

MIT Open Access Articles

Unique Airflow Visualization Techniques for the Design and Validation of Above-Plenum Data Center CFD Models

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Lloyd, Michael, and Leon Glicksman. "Unique Airflow Visualization Techniques for the Design and Validation of Above-Plenum Data Center CFD Models." ASHRAE Transactions, Vol. 117, Issue 1 (2011). © American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

As Published: <http://connection.ebscohost.com/c/articles/67359312/unique-airflow-visualization-techniques-design-validation-above-plenum-data-center-cfd-models>

Publisher: American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE)

Persistent URL: <http://hdl.handle.net/1721.1/90482>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Unique Airflow Visualization Techniques for the Design and Validation of Above-Plenum Data Center CFD Models

Michael Lloyd

Student Member ASHRAE

Leon Glicksman, PhD

Fellow ASHRAE

ABSTRACT

One cause for the substantial amount of energy used for data center cooling is poor airflow effects such as hot-aisle to cold-aisle air recirculation. To correct these and to investigate innovative designs that will notably increase efficiency requires a robust, well-verified computational fluid dynamics (CFD) model. Most above-plenum data center CFD models are only validated using temperature data. Although a temperature-only validation method can be useful, it does not confirm that the airflow patterns predicted by the CFD model are accurate. Since the airflow patterns above a raised-floor plenum should be confidently understood before they can be optimized, it is necessary to adopt a validation method that offers more than just a comparison of temperature data.

This paper summarizes the unique validation process of a CFD model for a small data center test cell located in Cambridge, Massachusetts. The validation method features point velocity and temperature measurements and the use of small neutrally-buoyant bubbles to visualize the airflow patterns above a raised-floor plenum.

The data center test cell was designed to emulate a standard hot-aisle and cold-aisle airflow configuration. The airflow visualization revealed that each perforated tile had a substantial non-uniform air velocity distribution leading to an unexpected three-dimensional flow pattern above the racks. When this surprising reality was properly accounted for in the CFD boundary conditions, good agreement was found with the observed airflow patterns. It is the purpose of this paper to show the difficulties and value of utilizing more robust validation techniques for data center CFD models.

INTRODUCTION

One report estimated that data centers in the United States account for 1.5% of the nation's annual electricity consumption (Brown 2007). Another report claimed that data centers can consume up to 100 times the amount of energy per square foot than a typical office building (ASHRAE 2008). Energy efficiency in data centers is extremely important because of the amount of energy they use, and one way to increase the energy efficiency in a data center is to optimize its airflow patterns.

Airflow optimization is typically accomplished using a computational fluid dynamics (CFD) model. The airflow results of data center CFD models, however, are difficult to validate. Above-plenum CFD models are usually validated using only point temperature measurements. However, temperature-only CFD validation does not confirm the accuracy of the airflow patterns. Temperature measurements in the cold-aisle can reveal the presence of recirculated air from the hot-aisle but they

Michael Lloyd was a M.S. student in the Building Technology Program, Massachusetts Institute of Technology (MIT), Cambridge, MA. He graduated in June of 2010. **Leon Glicksman** is a Professor of Mechanical Engineering and Building Technology at MIT, Cambridge, MA.

will not indicate the recirculated airflow patterns or give any information about possible non-uniformities in the cold air supply flow. As a result, better validation techniques are needed to properly verify the airflow patterns of CFD models. Two validation techniques are introduced in this paper: point velocity measurements and the use of neutrally-buoyant bubbles to visualize airflow patterns.

The utility of these techniques was demonstrated by comparing experimental data obtained from a controlled test cell with computational data from a CFD model of the test cell. The process of using these techniques to validate the CFD model led to a surprising discovery in the test cell. The airflow in the test cell was clearly non-uniform although the high resistance of the perforated tiles is generally assumed to ensure uniform airflow, and the initial widely spread out point velocity measurements suggested only minor flow deviations. It is highly unlikely that the complex airflow patterns in the test cell would have been discovered using standard validation techniques.

As a result of this discovery, the detailed boundary conditions of the CFD model were adjusted to best fit what is actually happening in the space; therefore, the CFD model was improved to be more accurate and better for airflow optimization. The experiment, the CFD model and the techniques used in the validation process are presented in the following sections.

THE EXPERIMENT

A full-size, controlled experiment in an operating data center was constructed to obtain experimental results that were compared to the results of a CFD model.

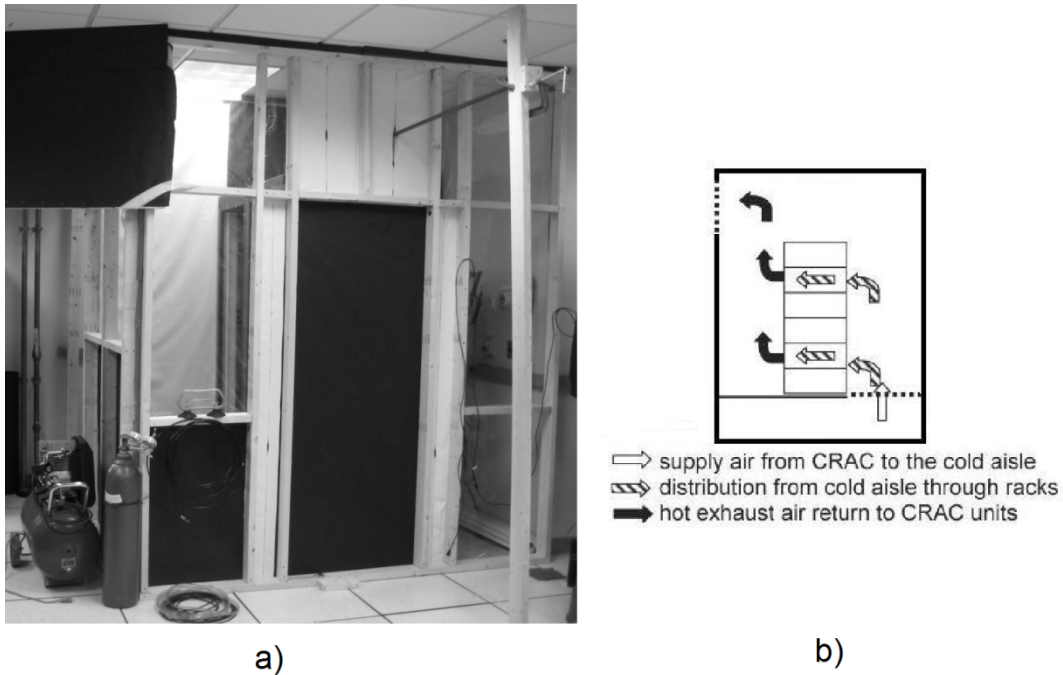


Figure 1 a) a photo showing the test cell, b) a schematic of the general airflow path

The test cell occupied a floor area of 36 square feet, and it was designed to emulate the airflow of a standard hot-aisle / cold-aisle configured data center. It was isolated from the larger room by custom partitions that went to the ceiling in order to eliminate stray air currents. The bolded lines in Figure 2 represent where the partitions were placed. These partitions were specifically designed to allow for easy and accurate measurements and observations.




Figure 2 A plan-view of the test cell

The room had a 6-inch (15.2 cm) raised-floor plenum and a ceiling height of 9.5 ft (2.9 m). The raised-floor plenum did not have any notable airflow obstructions, and the tiles were the standard size (2 ft by 2 ft ; 61 cm by 61 cm). Two perforated tiles emitted air into the experimental space from the plenum. Each perforated tile had 2,842 holes and each hole had a diameter of 0.236 inches; therefore, the open area of each tile was 21.6 percent. However, there were structural support pieces on the back of the tiles, which further reduced the actual open area to about 15 percent. The computer room air conditioner (CRAC) unit was not solely dedicated to the test cell as it supplied air for two perforated tiles located outside of the test cell.

Twenty servers were used for this experiment—ten in each rack. The racks were standard Electronics Industry Association (EIA) enclosures that were 78 inches (2.0 m) high, 24 inches (0.61 m) wide and 30 inches (0.76 m) deep. VA Linux Systems servers were used, and each server had a height of 2U or 3.5 inches (8.9 cm). A power meter was used to determine how much heat each server at steady-state operation generated. On average, each server generated about 75 W of heat, which is admittedly low per current industry practices. However, this study was not focused on the server inlet air temperatures or the amount of heat removed from the space. Rather, it was focused on better understanding the airflow patterns in order that they may be optimized to ultimately supply the servers with acceptably conditioned air. The specific server inlet air temperatures from this experiment are therefore not shared or discussed in this paper.

THE CFD MODEL

A CFD model of the test cell was created. The servers were modeled as open volumes with five features: an inlet, a face to apply a pressure drop, a face to apply a heat flux, faces to apply fan conditions and an outlet.

The server pressure drop was applied to a face inside the server to simulate the pressure drop through the component-packed server. Although most data center CFD models assume a fixed flow rate for each server, this model simulated the fans by applying a fan pressure-velocity relationship to five faces located along the back of the server. These fan faces were in series with the pressure drop face. The pressure drop versus flow rate was measured for one of the servers and the resulting curve fit to the data was used as the pressure drop boundary condition. Furthermore, to simulate the heat generated by the server components, a heat flux was applied to a face at the bottom of each server. The dimensions of the inlets and outlets of the server correspond to the actual areas measured from the servers in the test cell.

Velocity inlet boundary conditions were used to define the airflow velocity and temperature at the flow inlets. Initially, the airflow from the tiles was modeled as a uniform one-dimensional velocity. This velocity was the average of the measured readings for each tile. After the detailed flow measurements indicated significant non-uniformities, each perforated tile was

then modeled as sixteen small square velocity inlet faces. The summation of the inlet face areas in the CFD model equals the open area of the perforated tiles.




Figure 3 An isometric-view of the CFD model.

The standard k-epsilon turbulence model was used to account for the turbulent airflow. The final mesh had 1.3 million tetrahedral cells and approximately 250,000 nodes. The solution was considered converged when the solution residuals reached the default values of less than 0.001 for the flow equations and less than 10^{-7} for the energy equation. The simulations converged after about 2 hours and approximately 1,000 iterations using a computer with eight 2.66 GHz processors and 3.25 GB of RAM.

THE VALIDATION PROCESS & TECHNIQUES

A unique validation process was adopted to validate the CFD model from the experimental data. Most validation processes only compare point temperature data from the data center with the model. This process, however, offers very little to a researcher who wants to confirm that the cold-aisle airflow patterns predicted by the CFD model are accurate. As a result, two techniques were employed for the validation of this CFD model: point velocity measurements and the use of neutrally-buoyant bubbles to visualize airflow patterns.

Point Velocity Measurements

An array of point velocity measurements can be used to better understand the velocity patterns of a selected plane. They can show regions with high or low velocities in the airflow path. The point velocity measurements for this experiment were obtained using a stand which held a hot-wire anemometer with an accuracy of $\pm 3\%$ of the reading ± 0.015 m/s.

The airflow paths in the entire test cell had non-uniform airflow patterns. This was surprising as the large resistance of the perforated tiles in the airflow path hinted that the airflow should be fairly uniform; however, it was quickly discovered by taking a traverse (or array) of point velocity measurements that the airflow was non-uniform. The extent of the non-uniformity and the existence of a high speed jet-like flow were only found after the bubbles indicated there could be a high velocity region. It is highly unlikely that these patterns would have been discovered or confirmed if only temperature data or just a wide spaced velocity array were used for validation.

Figure 4 is an example of an array of point velocities above the racks, which shows the non-uniform velocity pattern.




Figure 4 The average velocity measurements for a traverse of points in a vertical plane above the racks. The bolded numbers are the average velocities in m/s (+ direction is from the back to the front of the racks) and the italicized numbers are the standard deviations.

The non-uniform airflow was not originally predicted by the CFD model using the original boundary conditions with a uniform velocity supplied above the tiles. So, after the non-uniformity discovery, the boundary conditions of the model were changed and the model then properly predicted the airflow patterns.

Neutrally-Buoyant Bubbles

Neutrally-buoyant bubbles were used to visualize the airflow patterns in the experimental space. A research-specific bubble generator was used to create the bubbles, which were capable of tracing intricate air motions.

The bubbles indicated which sections of the data center test cell should be examined in more detail with a point velocity measurement array. It was obvious after observing the bubbles in the test cell that there were some high velocity regions. Time-lapse photography of the bubbles in these regions indicated a velocity much higher than the average vertical velocity. This pattern was then confirmed using an array of point velocity measurements above the racks as seen in Figure 4.

Figure 5 is a computer visualization of an airflow pattern above the racks that was first discovered through the use of bubbles. The non-uniformity above the racks was caused by the jet-like flow in one corner of the perforated tiles.




Figure 5 A computer simulation of the airflow pattern that the bubbles showed.

The non-uniform velocities above the tiles are due, in part, to the very small height of the raised floor. The dynamic pressure of the flow in this space is of the same order of magnitude as the pressure drop through the tiles. Although this raised floor height is well below the recommended level, similar effects can be caused by obstructions within conventional raised floors.

It is recognized by the authors that the use of bubbles in an operating data center should be used with caution because the bubble film solution could potentially harm the server components.

CONCLUSION

A validation process that only utilizes temperature data is inadequate for accurately understanding the airflow patterns in a data center or for the verification of data center airflow patterns predicted by a CFD model. To have confidence in the airflow pattern predictions of a CFD model, new airflow-focused validation techniques should be utilized.

The usefulness of two airflow-focused validation techniques was introduced in this paper. Point velocity measurements and neutrally-buoyant bubbles are useful for gaining a better understanding of velocity airflow patterns in data centers.

ACKNOWLEDGMENTS

The authors would like to thank the Toda Foundation for sponsoring this research, and the MIT Laboratory for Nuclear Science for graciously allowing the use of a portion of one of their data centers for the experiment.

REFERENCES

- ASHRAE TC 9.9. *Best Practices for Datacom Facility Energy Efficiency*. Atlanta, GA: ASHRAE, 2008. ISBN 978-1-933742-27-4.
- Brown, R., Masanet, E., Nordman, B., Tschudi, W., Shehabi, A., Stanley, J., Koomey, J.,. *Report to Congress on Server and Data Center Energy Efficiency*. Berkeley, CA: Lawrence Berkeley National Laboratory, 2007. LBNL-363E.

Copyright of ASHRAE Transactions is the property of American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc. and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.