

DE-05-11-2

Rack Cooling Effectiveness in Data Centers and Telecom Central Offices: The Rack Cooling Index (RCI)TM

Magnus K. Herrlin, Ph.D.
Member ASHRAE

ABSTRACT

Managing today's dense equipment environments is a challenge; over-temperatures may not only harm expensive electronic equipment but also interrupt critical and/or revenue generating services. This paper is proposing a methodology and an index for analyzing the rack cooling effectiveness in data centers, telecom central offices, and other mission critical facilities.

Although Computational Fluid Dynamics (CFD) modeling allows visualization of temperatures throughout equipment rooms, to sort out the cooling effectiveness of different design options can be challenging. By using the same CFD technology, however, the rack intake temperatures can be established. They provide a "complete" picture of the thermal environment for air-cooled equipment. Still, there is no yardstick for interpreting the modeled (or measured) temperatures. This paper proposes a dimensionless index as the vehicle for such a common standard.

The Rack Cooling Index (RCI) is designed to be a measure of how effectively equipment racks are cooled and maintained within industry thermal guidelines and standards. By testing the methodology and the index on a typical data center environment with two contrasting cooling systems, it is shown that the index indeed provides a meaningful measure of the rack cooling effectiveness. The index has the capacity to help evaluate the equipment room health for managing existing environments or designing new environments.

INTRODUCTION

High electronic equipment heat dissipation and equipment concentration contribute to challenges providing an adequate thermal environment in equipment rooms. Energy consumption and heat release from servers and networking devices have increased sharply over the past years. As the heat dissipation continues to soar (per unit volume), service providers and equipment manufacturers find it increasingly difficult to provide adequate equipment cooling. The concentration of data processing equipment is driven by the demand for new and improved voice, data, and video services. A completely filled equipment rack with blade servers can generate 20 kW of heat or more—a tremendous heat density in historic terms (Uptime, 2000).

Magnus K. Herrlin is President of ANCIS Incorporated, San Francisco, Calif.

The Rack Cooling Index (RCI) is a Registered Trademark of ANCIS Incorporated (www.ancis.us)

In high density equipment rooms, thermal management relies heavily on a seamless integration of equipment-cooling design and room-cooling design, which cannot be viewed as two isolated parts of the overall thermal challenge. Managing dense environments are imperative; over-temperatures may not only harm expensive electronic equipment but may also cause downtime. The cost of downtime may be more than a million dollars per hour depending on industry (META Group, 2004).

The thermal environment for air-cooled equipment is primarily defined by the air temperature at the air intakes. The number of intakes that is above (or below) a certain temperature limit could constitute a simple “index.” Other more designed indices have been introduced for thermal design and performance; for example, the Supply Heat Index (SHI) and the Return Heat Index (RHI) (Sharma, et al, 2002). These dimensionless indices not only provide a tool to understand convective heat transfer in equipment rooms but also suggest means to improve the energy efficiency. Energy consumption can be impacted by inadequate cooling systems and equipment room configurations that allow hot and cold air to mix.

Although Computational Fluid Dynamics (CFD) modeling allows visualization of temperatures throughout equipment rooms, to sort out the cooling effectiveness of different design options can be problematic. By using CFD technology, however, the rack intake temperatures can accurately be established. Although these temperatures provide a rather complete picture of the thermal environment for air-cooled equipment, there is no existing criterion for interpreting the results.

In this paper, the dimensionless Rack Cooling Index (RCI)—or rather two of its siblings—is proposed as the vehicle for such a standard. These Rack Cooling Indices are designed to be measures of how effectively racks are cooled and maintained within industry thermal guidelines and standards. The indices help evaluate (rate) the equipment room health for managing existing environments or designing new environments, which could include short-term solutions in mature installations, transition to high average heat loads, and long-term solutions in new installations.

A thermal management rating can be compared with other performance “snap-shots” such as efficiency rating of power supplies or power rating of automobile engines. Such ratings are handy for evaluating key performance characteristics of the products. The overall methodology (and the outline of this paper) for developing and using the RCIs in data centers, telecom central offices, and other mission critical facilities is as follows:

1. Use Fluid Dynamics Modeling (CFD) for estimating temperature conditions at the equipment rack air intakes
2. Suggest a rack cooling index that will quickly and accurately compare different thermal environments and reveal trends
3. Tie the rack cooling index to leading industry thermal guidelines (e.g., ASHRAE (2004)) and standards (e.g., NEBS (2002, 2001))
4. Apply (test) the methodology on two equipment environments.

CFD MODELING

Computational Fluid Dynamics (CFD) modeling allows three-dimensional analyses of key environmental variables such as temperature and airflow. In this section, two realistic examples will demonstrate how CFD modeling can be used to estimate the rack intake temperatures in a conventional data center with two contrasting cooling systems: a bottom-up system widely used in data centers and an alternative modular top-down system. For both systems, the supply airflow matches the equipment cooling airflow.

Cooling Systems

The bottom-up system (left half of Figure 1) is based on an access floor with perforated floor tiles. Air is generally supplied by down-flow Computer Room Air Conditioning (CRAC) units. The tile airflow is

uniform assuming an even pressure distribution in the under-floor plenum; the actual plenum airflow and pressure distribution were not modeled. The supply temperature is 60°F (16°C).

The top-down system (right half of Figure 1) consists of modular cooling units installed above selected equipment racks. The number of units depends on the heat release in the equipment lineup. Built in fans and cooling coils move and condition the air; a refrigeration loop controls the cooling coil temperature. The supply temperature is 65°F (18°C) to ensure dry operation.




Figure 1. Selected Cooling Systems
(Left: Bottom-up, Right: Top-down)

Modeled Data Center

The two cooling systems are applied in a typical data center with hot and cold aisles (Figure 2). The ceiling is 12 ft (3.7 m) above the raised floor. The hot (3 ft (0.9 m)) and cold (4 ft (1.2 m)) aisles are sized per ASHRAE (2004) recommendations. For the bottom-up system, the entire cold aisles have perforated floor tiles for supplying cold air into the space. The modeled top-down system does not have a raised floor.

For modeling purposes, each 80 in. (2.0 m) electronic equipment rack is subdivided into 10 shelves, each having six monitor points to establish the mean temperature at the air intake (ASHRAE, 2004). Each shelf is individually fan cooled; the equipment has front-to-rear cooling with an equipment-cooling class “F-R” (Telcordia, 2001). The temperature rise across the shelves is 27°F (15°C).




Figure 2. Modeled Data Center with Equipment Racks and Hot/Cold Aisles

CFD Temperature Visualization

CFD modeling allows visualization of a number of environmental parameters throughout the equipment room, including airflow and temperature. Figure 3 depicts the temperature distributions for the cross

sections in Figure 1. The perspective is looking down a cold aisle; the dark rectangles are the equipment racks with 10 shelves each. Different temperatures are represented by different colors, from purple (cold) to red (hot). It is clear that the two cooling systems result in different temperature conditions in the cold aisle.

For the bottom-up system, re-circulation occurs at the top of the racks due to insufficient supply flow rate (see superimposed arrows); the interface between cold and hot air is distinct (see superimposed horizontal bar and selected intake temperatures). Equipment shelves above the interface are exposed to significant over-temperatures. The top-down system, on the other hand, produces a well-mixed cold aisle and the servers draw air with even temperature conditions. Note that only the right lineup is equipped with a cooling unit at the shown cross section.

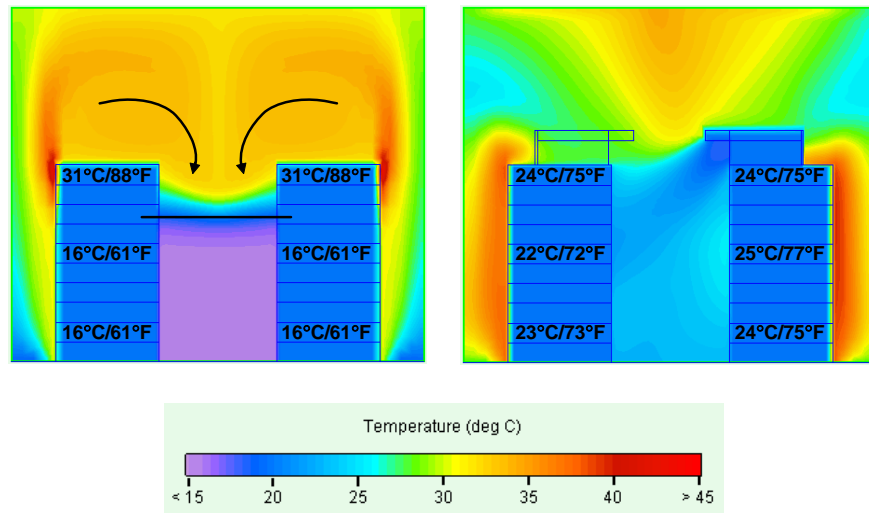


Figure 3. Temperature Distribution for Selected Cooling Systems @ 150 W/ft²
(Left: Bottom-Up, Right: Top-Down)

CFD modeling has the capacity to produce almost unlimited output data. Different parameters can be shown at different cross sections and different parameters can be animated. However, what matters most for the health of the equipment room is the equipment rack intake temperature—the temperature air-cooled equipment senses and depends on.

CFD Modeling of Intake Temperatures

Although Computational Fluid Dynamics (CFD) modeling allows visualization of temperatures across the equipment room, to sort out the cooling effectiveness of different design options can be challenging. By using the same CFD technology, however, the rack intake temperatures can accurately be established by using monitor points much the same way as taking actual measurements. Figures 4 and 5 show the resulting rack intake temperature plots for the two cooling systems at four heat densities: 75, 150, 225, and 300 W/ft². The intake temperatures are arranged in order of increasing temperature and the intakes are numbered accordingly.

For the bottom-up system (Figure 4) the temperature “steps” correspond to the upper equipment shelf levels. Obviously, the matched airflow does not manage to submerge the top shelves with the cold supply air due to loss of cold air at the end of the cold aisles. Few guidelines and standards recommend an equipment intake temperature of 60°F (16°C). However, a low supply temperature provides some over-temperature protection should the temperature interface sink below the top shelf. With adequate supply air, however, the bottom-up system generally performs well. Finally, Figure 4 shows that the temperature conditions improve at higher heat densities since the supply air is introduced with a higher velocity (the same number and type of tiles are used for all densities). Clearly, the tile selection is not optimal for the lower heat densities.

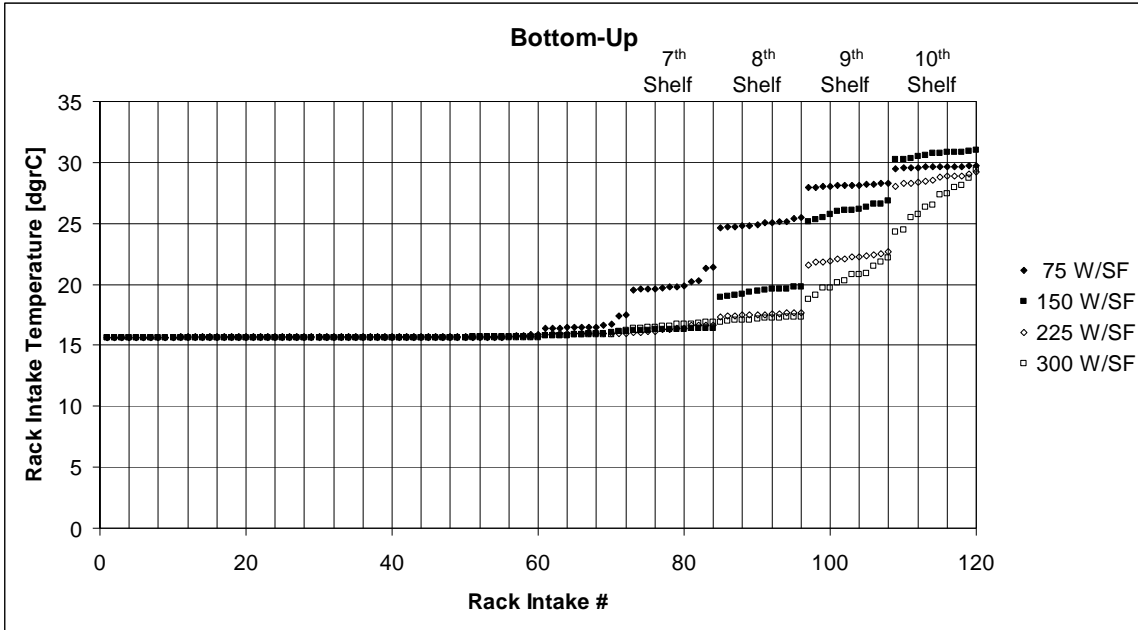


Figure 4. Rack Intake Temperature for Bottom-Up System

For the top-down system (Figure 5) the temperature distribution is relatively flat; the temperature variation from equipment intake to intake is small (Figure 3 also suggests that the aisle temperature is even). By supplying the cold air from the top, its high density promotes mixing in the aisle. Indeed, this gravity-assisted mixing may be the most important difference between the top-down system and the bottom-up system. Note that although the supply temperature is higher than for the bottom-up system, the peak temperature is lower. Since the top-down system is modular with a certain cooling capacity per unit, the spacing of the units is important; an even distribution promotes even intake temperatures. At high heat densities, therefore, the temperature distribution improves further.

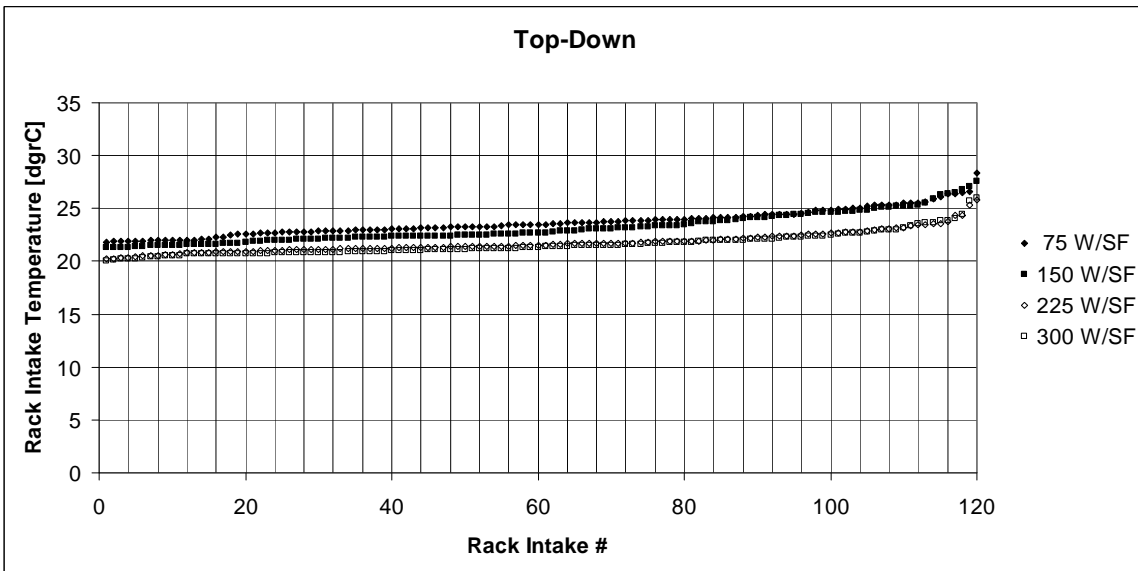


Figure 5. Rack Intake Temperature for Top-Down System

These temperature distributions provide a “complete” picture of the thermal rack environment. But, which system provides the best electronic equipment protection? A well selected index would help compare the systems based on the very same criteria.

RACK COOLING INDICES

For defining an index in general, there is a struggle between conserving details and maintaining simplicity. The overall goal is to express key characteristics in a simplified form without losing too much detail. At the outset of defining the proposed Rack Cooling Indices, the following desirable characteristics were compiled.

- Meaningful measure of rack cooling that also can be represented graphically. Overall measure of how effectively racks are cooled and maintained within leading thermal guidelines (e.g., ASHRAE) and standards (e.g., NEBS).
- Focus on over-temperature, the temperature above the max recommended temperature. Help evaluate equipment room health for managing existing environments or designing new environments by avoiding temperature-related equipment failures.
- Easily understood numerical scale. The indices should have an easily understood scale from 0% to 100%, where 100% means that all racks are cooled per some guideline or standard
- Over-cooling of some racks should not compensate for under-cooling of others—the indices should stay unchanged; no “credit” should be given to intake temperatures below the min recommended
- Indication of harmful conditions. Indication of potentially harmful thermal conditions should be provided in a simple but meaningful manner
- The indices should be suited for either computer modeling or direct measurements for establishing the input data needed for determining the indices
- Portable between platforms and non-dimensional. The indices should be independent of unit of measurement (SI/I-P units). Also, they should work with any guideline or standard that specifies max/min recommended and max/min allowable temperatures (see definitions below). In this way, the indices become a relative measure to the guideline or standard used.

Overall Thermal Conditions

The thermal conditions that may occur in an equipment room are shown 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. Per ASHRAE (2004), 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 failures.

The recommended range and the allowable range vary depending on which guideline or standard is utilized. Generally, telecom equipment complying with the NEBS requirements (Telcordia, 2001 and 2002) can withstand wider extremes than equipment designed for traditional data centers. For the recommended temperatures, NEBS (Telcordia, 2001) suggests 65°- 80°F (18°-27°C) whereas ASHRAE Thermal Guideline (ASHRAE, 2004) lists 68°- 77°F (20°-25°C) for environmental Class 1.




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

Over- and Under Temperature Conditions

Over-temperature conditions exist when one or more intake temperatures exceed the max recommended temperature per some guideline or standard. Similarly, under-temperature conditions exist when one or more intake temperatures drop below the min recommended temperature. The size—not shape—of the areas highlighted in Figure 7 is a measure of the total over-temperature and total under-temperature. These areas represent a summation of all over-temperatures and under-temperatures, respectively. The actual distributions, however, are of importance when temperatures exceed the max allowable temperature and/or drop below the min allowable temperature. As discussed below, an indication of such potentially harmful thermal conditions will be provided in a simple but meaningful way.

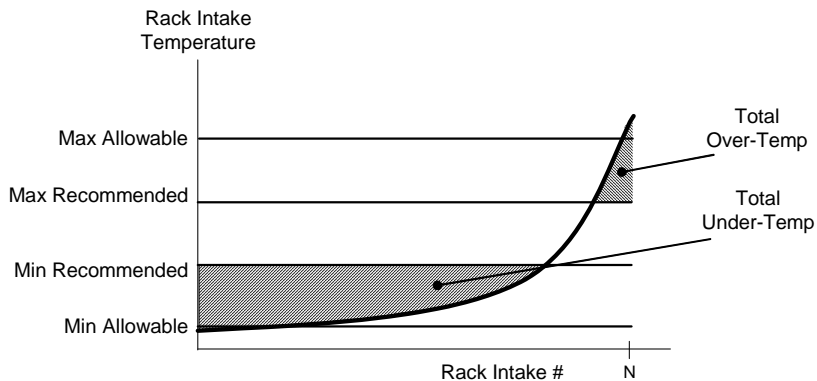


Figure 7. Definition of Total Over-Temperature and Total Under-Temperature

Max Allowed Over- and Under-Temperature

Although the total over-temperature and total under-temperature are logical measures of the overall thermal exposure, they are somewhat difficult to interpret. Thus, suitable references are introduced to make the indices easier to understand as well as provide a number of other desirable characteristics. The areas highlighted in Figure 8 define the max allowable over-temperature and max allowable under-temperature, respectively.

For a constant temperature distribution (identical equipment intake temperatures) the max allowable over-temperature results in the max allowable temperature. Although this particular distribution is uncommon (but highly desirable), the selected reference also serves other distributions well by suggesting a practical upper limit for the allowable over-temperature. An analogous argument can be made for the max allowable under-temperature.

In defining the two indices, as outlined in the next paragraphs, the max allowable over-temperature will be used with the total over-temperature whereas the max allowable under-temperature will be used with the total under-temperature.




Figure 8. Definition of Max Allowable Over-Temperature and Max Allowable Under-Temperature

Definition of RCI_{HI} and RCI_{LO}

The preceding two figures help provide a graphic representation of the Rack Cooling Indices RCI_{HI} and RCI_{LO}. The two indices are measures of the equipment room health at the high (HI) end and at the low (LO) end of the temperature range, respectively.

The RCI_{HI} definition is as follows:

$$RCI_{HI} = \left[1 - \frac{\text{Total Over-Temp}}{\text{Max Allowable Over-Temp}} \right] 100 \% \quad (1)$$

The interpretation of the index is as follows:

- RCI_{HI} = 100% All intake temperatures ≤ max recommended temperature
- RCI_{HI} < 100% At least one intake temperature > max recommended temperature

The RCI_{HI} is a measure of the absence of over-temperatures; 100% means that no over-temperatures exist. The lower the percentage (it may go negative), the greater probability (risk) that equipment experiences temperatures above the max allowable temperature. Indication of such potentially harmful thermal conditions is provided by a warning flag “*” appended to the index, generally indicating localized overheating. More work is needed to establish adequate levels of the index to meet certain design criteria.

Note that the numerical value of the index depends on the guideline or standard utilized for setting the four temperature limits per Figure 6. Even within the same guideline or standard, the numerical value may change due to temperature de-rating for elevation. Therefore, a more practical expression of the index definition is the following:

$$RCI_{HI} = \left[1 - \frac{\sum (T_x - T_{\max\text{-rec}})_{T_x > T_{\max\text{-rec}}}}{(T_{\max\text{-all}} - T_{\max\text{-rec}}) n} \right] 100 \% \quad (2)$$

- 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]

For recording the mean intake temperature (modeled or measured) at each equipment air intake, Sections 3.2 and 3.3 of ASHRAE Thermal Guideline (ASHRAE, 2004) are recommended for determining suitable monitor points. In the future, liquid-cooled systems may be introduced. Then, the RCI_{HI} is simply applied to the air-cooled equipment in the equipment space.

An analogous index can be defined for temperature conditions at the low end of the temperature range. The RCI_{LO} is a complement to the previously defined index especially when the supply condition is below the min recommended temperature. Under such circumstances the two indices can preferably be used in tandem. On the other hand, if an under-temperature is of less concern, the focus should be on maximizing the RCI_{HI}. The RCI_{LO} is defined as follows (fully parallel with the RCI_{HI} definition):

$$RCI_{LO} = \left[1 - \frac{\text{Total Under-Temp}}{\text{Max Allowable Under-Temp}} \right] 100 \% \quad (3)$$

RCI_{LO} = 100% All intake temperatures ≥ min recommended temperature
 RCI_{LO} < 100% At least one intake temperature < min recommended temperature

$$RCI_{LO} = \left[1 - \frac{\sum (T_{\min\text{-rec}} - T_x)_{T_x < T_{\min\text{-rec}}}}{(T_{\min\text{-rec}} - T_{\min\text{-all}}) n} \right] 100 \% \quad (4)$$

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]

The indices for the hypothetical data shown in Figures 7 and 8 are approximately RCI_{HI} = 95%* and RCI_{LO} = 40%*. The low RCI_{LO} value is due to the low supply temperature. Note that the “*” indicates that intake temperatures exceed the max allowable temperature and drop below the min allowable temperature, respectively. Especially the former situation may cause potentially catastrophic conditions. The next section includes an application of the indices.

APPLICATION OF THE RCIs

An application of the RCIs is performed to validate that the indices provide meaningful information that help evaluate the equipment room health. For this purpose, the comprehensive intake temperature data shown in Figures 4 and 5 are fed into Equations 2 and 4; the results are plotted in Figure 9. A close comparison and examination of Figures 4 and 5 reveal that the indices indeed condense and highlight (the hall mark of a well selected index) the differences in rack cooling effectiveness between the two cooling systems. All of a sudden, it is apparent which system that provides the healthier equipment environment.

The thermal conditions (and RCI) depend on a number of space attributes, including equipment layout, equipment cooling protocol, and space cooling system. These attributes describe the overall equipment environment. A cooling system that performs poorly in one environment may perform well in another. Therefore, the results depicted in Figure 9 are solely used here to demonstrate the usefulness of the indices, and not to demonstrate that one cooling system is universally better than the other.

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

- Air is supplied from the top down and the heavy cold air promotes well-mixed conditions in the aisle and—in turn—even 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 min recommended limit per ASHRAE Class 1, avoiding under-temperatures and under-cooling of racks
- Small, flexible cooling modules allow for even distribution of the cool air, promoting even temperature conditions in the aisle.




Figure 9. The Rack Cooling Indices (RCIs) for the Selected Cooling Systems With Environmental Class 1 (ASHRAE, 2004) as the Guideline (T-D = Top-Down, B-U = Bottom-Up)

The bottom-up system needs supply flow rates higher than the rack airflow to limit the total over-temperature—matched airflows do not suffice; the RCI_{HI} clearly helps point this out. The conditions improve with higher heat densities since the supply air is introduced with higher velocities. The longer throw limits loss of cool air at the end of the cold aisles. Due to the relatively low supply temperature of 60°F (16°C) and the lack of mixing in the aisle, the RCI_{LO} is below 50% for all heat densities studied.

Although the RCI_{HI} would improve with additional supply air, the RCI_{LO} would deteriorate further. The solution is to use a slightly 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 in the cold aisle is above the top equipment shelf (high RCI_{HI}). But, the total supply air required by the CRAC unit can be significant since more than matched airflow is needed and the raised floor leakage can be substantial. The latter is not only a design issue but also a maintenance item.

DISCUSSION AND CONCLUSIONS

This paper is proposing a methodology and two Rack Cooling Indices (RCIs) for analyzing the thermal environment in data centers, telecom central offices, and other mission critical facilities. The indices are designed to be rational measures of how effectively equipment racks are cooled and maintained within industry thermal guidelines and standards.

By using Computational Fluid Dynamics (CFD) modeling or traditional measurements, equipment intake temperatures can be recorded. Although they provide a rather complete picture of the thermal environment for air-cooled electronic equipment, they do not provide an easy and quick way of determining the overall health of the equipment room—there is no current yardstick for interpreting the results. This paper proposes the Rack Cooling Indices as a vehicle for such a common standard.

Over-temperatures exist when equipment intake temperatures exceed the max recommended temperature per some guideline or standard. Similarly, under-temperatures exist when intake temperatures drop below

the min recommended temperature. The introduced RCI_{HI} is a measure of the absence of over-temperatures; 100% means that no over-temperatures exist. This index is a gauge of the health of the equipment room at the high end of the temperature range; an analogous index (RCI_{LO}) is defined at the low end of the range.

Based on initial analyses, the methodology is easy to implement/automate and the indices indeed condense and highlight differences in rack cooling effectiveness between equipment environments or cooling systems. The indices provide meaningful information on how effectively the equipment is maintained within the recommended temperature range. They have the capacity to help manage existing environments or design new environments by facilitating standardized comparisons between various design options. Nevertheless, additional work is needed to establish adequate levels of the indices to meet certain design criteria.

The rack cooling effectiveness depends on the overall equipment room environment. A cooling system that performs less favorably in one environment may perform well in another. Consequently, a holistic approach is required to understand the thermal management challenges in data centers and telecom central offices. An important future research task is to analyze these inter-dependencies, and the Rack Cooling Indices (RCIs) have the capacity to help.

REFERENCES

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

META Group, 2004. META Group Study as reported online in DM Review at www.dmreview.com home of DM Review magazine.

Sharma, R. K., C. E. Bash, and C. D. Patel. 2002. Dimensionless Parameters for Evaluation of Thermal Design and Performance of Large-Scale Data Centers. American Institute of Aeronautics and Astronautics, AIAA-2002-3091.

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

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

Uptime, 2000. Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment, www.uptimeinstitute.org