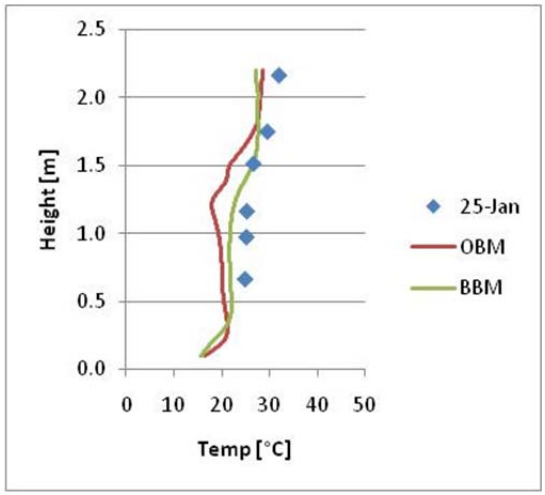

(d) Left 2

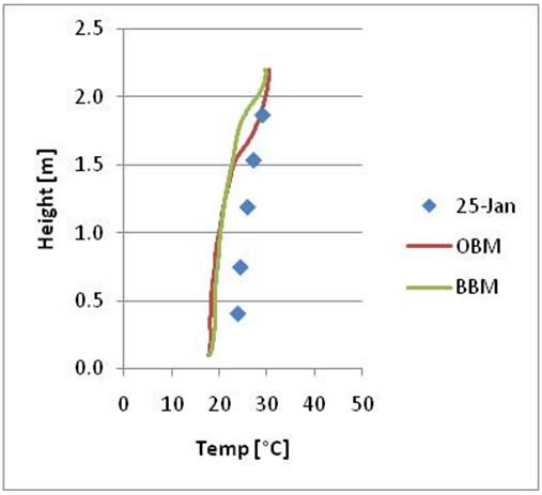

(e) Right 1

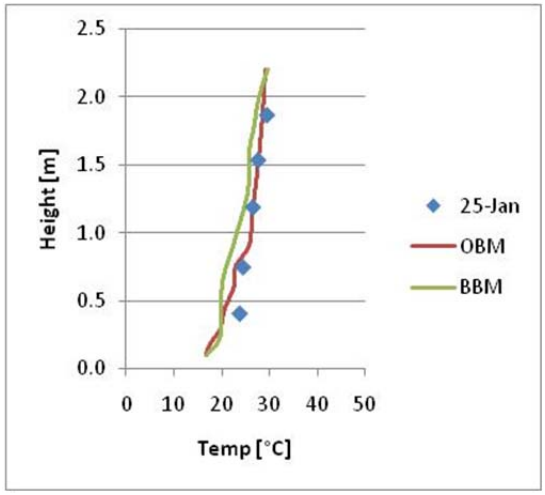

(f) Right 2

Figure 0–80: Temperature Profiles for Experiment 2

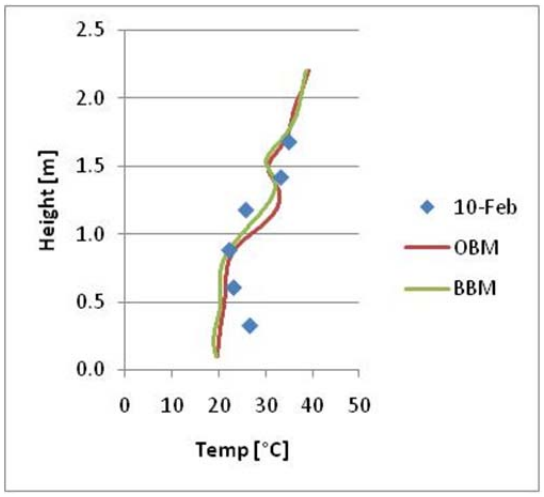

(a) Front Average

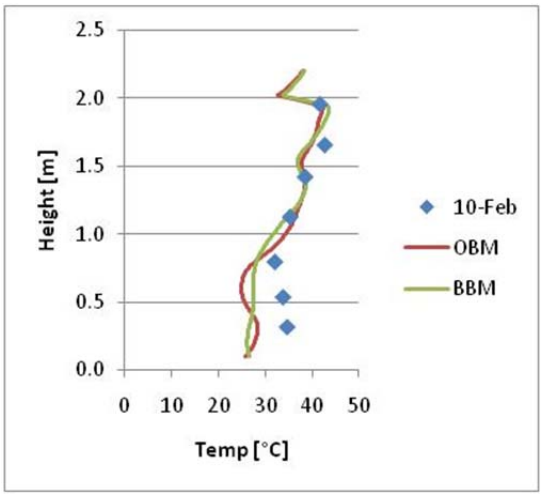

(b) Back Average

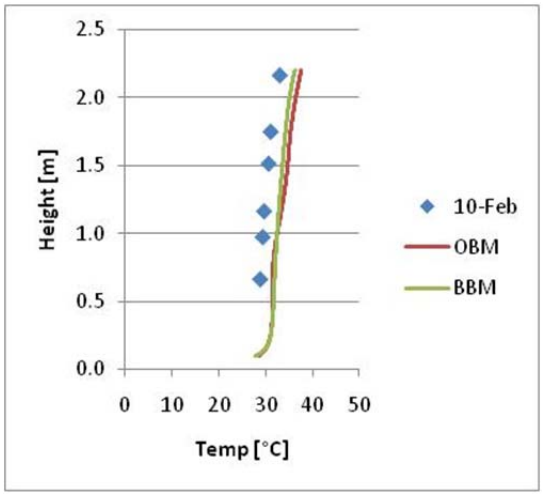

(c) Left 1

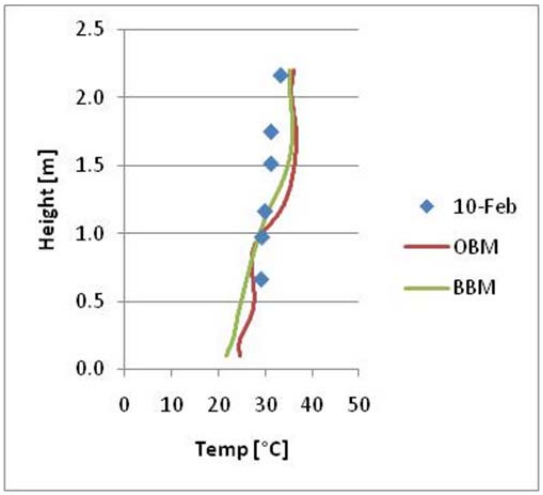

(d) Left 2

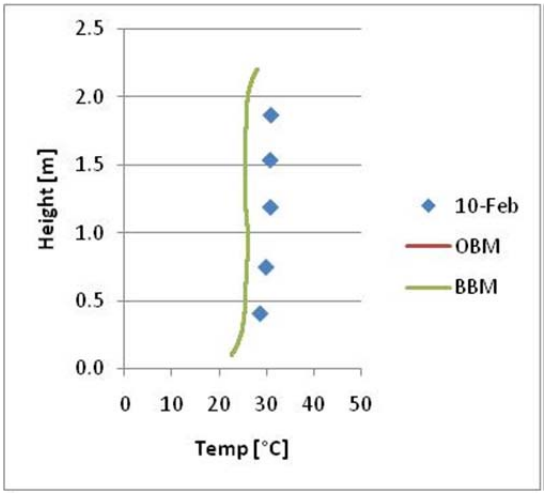

(e) Right 1

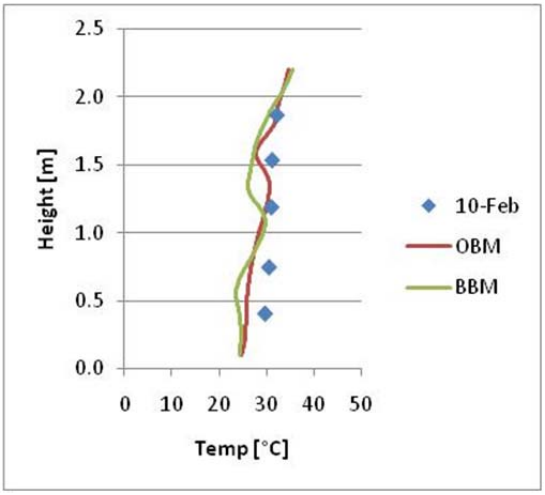

(f) Right 2

Figure 0–81: Temperature Profiles for Experiment 3

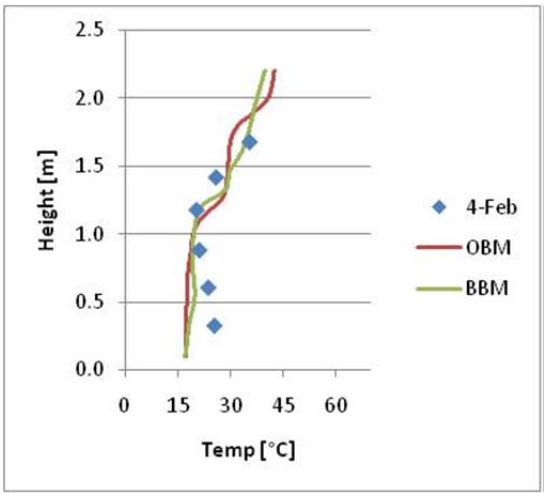

(a) Front Average

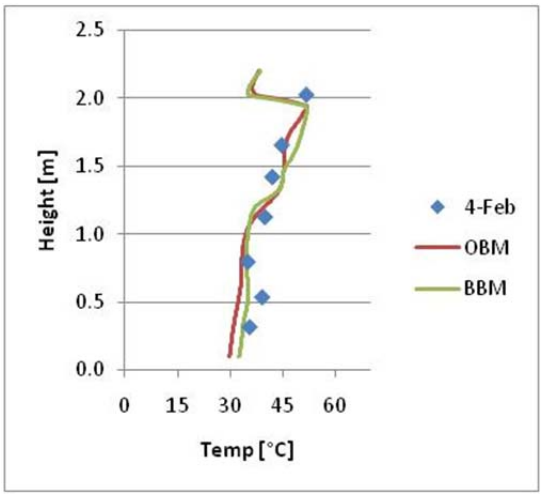

(b) Back Average

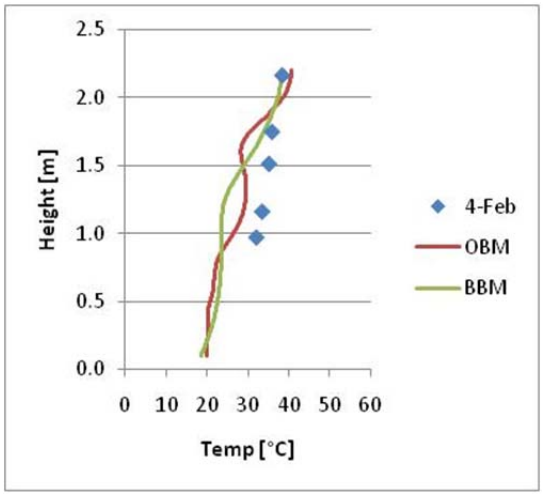

(c) Left 1

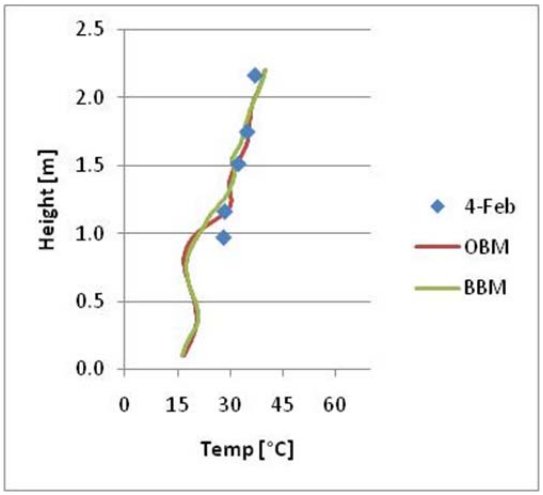

(d) Left 2

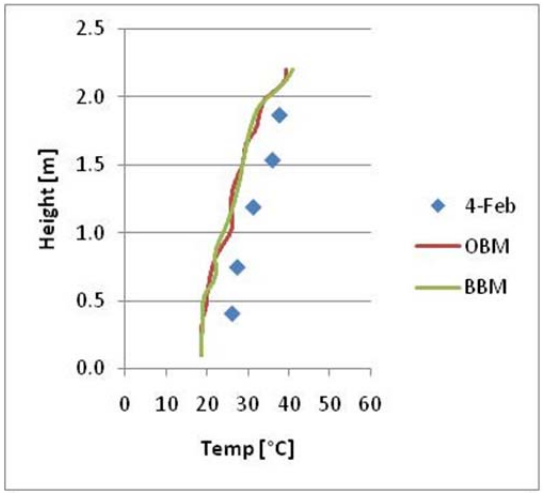

(e) Right 1

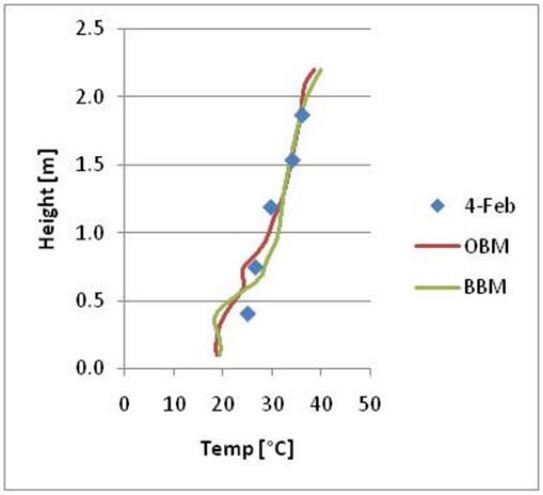

(f) Right 2

Figure 0–82: Temperature Profiles for Experiment 4

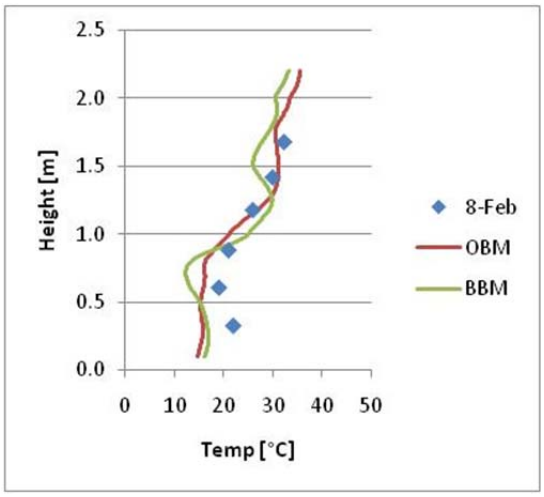

(a) Front Average

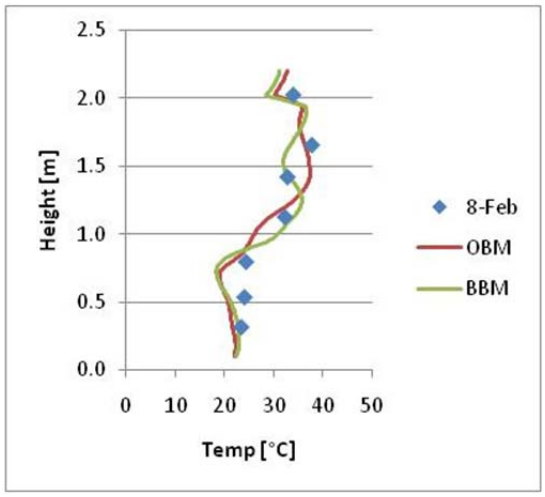

(b) Back Average

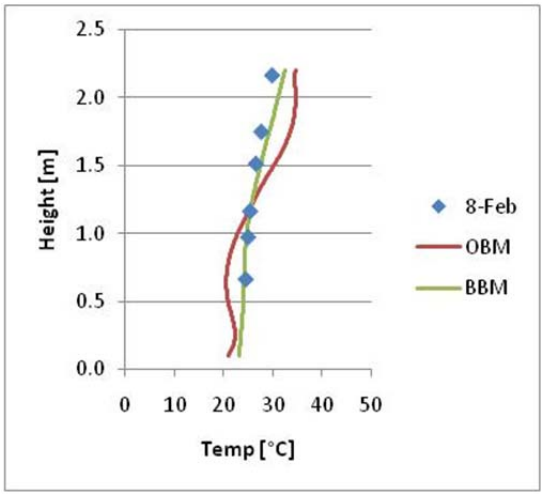

(c) Left 1

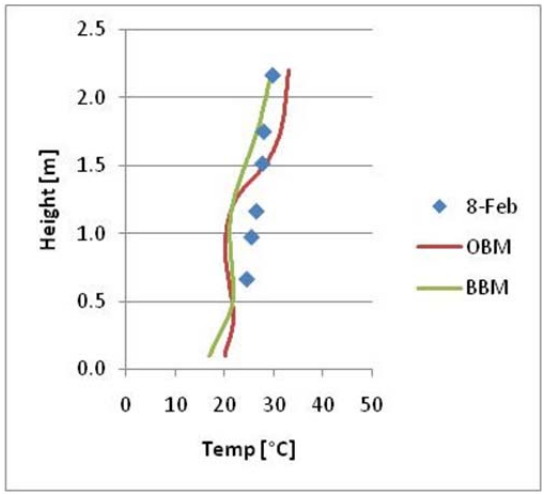

(d) Left 2

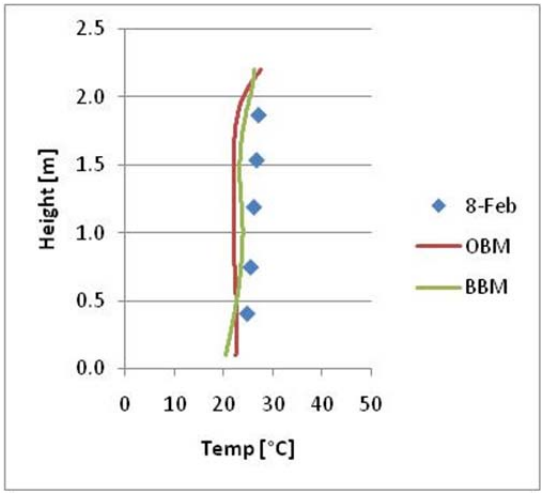

(e) Right 1

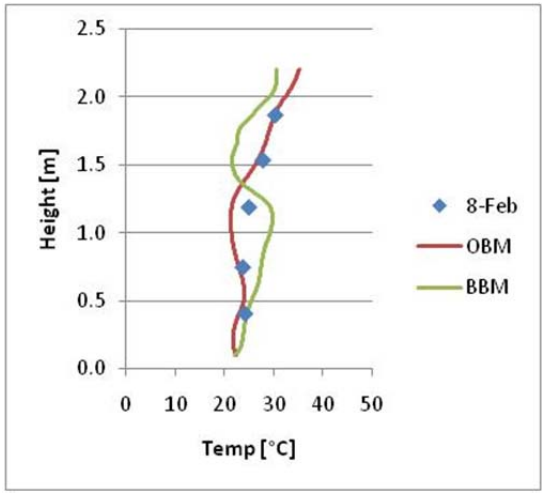

(f) Right 2

Figure 0–83: Temperature Profiles for Experiment 5

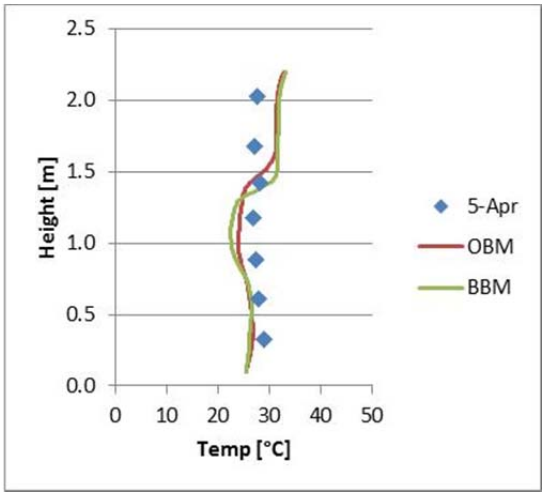

(a) Front Average

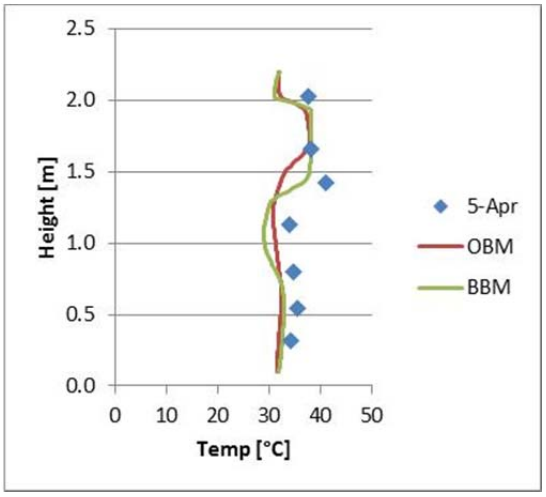

(b) Back Average

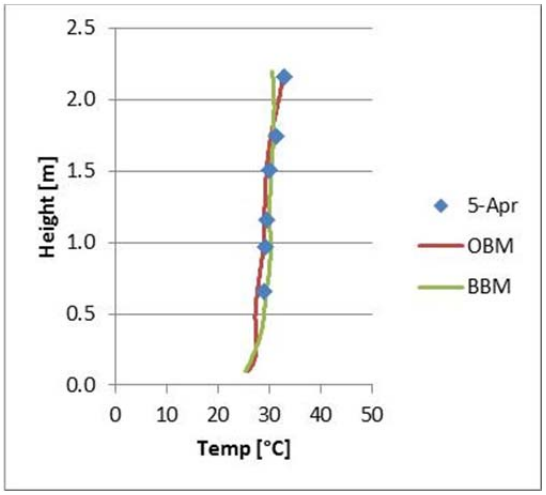

(c) Left 1

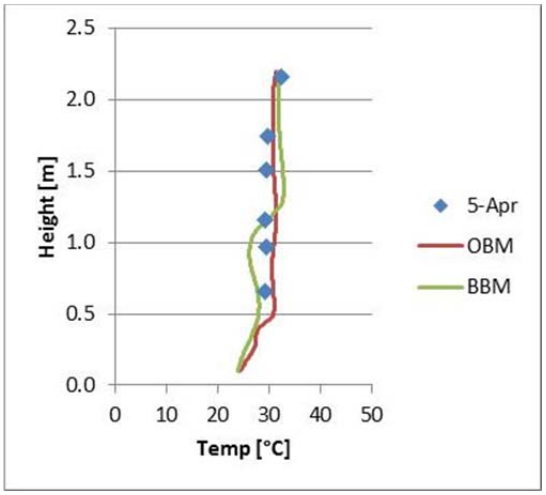

(d) Left 2

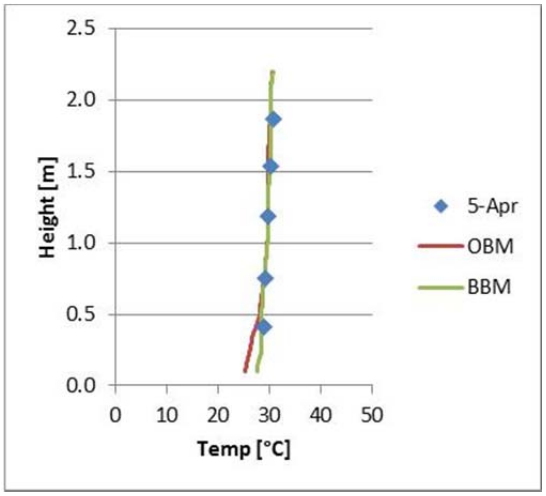

(e) Right 1

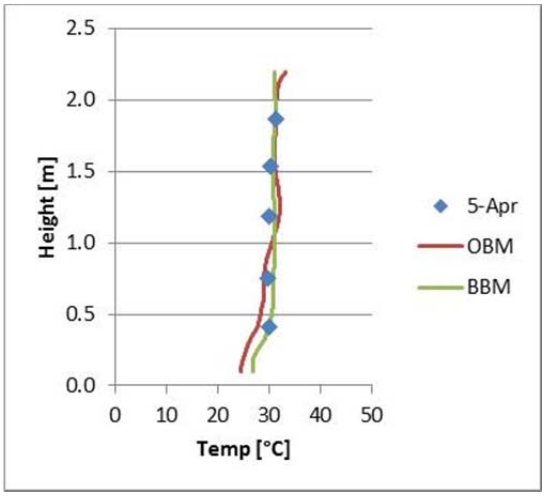

(f) Right 2

Figure 0–84: Temperature Profiles for Experiment 6

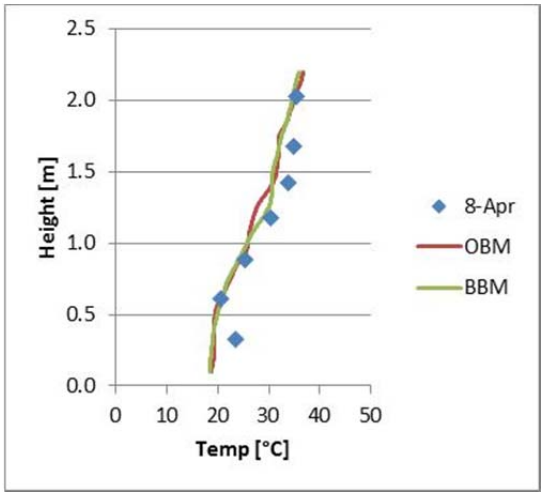

(a) Front Average

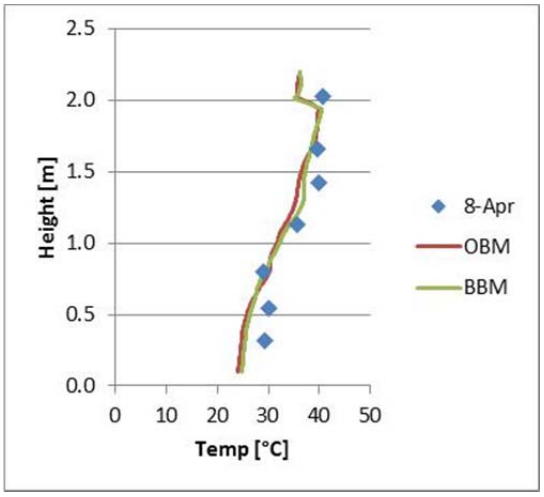

(b) Back Average

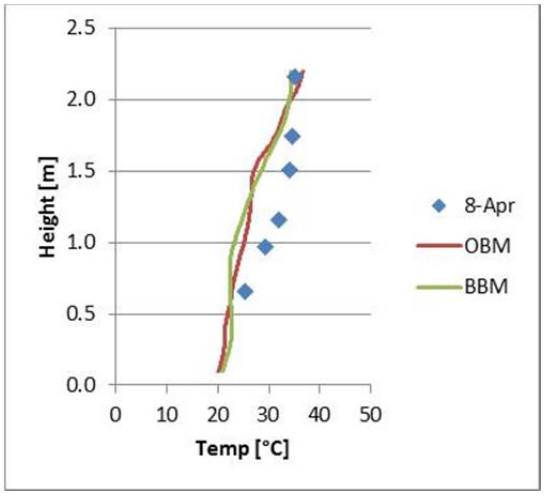

(c) Left 1

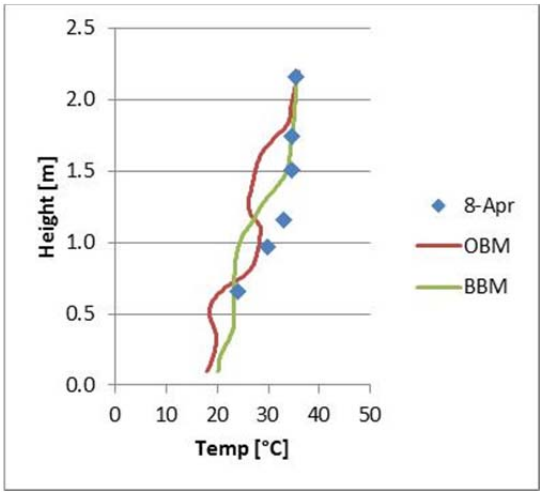

(d) Left 2

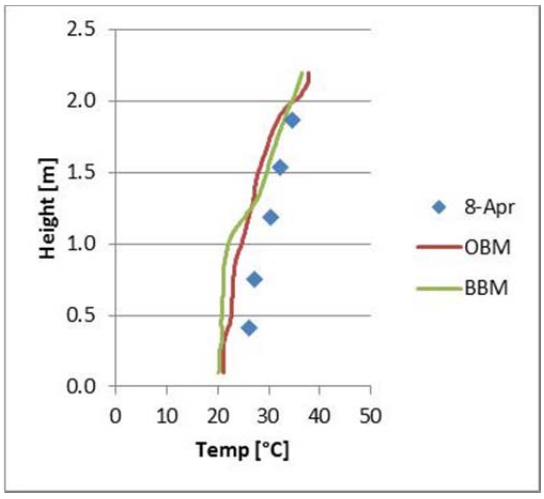

(e) Right 1

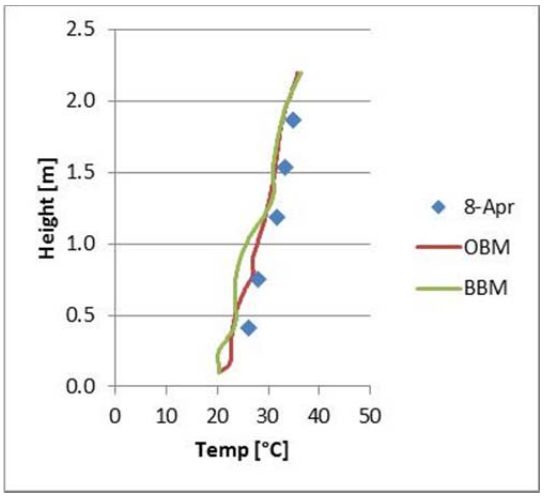

(f) Right 2

Figure 0–85: Temperature Profiles for Experiment 7

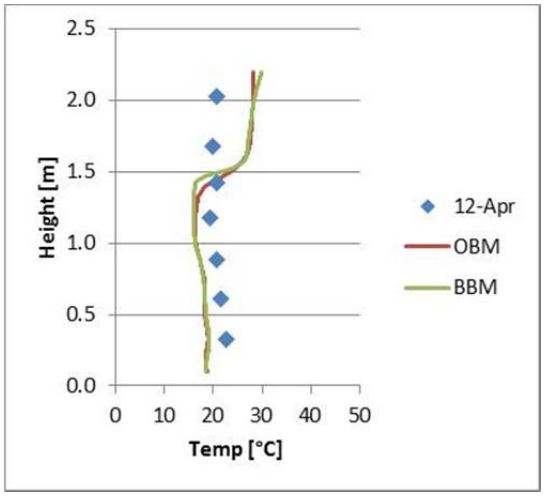

(a) Front Average

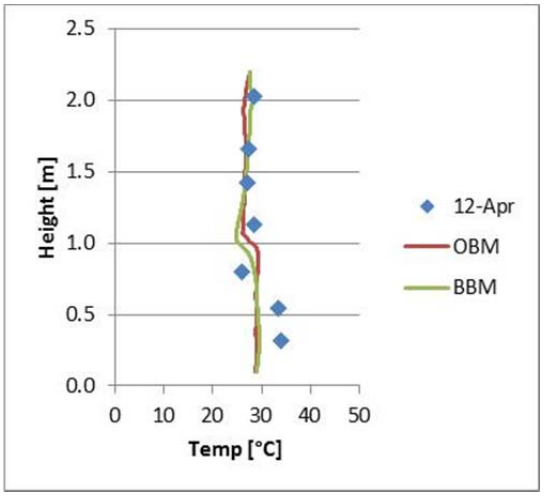

(b) Back Average

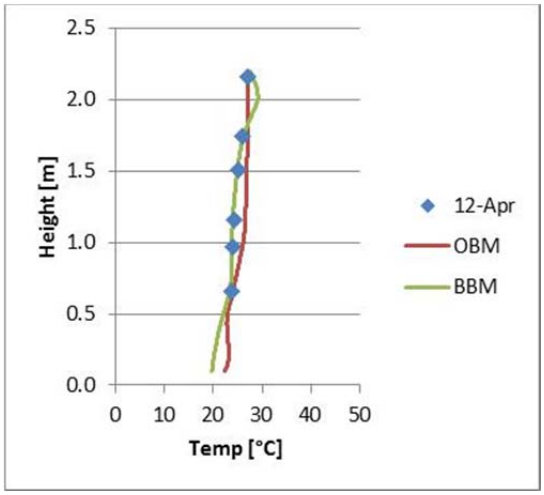

(c) Left 1

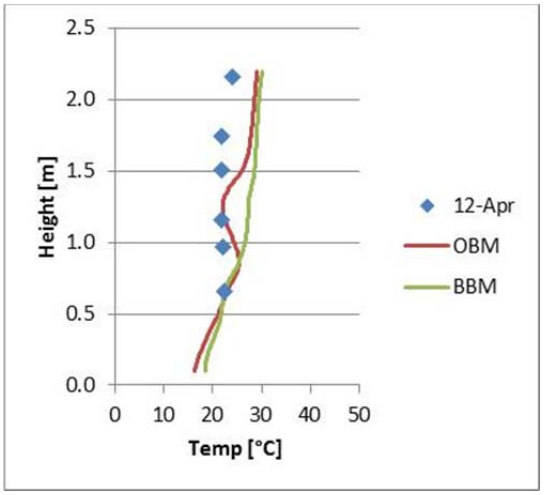

(d) Left 2

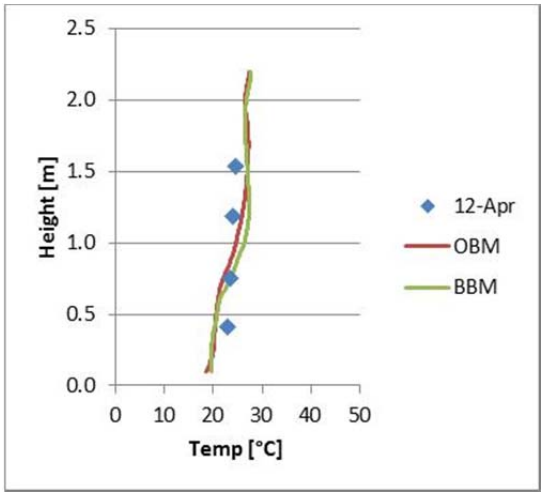

(e) Right 1

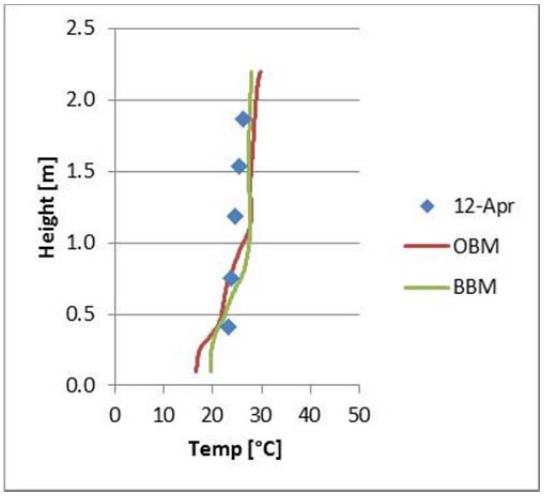

(f) Right 2

Figure 0–86: Temperature Profiles for Experiment 8

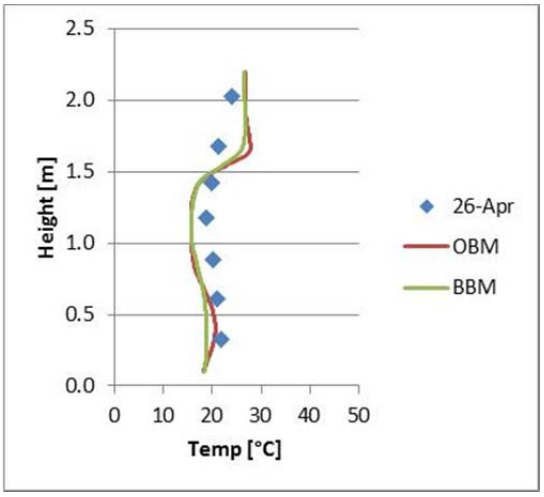

(a) Front Average

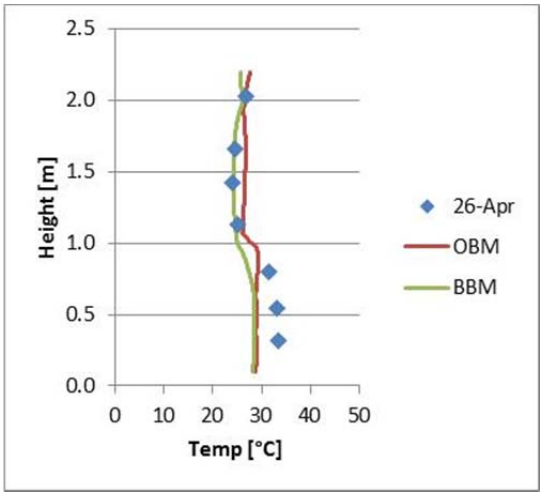

(b) Back Average

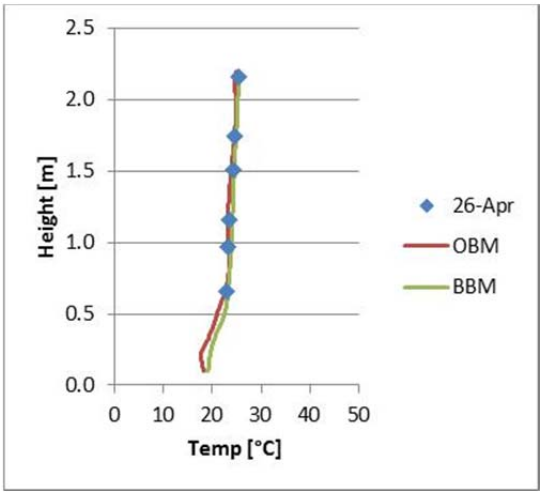

(c) Left 1

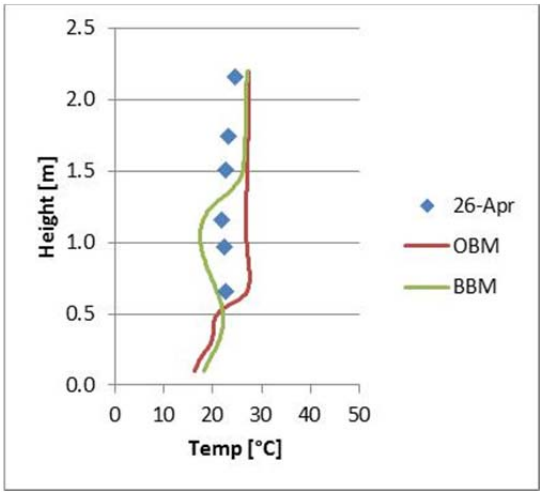

(d) Left 2

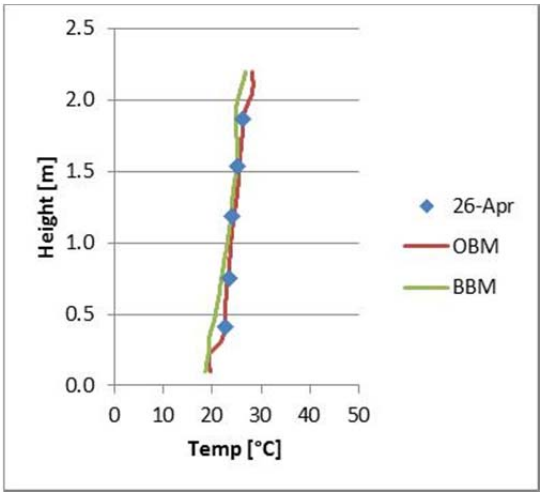

(e) Right 1

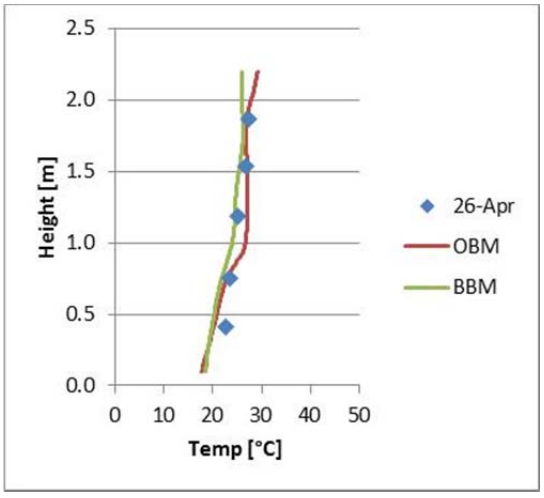

(f) Right 2

Figure 0–87: Temperature Profiles for Experiment 9

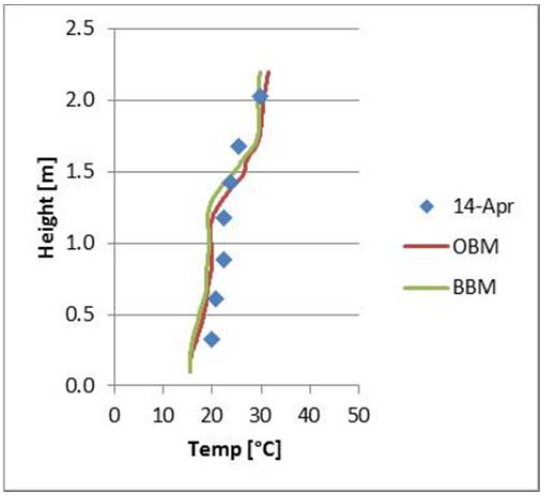

(a) Front Average

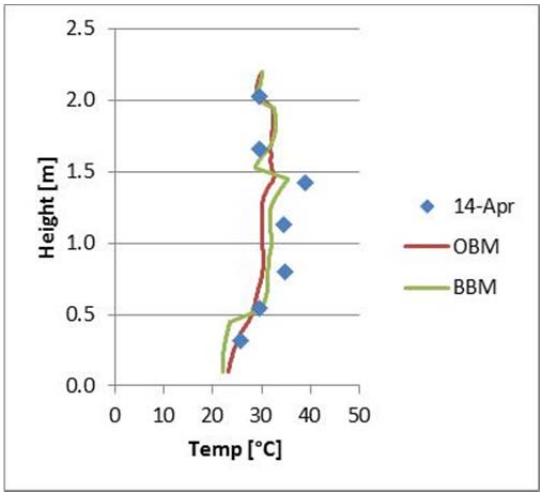

(b) Back Average

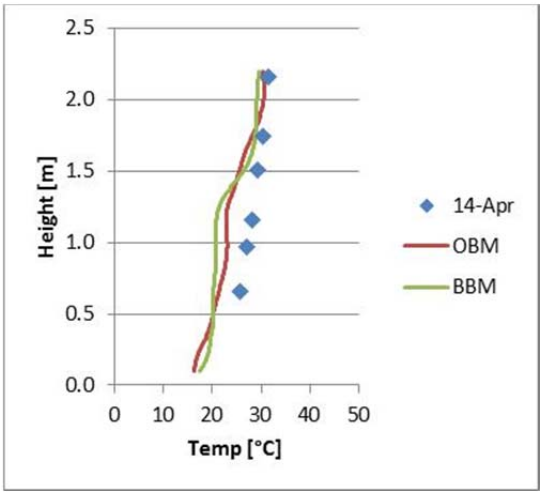

(c) Left 1

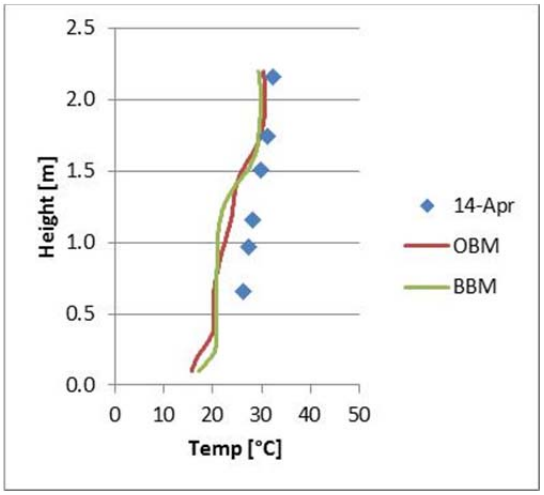

(d) Left 2

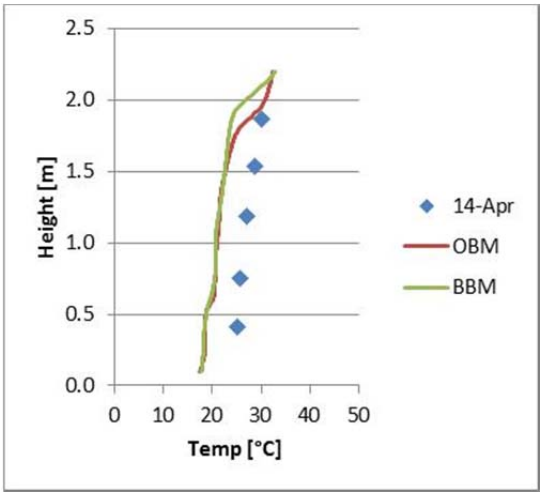

(e) Right 1

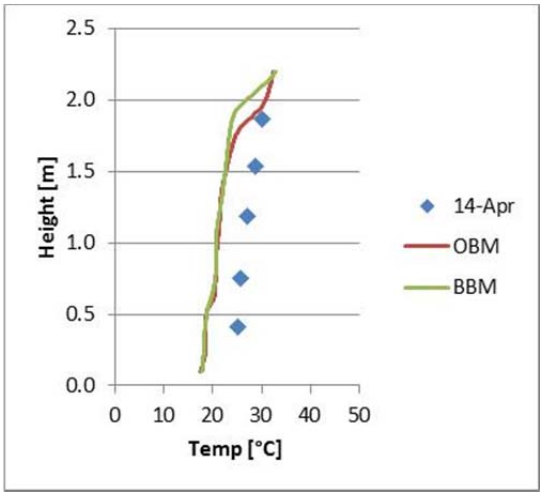

(f) Right 2

Figure 0–88: Temperature Profiles for Experiment 10