GRAVITY-ASSISTED AIR MIXING IN DATA CENTERS AND HOW IT AFFECTS THE RACK COOLING EFFECTIVENESS

Magnus K. Herrlin, Ph.D.
ANCIS Incorporated
17 Wood Street, Suite 201
San Francisco, CA 94118
USA
415-379-1947
mherrlin@ancis.us
www.ancis.us

Christian Belady, P.E.
Hewlett-Packard Company
3000 Waterview Parkway
Richardson, TX 75080
USA
972-497-4049
christian.belady@hp.com
www.hp.com

ABSTRACT

As power and heat densities continue to increase [1], the behavior of the data-center environment grows in importance, especially when developing new in-room cooling technologies. As a result, there has been a continued interest in understanding the behavior of various technologies in data centers. One issue is whether gravity plays a role in air cooled high density environments with an open architecture. A previous paper [2] suggests that technologies working with gravity perform better. This poses the question: Does gravity indeed play a key role in forced convection environments such as data centers?

The purpose of the present paper is to try to answer this question by further analyzing what we call "gravity-assisted" air mixing. Computational Fluid Dynamics (CFD) modeling in conjunction with the Rack Cooling Index (RCI) are used to demonstrate that such mixing is central to creating an adequate and "forgiving" thermal equipment environment. The paper also describes why this is an important finding for designing new high-performance cooling systems.

KEY WORDS: Central Office, CFD, Gravity, High Density, Index, RCI, Telecom, Thermal

INTRODUCTION

Thermal management of electronic equipment relies heavily on how well cool air is distributed in the data center. Over-temperatures or hot-spots may disrupt data processing equipment and cause downtime. In some industries this may mean losses that could be millions of dollars per hour.

Generally, telecom central offices use over-head air distribution whereas data centers use under-floor cooling with raised floors. Naturally, both systems have benefits and drawbacks. However, are there intrinsic cooling effectiveness differences between the systems? In other words, for well-designed systems, is cool air best supplied from the "top-down" or from the "bottom-up"? Furthermore, are the systems' performances equally sensitive to deviations from the supply airflow design point? Or, does the gravity-assisted mixing contribute to less sensitivity and better cooling performance over time?

Providing a generalized answer to this question is outside the scope of this paper. The goal here is more modest and, thus, introduces the importance of gravity in data centers with an open architecture (no physical separation of hot and cold air) by showing some examples as well as outlining an explanation to our findings. Furthermore, there is a brief discussion how these results can be applied to design high performance data center cooling solutions.

In this study, we compare a bottom-up system widely used in data centers, an alternative modular top-down system, and a hybrid system using Computational Fluid Dynamics (CFD) modeling, which has been shown to be an effective tool for analyzing data center environments [3]. In addition, the Rack Cooling Index (RCI) [2] is used to measure how effectively the equipment racks are cooled and maintained within the ASHRAE Thermal Guideline [4].

EXAMPLE 1 – REVERSE GRAVITY

To demonstrate to impact of gravity on a system seemingly dominated by forced convection, the following is a discussion centered round "reverse gravity." In this example, conventional raised floor cooling is modeled with normal gravity and then with reverse gravity. All other aspects of the two models are identical. Although this is not something that can be done in the real world and is grossly simplified, it is a simple idea that demonstrates the importance of gravity. If the results of the two models are virtually the same at the equipment air intakes, then that would indicate that gravity has little or no impact in forced convection cooled data centers. On the other hand, if the results of the models are different, then we may need to consider gravity in our design and analysis.

Figure 1 shows a cross-section along an equipment aisle in a data center with raised-floor cooling. CFD modeling allows visualization of the air temperatures. A characteristic distinct interface between hot and cold air is developed towards the top of the server racks (highlighted). One would think that this environment is dominated by the forced air from the floor tiles. However, by simply reversing the gravity field, the temperature distribution looks quite different, not only in the room at large but also in the equipment aisle. Clearly, gravity plays a significant role. This is a key finding, but it is shown under idealized conditions. How does this finding relate to more realistic set-ups? We explore this further in Example 2 where three different cooling systems are compared.

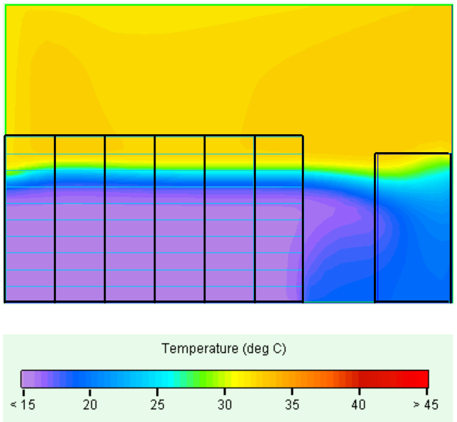

Figure 1. Normal Gravity

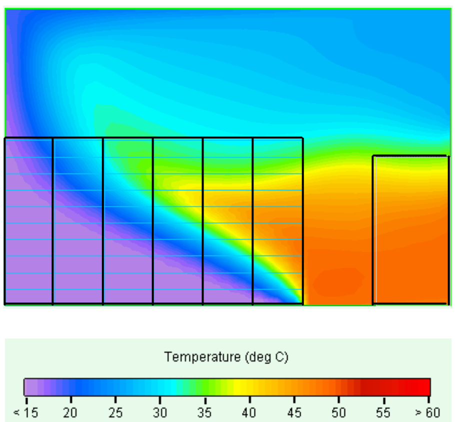

Figure 2. Reverse Gravity

EXAMPLE 2 – COMPARING TECHNOLOGIES

The bottom-up system (Figure 3) is based on a raised floor with perforated floor tiles. Air is generally supplied by down-flow Computer Room Air Conditioner (CRAC) units located along the perimeter of the equipment room. The tile airflow is uniform assuming an even pressure distribution in the under-floor plenum; the actual plenum pressure distribution was not modeled. The supply temperature is 60°F (16°C).

The top-down system (Figure 4) consists of modular cooling units installed above selected equipment racks. The number of units depends on the heat dissipation in the equipment lineup. Built-in fans and cooling coils moves and conditions the air, and a refrigeration loop controls the cooling coil temperature. The supply temperature is 65°F (18°C) to ensure dry coil operation.

The reverse-tile system (Figure 5) is essentially a reversed raised floor environment. Cool air is supplied from ceiling tiles above the cold equipment aisle. The supply temperature is identical to the bottom-up system, that is, 60°F (16°C).

The three systems are applied to a data center with hot and cold aisles. For the bottom-up system and the reverse-tile system, the entire cold aisles have perforated floor tiles for supplying cold air. The reverse-tile system and the top-down system do not require a raised floor. Each equipment rack holds ten servers (shelves), which are individually fan cooled. The equipment has front-to-rear cooling with an Equipment-Cooling Class "F-R" [5], the temperature rise across the servers is 27°F (15°C), and the heat dissipation is 3kW per rack. For all three systems, the supply airflow matches the equipment cooling airflow.

The perspective in Figures 3 through 5 is looking down a cold aisle. Clearly, the systems produce significantly different temperature conditions.

For the bottom-up system, re-circulation occurs at the top of the racks; the interface between cold and hot air is distinct (highlighted by a black bar). The top-down system, on the other hand, produces a well mixed cold aisle and the servers draw air with nearly uniform temperature. Note that only the right lineup is equipped with a cooling unit at the shown cross section. Finally, the reverse-tile system also results in fairly uniform conditions.

Although CFD modeling allows visualization of temperatures, determining the cooling effectiveness of the systems can be challenging. What matters most for the health of the air-cooled equipment is the intake temperature. The temperature in the middle of the aisle has little to do with the cooling effectiveness. The intake temperature is typically monitored by the equipment to ensure that air is provided within the specified limits for reliable and proper operations. However, these sensors are not easily available for feedback to the in-room cooling system.

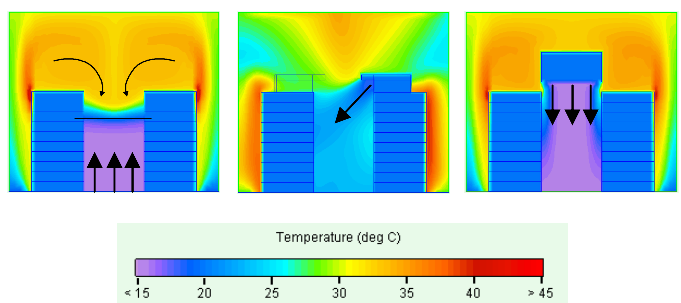

Figures 3-5. Bottom-up, Top-down, and Reverse-Tile Systems

For the bottom-up system (Figure 3), hot intake temperatures are limited to the upper servers; the matched airflow is not enough to "submerge" the top shelves with cold supply air due to loss of cold air at the end of the equipment aisles. Equipment shelves above the interface are exposed to significant over-temperatures. This temperature distribution supports the perception that equipment failures are more common for the top servers.

Although the ASHRAE thermal guideline [4] and GR-3028-CORE [5] do not recommend an equipment intake temperature of 60°F (16°C), a low supply temperature provides some over-temperature protection for the top shelves. It may also, however, introduce some relative humidity concerns for the lower shelves. In high-density environments, airflow limitations for perforated tiles may pose a challenge for achieving matched airflow rates. The issues discussed here would only be amplified had we not assumed matched airflows.

For the top-down system (Figure 4), the intake temperature distribution is relatively uniform. By supplying the cold air from the top, its high density promotes air mixing in the aisle. Indeed, this "gravity-assisted" mixing may be the most important difference between the bottom-up system and the top-down system. Although the supply temperature is higher than for the bottom-up system, the peak intake temperature is lower. Since the top-down system is modular with a certain cooling capacity per unit, the spacing of the cooling units is critical.

Finally, the reverse-tile system can be viewed as a bottom-up system turned upside-down. The temperature distribution is significantly different, however. Again, by supplying cold air from the top, its high density promotes air mixing in the aisle. Although some elevated temperatures can be expected for the top servers due to entrainment of hot air, this phenomenon could be limited by proper diffuser design.

RECOMMENDED THERMAL CONDITIONS

The thermal conditions that may occur in an equipment room are depicted in Figure 6. First, facilities should be designed and operated to target the recommended range. Second, electronic equipment should be designed to operate within the extremes of the allowable operating environment. Prolonged exposure to temperatures outside the recommended range can result in decreased equipment reliability and longevity; exposure to temperatures outside the allowable range may lead to catastrophic equipment failure.

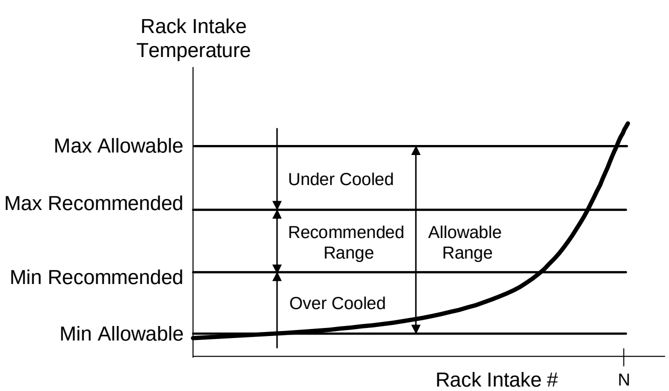

Figure 6. Temperature Distribution (hypothetical), Thermal Ranges, and Thermal Limits

The recommended range and the allowable range vary with the guideline or standard used. Generally, telecom equipment complying with the NEBS requirements [5 and 6] can withstand wider extremes than equipment designed for traditional data center environments. For the recommended temperatures, NEBS [5] suggests 65°-80°F (18°-27°C) whereas ASHRAE Thermal Guideline [4] lists 68°- 77°F (20°-25°C) for a "Class 1" environment.

THE RACK COOLING INDEX (RCI)

The Rack Cooling Index is a gauge of the thermal health of the electronic equipment; it is a measure of how effectively the racks are cooled. Specifically, the RCI_{HI} is a measure of the absence of over-temperatures; 100% means that no over-temperatures exist, and the lower the percentage, the greater probability (risk) that equipment experience excessive intake temperatures.

In other words, the RCI_{HI} is a measure of the equipment health at the high (HI) end of the temperature range. Over-temperature conditions exist when one or more equipment intake temperatures exceed the max recommended temperature.

The RCI_{HI} is defined as follows [2]:

$$RCI_{HI} = \left[1 - \frac{\sum (T_x - T_{max\text{-}rec})_{T_x > T_{max\text{-}rec}}}{(T_{max\text{-}all} - T_{max\text{-}rec})\, n}\right] 100\ \%$$

where
- T_x Mean temperature at intake x [°F or °C]
- n Total number of intakes [−]
- $T_{max\text{-}rec}$ Max recommended temperature per some guideline or standard [°F or °C]
- $T_{max\text{-}all}$ Max allowable temperature per some guideline or standard [°F or °C]

An analogous index can be defined at the low (LO) end of the temperature range [2]. The RCI_{LO} is a complement to the previously defined index especially when the supply condition is below the minimum recommended temperature. If under-temperatures are of less concern, the focus should be on maximizing the RCI_{HI}.

$$RCI_{LO} = \left[1 - \frac{\sum (T_{min\text{-}rec} - T_x)_{T_x < T_{min\text{-}rec}}}{(T_{min\text{-}rec} - T_{min\text{-}all})\, n}\right] 100\ \%$$

where
- T_x Mean temperature at intake x [°F or °C]
- n Total number of intakes [−]
- $T_{min\text{-}rec}$ Min recommended temperature per some guideline or standard [°F or °C]
- $T_{min\text{-}all}$ Min allowable temperature per some guideline or standard [°F or °C]

RCI COMPARISON

A cooling system that performs well in one environment may perform poorly in another. The overall equipment room configuration must be considered; namely, the combination of equipment layout, equipment cooling protocol, and in-room cooling system. A holistic approach is required to understand the thermal management challenges in data centers and telecom central offices.

Comprehensive CFD data generated by computer simulations are condensed by the RCI and the results are shown in Figure 7 using ASHRAE Class 1 recommended and allowable temperature ranges. The indices clearly highlight differences in the rack cooling effectiveness of the three systems.

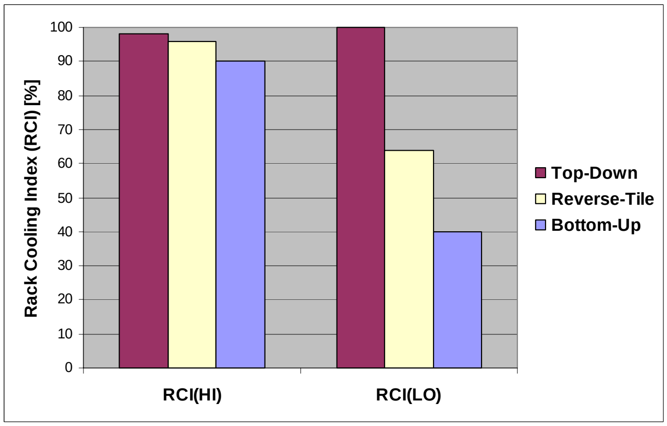

Figure 7. RCI Comparison of the Three Systems at 150 W/ft^2 (1,615 W/m^2)

The significance of a low RCI_{LO} is the potential for harmful relative humidity (RH) levels—RH is strongly correlated to temperature—and that the equipment may not be qualified at low temperatures. Timing (internal timing of data packages) may be affected and thus contribute to data corruption.

With the given design assumptions, the top-down system provides intake temperatures that match closely ASHRAE environmental Class 1 (RCIs near 100%). Contributing factors are the following:

- Air is supplied from the top down and the heavy cold air promotes well-mixed conditions in the aisle and—in turn—uniform intake temperatures
- Air is supplied in proximity of the electronic equipment, effectively avoiding entrainment of hot ambient air and loss of thermal/energy efficiency
- The supply temperature is only modestly cooler than the minimum recommended limit per ASHRAE Class 1, avoiding under-temperatures and sub-cooling of servers
- Small, flexible cooling modules allow for even distribution of the cool air, promoting even temperature conditions along the aisle.

The bottom-up system needs supply flow rates higher than the rack airflow to limit over-temperatures; matched airflows do not suffice. The reason is that some of the cold air flows out at the end of the aisles. Due to the relatively low supply temperature of 60°F (16°C) and the lack of air mixing in the aisle, the RCI_{LO} is 40%. Although the RCI_{HI} would improve with additional supply air, the RCI_{LO} would deteriorate further. The solution is to use a higher supply temperature and ensure enough supply airflow. A higher supply temperature improves the RCI_{LO} whereas an adequate airflow ensures that the hot-cold interface is above the highest equipment shelf. Another option is to put doors at the end of the cold aisles to contain the cold air. Such solutions, however, are outside the scope of this paper.

The performance of the reverse-tile system—as measured by the RCI—is in between the two other systems. This result is not completely unexpected since the system physically is a hybrid. The low RCI_{LO} of 62% is mainly due to the supply temperature of 16°C (60°F).

In-room cooling systems that are perceived as dominated by forced convection may be so only in areas nearest the air supply devices. The bottom-up system, for example, is dominated by such convection near the floor tiles. Traveling upwards, however, the air loses its momentum as the upward velocity drops in the aisle. This becomes more pronounced at the end of the equipment lineups. At a certain point, the upwards momentum is not large enough to overcome the gravity effects on the cold heavy air. Figure 3 also shows that there is a stable interface between cold and hot air. This interface, in turn, contributes to the loss of cold air at the end of the aisle. To ensure adequate equipment cooling, this interface needs to be moved above the top servers.

The top-down system, on the other hand, demonstrates the effect when the forced supply air is assisted by gravity, increasing the mixing instead of stopping it. The effect is truly significant in terms of rack cooling effectiveness—the RCI_{HI} and RCI_{LO} are both near 100% (ideal).

DISCUSSION

Gravity-assisted mixing is a combination of a longer jet throw due to gravity and natural convection due to buoyancy. For top-down systems, the cold downward air jets into the cold equipment aisles will be subjected to the benefits of gravity-assisted mixing.

The jet will initially entrain some surrounding air while traveling downwards from the diffuser. Towards the end of the jet throw, colder air is located above hotter air which results in natural (free) convection. Such convection is driven by gravity and differences in air densities, which, in turn, depend on temperature differences. The end result is a well-mixed cold aisle without a stagnant zone.

As was shown above, the resulting temperature distribution within the cold aisle looks significantly different for a typical bottom-up system. Cold air is supplied upwards from perforated floor tiles. While this plug of cold air is traveling upwards, air is drawn into the electronic equipment and some air “floats” out of the aisles at the end of the equipment rows. So, the velocity is decreasing to a point where the upward motion is zero.

In this system, the throw is reduced rather than assisted by buoyancy and no natural convection will develop due to the fact that cold heavier air is located below hotter and lighter air. Heat transfer between the cold plug and the upper hotter air is by conduction only. The developed thermal interface between cold and hot air is stable, which explains the distinct interface. Furthermore, the hot area is a stagnant zone with relatively low air motion.

Another benefit of the in-aisle air mixing characteristics of the top-down system is its low sensitivity to changes in the ratio of equipment airflow to supply airflow. This is significant since high cooling performance can be expected not only initially but also over time.

What all this demonstrates is that—in fact—gravity is a major contributor in how data center cooling solutions perform and cannot be neglected when developing new effective solutions. This conclusion is supported both by fluid dynamics theory as well as our CFD/RCI findings.

SUMMARY

This paper is comparing the equipment rack cooling effectiveness for contrasting cooling systems intended for data centers with an open architecture: a typical bottom-up system, a modular top-down system, and a hybrid system. To facilitate a standardized comparison, the Rack Cooling Index (RCI) is applied to the temperature data generated by CFD modeling. The RCI is a measure of the absence of over- and under temperatures at the air intakes of the equipment.

With the given design assumptions, the modular top-down system provides intake temperatures that match closely the ASHRAE environmental Class 1 (RCI near 100%). Strongly contributing factors are "gravity-assisted" mixing in the cold aisle, avoidance of entrainment of hot ambient air, modestly cool supply air, and a highly modular system. The system is also in-sensitive to deviations from the correct balance between equipment airflow rate and supply airflow rate.

The conventional raised-floor system requires significant air volumes to avoid re-circulation and elevated intake temperatures. Since gravity works against the velocity pressure, a stable hot-cold interface develops in the cold aisle, which, in turn, causes a substantial loss of cold air at the end of the aisle. A higher than typical supply temperature and supply air volume would improve the conditions. The hybrid system has a cooling effectiveness in between the two previous systems.

These observations are pointing to some significant differences in rack cooling effectiveness when gravity-assisted mixing is allowed to thrive. This leads to our hypothesis that there are intrinsic performance differences between over-head cooling and under-floor cooling. The authors encourage others to focus on this important area; today's high density data centers need the best possible in-room cooling solutions.

REFERENCES

1. ASHRAE Special Publication, Datacom Equipment Power Trends and Cooling Applications, American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., Atlanta, GA, 2005.

2. M. K. Herrlin, "Rack Cooling Effectiveness in Data Centers and Telecom Central Offices: The Rack Cooling Index (RCI)," ASHRAE Transactions, Vol. 111, Part 2, American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., Atlanta, GA, 2005.

3. C. Patel et al., "Computational Fluid Dynamics Modelling of High Compute Density Data Centers to Assure System Inlet Air Specification," Proceedings of International Electronic Packaging Technical Conference and Exhibition. pp 821-829, 2001.

4. ASHRAE Special Publication, *Thermal Guidelines for Data Processing Environments*, American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., Atlanta, GA, 2004.

5. Telcordia Generic Requirements, *NEBS GR-3028-CORE Thermal Management in Telecommunications Central Offices*, Issue 1, December 2001, Telcordia Technologies Inc., Piscataway, NJ, 2001.

6. Telcordia Generic Requirements, *NEBS GR-63-CORE, NEBS Requirements: Physical Protection*, Issue 2, April 2002, Telcordia Technologies Inc., Piscataway, NJ, 2002.