Challenges of Data Center Thermal Management

Roger Schmidt
IBM Corporation
Poughkeepsie, N.Y.

Key Words: data centers, cooling, numerical modeling

Abstract

The insatiable desire of consumers to want more and more performance from computer equipment has driven the powers of this equipment to levels that are putting a strain on the thermal management of data centers housing this equipment. Equipment powers have been rising steadily over the past 10 years at a rapid rate. When the industry switched from bipolar to CMOS back in the early 90's industry experts had thought that the low power CMOS technology would resolve all problems associated with power and heat. Little did they know that now the problems associated with the CMOS equipment has surpassed anything installed with the bipolar technologies 10 to 15 years ago. Data centers are being designed with 15 to 20 years life spans and customers are asking how to plan for the power and cooling within these data centers. This paper addresses some of the current issues with cooling of equipment in data centers and describes some of the on-going efforts to under the thermal environment. To set the stage for describing the data center thermal management issues the power trends from the microprocessor to the rack will be described.

Introduction

Due to technology compaction, the Information Technology (IT) industry has seen a large decrease in the floor space required to achieve a constant quantity of computing and storage capability. However, the energy efficiency of the equipment has not risen at the same rate. This has resulted in a significant increase in power density and heat dissipation within the footprint of computer and telecommunications hardware. The heat dissipated in these systems is exhausted to the room and the room has to be maintained at acceptable temperatures for reliable operation of the equipment. Datacenter equipment such as those depicted in Figure 1, house several hundred, and sometimes several thousand micro-processors. The cooling of computer and telecommunications equipment rooms is thus becoming a major challenge.

Physical Design of Data Center Systems

Data centers are typically arranged into hot and cold aisles as shown in Figure 2. This arrangement accommodates most rack designs which typically employ front to back cooling and somewhat separates the cold air exiting the perforated tiles (for raised floor designs) and overhead chilled airflow (for non-raised floor designs) from the hot air exhausting from the rear of the racks. The racks are positioned such that the fronts of the racks face the cold aisle. Similarly, the rear of the racks face each other and provide a hot air exhaust region. This layout allows the chilled air to wash the fronts of the data processing (DP) equipment while the hot air from the racks exit into the hot aisle as it returns to the inlet of the air conditioning (A/C) units.

With the arrangement of DP equipment in rows within a data center there may be zones where all the equipment within that zone dissipates very high heat loads. This arrangement of equipment may be required in order to achieve the performance desired by the customer. These high performance zones can provide significant challenges in maintaining an environment within the manufacturers' specifications.

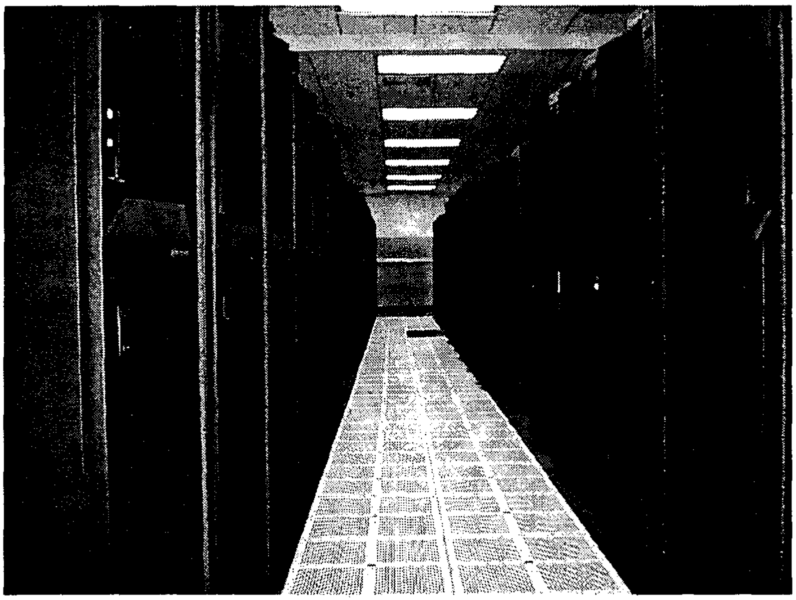

Figure 1. Photograph of data center computer equipment

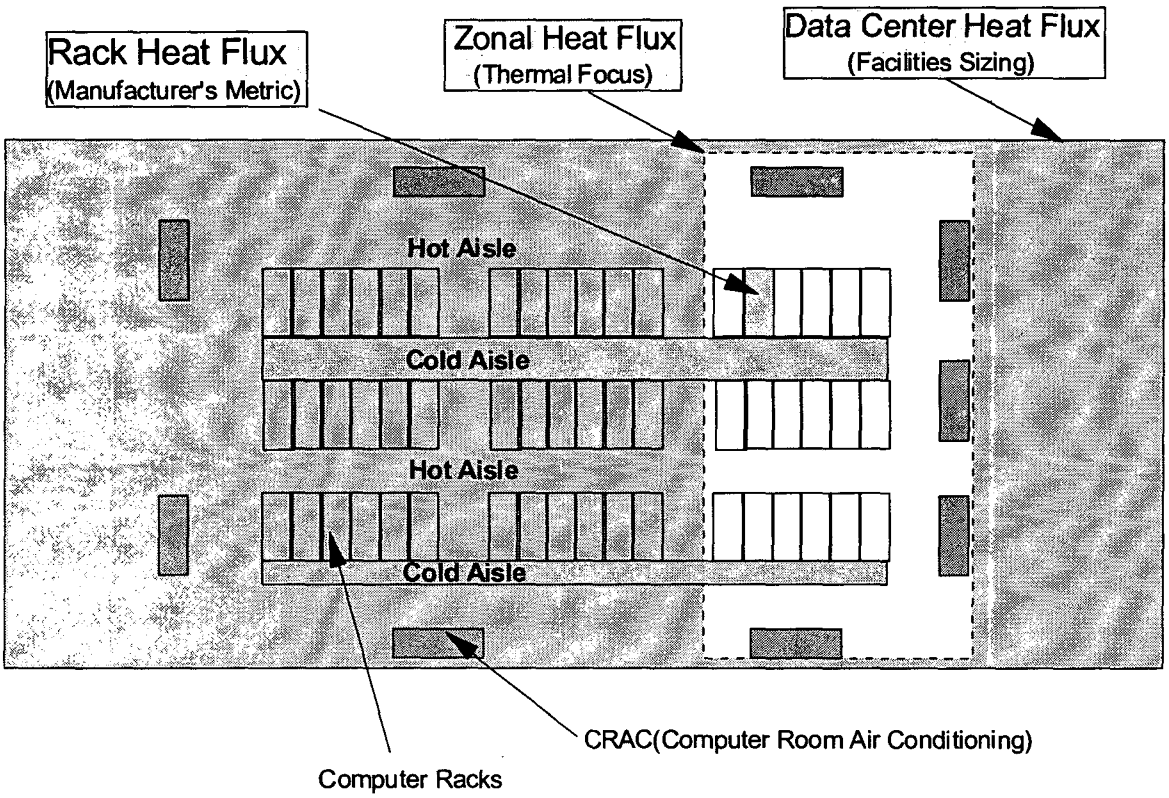

Figure 2. Organization of computer racks in a hot aisle-cold aisle layout

Air Distribution Configurations

Airflow distribution within a data center has a major impact on the thermal environment of the data processing (DP) equipment located within these rooms. A key requirement of manufacturers is that the inlet temperature and humidity to the electronic equipment be maintained within their specifications. To provide such a cool and controlled humidity environment, customers of such equipment commonly utilize two types of air distribution configurations, namely, the raised floor and non-raised floor layouts. Whether the chilled air is from a raised floor tile as shown in Figure 3(a) or via diffusers from the ceiling as shown in Figure 3(b), it is only a fraction of the rack airflow rate. This is due to the limitation of tile or diffuser flowrate. The remaining fraction of the supply side air is made up by ambient room air through recirculation. This recirculating flow is often extremely complex in nature, and can lead to significantly higher rack inlet temperatures than one might expect. There is significant hot exhaust air recirculation, which can be detrimental to the performance and reliability of the computer equipment. DP equipment is typically designed to operate for rack air inlet temperatures in the 10-35°C range. Due to recirculation, there could be a wide range in inlet air temperature across the rack. For a raised floor layout as shown in Figure 3(a), the inlet air temperature can range from 10-15°C at the bottom of the rack close to the input chilled air to as much as 30-40°C at the top end of the rack where the hot air can form a self-heating recirculation loop. Since the rack heat load will be limited by the rack inlet air temperature at the "hot" part, this temperature distribution correlates to inefficient utilization of available chilled air. Also, data center equipment almost always represents a high expenditure capital investment to the customer. Thus, it is of paramount importance from a product reliability, performance, and customer satisfaction that the temperature of the inlet air to a rack be within the desirable range.

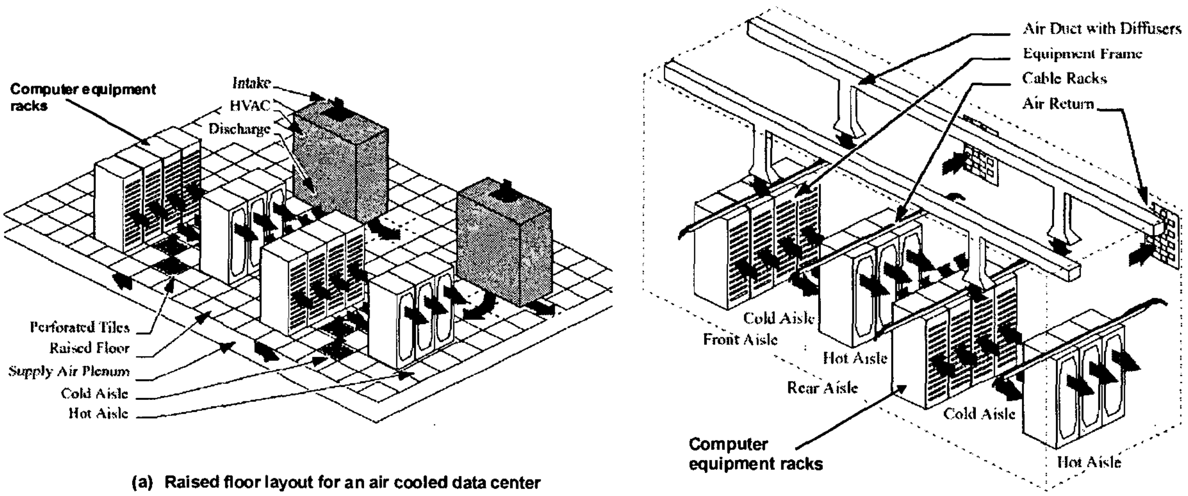

(a) Raised floor layout for an air cooled data center

(b) Non-raised floor layout for an air cooled data center

Figure 3. Configurations for Air Distribution in a data center [8], (a) raised floor, (b) non-raised floor

Thermal Profiling of High Density Data Centers

To aid in the advancement of data center thermal management, it is of utmost importance to understand the current situation in high density data centers in order to build on this understanding to further enhance the thermal environment in data centers. In this effort Schmidt [1] published the first paper of its kind to completely thermally profile a high density data center. The motivation for the paper was twofold. Firstly, the paper provided some basic information on the thermal/flow data collected from a high density data center. Secondly, it provided a methodology by which others can follow in collecting thermal and airflow data from data centers so that the information can be assimilated in order to make comparisons. This database can then provide the basis for future data center air cooling design and aid in the understanding of deploying racks of higher heat loads in the future. This data needs to be further expanded to further optimize the data center design from an air cooled viewpoint.

Data Center Numerical Modeling

Raised Floor Configuration

Schmidt [2] and Schmidt and Cruz [3-7], in a series of papers using numerical modeling techniques, studied the effects of various parameters on the inlet temperature of a raised floor data center. The numerical models were run using Flotherm, a commercially available finite control volume computer code. The Computational Fluid Dynamics (CFD) code used a k-ε turbulence model in the flow solver. The data center consisted of 40 racks of data processing units (DP) with front to back cooling which were placed in rows in a hot aisle cold aisle arrangement located in the center of the floor with four CRAC units at the perimeter of the 12.1 m wide by 13.4 m long room. The chilled cooling air provided by the CRAC units was delivered through perforated tile openings via the air plenum under the floor. Only the above floor (raised floor) flow and temperature distributions were analyzed in order to reduce the computation domain and subsequently time. By taking advantage of the symmetry in the model, only one-half of the data center needed to be modeled. The results of these papers are presented in order to provide some guidance on the design and layout of a data center.

Future Data Centers

Each data center is unique and has its own limitations with some having much higher or lower power density limits than others. To resolve these environmental issues, in some data centers today manufacturers of HVAC equipment have begun to offer liquid cooling solutions to aid in data center thermal management. The objective of these new approaches is to move the liquid cooling closer to the source of the problem, that is the electronic equipment that is producing the heat. Placing the cooling near the source of heat shortens the distance that air must be moved and minimizes the required static pressure. This increases the capacity, flexibility, efficiency, and scalability of the cooling solutions. Several viable options based on this strategy have been developed: (i) rear mounted fin and tube heat exchangers, (ii) internal fin and tube heat exchangers either at the bottom of a rack of electronic equipment or mounted to the side of a rack, and (iii) overhead fin and tube heat exchangers. Although each one of these is a liquid cooled solution adjacent to the air cooled rack, the liquid can be either water based or refrigerant based. These solutions and others will continue to be promoted as the power densities being shipped and the projected heat loads by the manufacturers of Datacom equipment continue to increase.

Summary

With the increased performance required by customers of computer equipment and the resulting heat dissipated by this equipment a significant strain is impacting the data center and its environment. Air cooling for some data centers has reached its limit whereby racks of equipment of high density loads are being divided amongst more racks in order to spread out the heat load. Of course the value proposition that racks are packed as tightly as possible for maximum performance is lost but the air conditioning capability within a data center is not allowing this to be done. These cases need to be closely examined and the understanding of the flow/temperature dynamics within a data center need to be better understood. Because of the extreme complexity of this problem and uniqueness of every data center this effort is only just beginning. Much more thermal data from data centers needs to be collected, examined, correlated and understood along with detailed fluid dynamic models to aid in this understanding. Non-dimensional metrics need to be discovered to provide guidance for data center thermal management. As more is understood of the data center dynamics then energy optimization and savings may follow. Since data centers typically expend 30 to 50% on cooling energy a large amount of energy can be saved when addressing the 1000's of data centers around the world. Simple math shows that significant savings can be achieved if only a 5% improvement can be achieved in energy related improvements.

With the increases that have occurred and will continue to occur in equipment power densities the HVAC industry as well as the server manufacturers will have to embark on liquid cooling solutions that resolve some of the temperature problems that will be more pervasive in the future if air cooling is continued. Only with the combined effort of these two industries can a consistent set of solutions that cut across all platforms be realized.

References

1. . Schmidt, "Thermal Profile of a High Density Data Center-Methodology to Thermally Characterize a Data Center-", Proceedings of the ASHRAE Nashville Conference, May (2004).

2. R. Schmidt, "Effect of Data Center Characteristics on Data Processing Equipment Inlet Temperatures", Advances in Electronic Packaging, Proceedings of IPACK '01, The Pacific Rim/ASME International Electronic Packaging Technical Conference and Exhibition, Vol. 2, Paper IPACK2001-15870, Kauai, Hawaii, July 8-13, pp. 1097-1106 (2001).

3. R. Schmidt, and E. Cruz, "Raised Floor Computer Data Center: Effect on Rack Inlet Temperatures of chilled Air Exiting both the Hot and Cold Aisles", Proceedings of the Eight Inter-Society Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), San Diego, California, June, pp. 580- 594 (2002).

4. R. Schmidt, and E. Cruz, "Raised Floor Computer Data Center: Effect on Rack Inlet Temperatures when Rack Flow rates are Reduced" Interpack Conference, Maui, Hawaii, July 6 - 13, (2003).

5. R. Schmidt, and E. Cruz, "Raised Floor Computer Data Center: Effect on Rack Inlet Temperatures when Adjacent Racks are Removed" Interpack Conference, Maui, Hawaii, July 6 - 13, (2003

6. R. Schmidt, and E. Cruz, "Raised Floor Computer Data Center: Effect on Rack Inlet Temperatures when High Powered Racks are Situated Amongst Lower Powered Racks" Proceedings of the ASME International Mechanical Engineering Congress (IMECE), New Orleans, La., November, (2002).

7. R. Schmidt, and E. Cruz, "Clusters of High Powered Racks within a Raised Floor Computer Data Center: Effect of Perforated Tile Flow Distribution on Rack Inlet Air Temperatures" Proceedings of the ASME International Mechanical Engineering Congress (IMECE), Washington, D.C., November (2003).

8. ASHRAE "Thermal Guidelines for Data Processing Environments", (2004).