Balance of Power: Dynamic Thermal Management for Internet Data Centers

Ratnesh K. Sharma, Cullen E. Bash, Chandrakant D. Patel,
Richard J. Friedrich, Jeffrey S. Chase[1]
Internet Systems and Storage Laboratory
HP Laboratories Palo Alto
HPL-2003-5
February 18th, 2003*

E-mail: ratnesh_sharma@hp.com

data center, thermal management, workload placement

The advent of Internet-based applications and their resulting multi-tier distributed architecture has changed the focus of design for large-scale Internet computing. Internet server applications execute in a horizontally scalable topology across hundreds or thousands of commodity servers in an *Internet data center*. Increasing scale and power density have a significant impact on the thermal properties of the data center. Effective thermal management is essential to the robustness of mission-critical applications. This paper shows how Internet service architectures can address multi-system resource management and thermal management at the granularity of data centers. It presents a framework for *thermal load balancing* by applying load monitoring and dynamic workload placement to manage the thermal distribution within a data center. We propose local and regional policies for thermal control and evaluate them using simulation results from a detailed computational fluid dynamics model of a typical data center with a raised floor plenum. The results demonstrate that dynamic thermal management based upon asymmetric workload placement can promote uniform temperature distribution that reduces local hot spots, quickly responds to thermal emergencies, reduces energy consumption costs, reduces initial cooling system capital costs and improves equipment reliability.

* Internal Accession Date Only Approved for External Publication
[1] Department of Computer Science, D306 Levine Science Research Center, Duke University, Durham, NC 27708-0219

Balance of Power:
Dynamic Thermal Management for Internet Data Centers

Ratnesh K. Sharma[1], Cullen E. Bash, Chandrakant D. Patel, Richard J. Friedrich
Hewlett-Packard Laboratories
1501 Page Mill Road, MS 1183
Palo Alto, California, 94304-1126

Jeffrey S. Chase
Department of Computer Science
D306 Levine Science Research Center
Duke University, Durham, NC 27708-0129

ABSTRACT

The advent of Internet-based applications and their resulting multi-tier distributed architecture has changed the focus of design for large-scale Internet computing. Internet server applications execute in a horizontally scalable topology across hundreds or thousands of commodity servers in an *Internet data center*. Increasing scale and power density have a significant impact on the thermal properties of the data center. Effective thermal management is essential to the robustness of mission-critical applications.

This paper shows how Internet service architectures can address multi-system resource management and thermal management at the granularity of data centers. It presents a framework for *thermal load balancing* by applying load monitoring and dynamic workload placement to manage the thermal distribution within a data center. We propose local and regional policies for thermal control and evaluate them using simulation results from a detailed computational fluid dynamics model of a typical data center with a raised floor plenum. . The results demonstrate that dynamic thermal management based upon asymmetric workload placement can promote uniform temperature distribution that reduces local hot spots, quickly responds to thermal emergencies, reduces energy consumption costs, reduces initial cooling system capital costs and improves equipment reliability.

Keywords: dynamic thermal management, thermal load balancing, Internet data center, row-wise and regional thermal management

[1] Please address further correspondence to this author at ratnesh_sharma@hp.com or call (650) 857-3835

1 Introduction

Thermal management is an increasingly prominent architectural consideration for high-performance computing. Thermal management challenges range across from chip, to server, rack and data center. The rise of new classes of applications, such as internet services and distributed scientific processing, have exacerbated the situation through the use of horizontal scaling techniques that result in high-density computing environments [6].

More powerful processors dissipate increasing amounts of waste heat for a given package size. Recent research demonstrates the potential of dynamic thermal management (DTM) as a means to respond to thermal conditions by adaptively adjusting the power consumption profile of a chip based on feedback from temperature sensors [5] [11]. Recent research also shows how to extend DTM to improve processor energy efficiency [7].

These techniques are key elements of thermal design for next-generation server environments, but they address only one aspect of the problem. Several factors are converging to drive a migration of computational power onto large arrays (i.e. racks) of servers aggregated in data centers, primarily for Internet-based application services. There has been tremendous growth in power dissipation in these data centers [9], and their architecture also presents difficult thermal engineering challenges stemming from increasing size and density [9]. Moreover, Internet data centers of the next decade may grow to thousands of such racks to exploit economies of scale in management, power supply, and security [6].

In this paper, we address this challenge by applying the concept of dynamic thermal management at the granularity of a complete data center rather than individual servers or chips. The key contribution of this paper is to propose policies for *workload placement* that can promote uniform distribution of temperatures through active thermal zones. We present a detailed computational fluid dynamics (CFD) model of a typical data center and use it to evaluate the effectiveness of thermal policies through fault injection simulations and dynamic variations in computational load. To our knowledge this is the first paper to deal directly with the interaction of data center architecture and software resource management functions.

2 Background and Overview

Our approach is one element of a broader framework to manage a data center as a utility whose server, router/switch and storage resources are automatically provisioned and sold according to demand, much as electricity is today. In previous research we described HP's architecture for large-scale Internet computing infrastructure on demand [10]. Instead of physically wiring each resource into a static topology, the programmable data center (PDC) treats compute, storage, and networking components as shared, virtualized resources aggregated into a pool that can be dynamically partitioned and allocated across hosted applications based on QoS requirements.. Example applications include Web-based services, computational batch schedulers, and database services for decision support, data mining, or other applications. Server and storage resources are generic and interchangeable, so the resource manager can reconfigure workload placement to respond to changing load and resource conditions [10]. Features such as service migration [1], security, and performance isolation [3] enable construction of automated data centers with dynamic and adaptive resource provisioning to ensure that service levels are maintained in the presence of changing per-application workload intensity and mix..

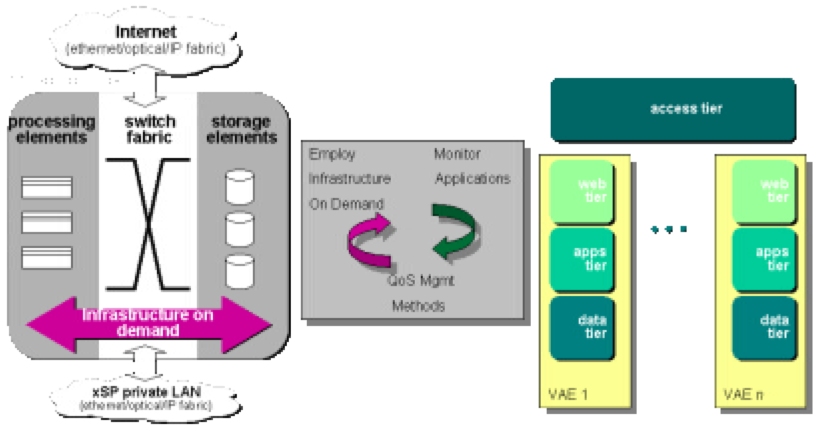

Figure 1: Programmable Data Center Architecture with QoS Management System

2.1 Thermo-Mechanical Architecture of a Programmable Data Center

PDCs are assembled in a regular, horizontally scalable topology made up of *service cores* comprising server and storage racks.. Each service core (see Figure 1) has on the order of 1000 compute nodes and a commensurate-sized storage subsystem connected via a high speed Ethernet (layer 2) switched fabric. Service cores can be aggregated to achieve vast computing and storage capabilities. The PDC architecture supports close to 50,000 rack-mounted servers, petabytes of storage, and tens of terabits per second of ingress and egress traffic.

Figure 2 depicts a typical small PDC with a standard cooling layout based on under-floor cold air distribution. The racks are arranged back-to-back and laid out in rows on a raised floor over a shared *plenum*. Modular computer room air conditioning (*CRAC*) units along the walls circulate warm air from the machine room over cooling coils, and direct the cooled air into the shared plenum where it enters the machine room through floor *vent tiles* in alternating aisles between the rows of racks. Aisles containing vent tiles are *cool aisles*; equipment in the racks is oriented so their intake draws inlet air from cool aisles. Aisles without vent tiles are *hot aisles* providing access to the exhaust air and, typically, rear panels of the equipment. There is growing interest in chilled water distribution to cool high-density data centers [8], but the simulations in this paper assume a raised floor with multiple compressor-driven CRAC units and a shared plenum.

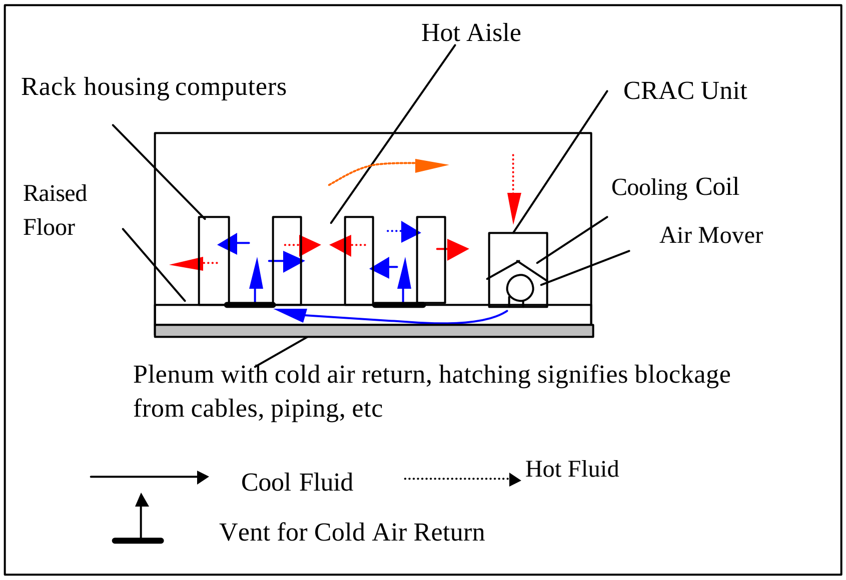

Figure 2: Typical Data Center schematic showing cooling infrastructure and EIA Racks on a raised floor

2.2 Thermal Load Balancing

Fundamentally, the effectiveness of a server thermal package depends on external environmental controls to maintain inlet air temperature within a safe operating range below the "redline" of approximately 25^{o}C. A variety of factors can cause temperature variations and thermal hot spots in the data center. Non-uniform equipment loads in the data center cause some areas to heat more than others, while irregular air flows cause some areas to cool less than others. In high heat density data centers (>2000W/m^2 (~200 W/ft^2)), mixing of hot and cold streams leads to complex airflow patterns that can create hot spots. Finally, even with a shared plenum, failure of a CRAC unit creates nonuniform cool air distribution across different regions in the plenum and machine room, resulting in high temperatures in thermal zones dominated by the failed CRAC. Thermal imbalances interfere with efficient cooling operation, and hot spots create a risk of *redlining* servers by exceeding the specified maximum inlet air temperature, damaging electronic components and causing them to fail prematurely.

Our approach is to correct thermal imbalances by incorporating thermal monitoring and workload placement policies into the PDC resource manager. The resource manager can predict thermal load by monitoring the utilization of the server and storage components, and determining the real time temperature distribution from the sensors dispersed throughout the PDC. Sections 3 and 4 use a detailed CFD thermal model to explore the formation of thermal imbalances and the effectiveness of thermal load balancing in an example data center similar to Figure 2.

2.3 Motivation and Related Work

Data center cooling designs are guided by rules of thumb and are frequently over-provisioned to provide a margin for safety, increasing both capital costs and operating costs. Recent work has investigated the use of CFD models for static provisioning of data center cooling systems and compared the results with physical measurements in a data center environment to verify accuracy of the modeling technique [8]. One study indicates that minor differences in layout can create hot spots and significantly reduce energy efficiency, and derives guidelines for effective layout [9]. Another study is aimed at the creation of fundamental metrics that can be scaled to various data center geometries and layouts. Principled layout and cooling design is a key element of a comprehensive approach to data center thermal management. However, these approaches to static design presuppose a fixed, known, balanced thermal load and stable behavior of the cooling system. In practice, thermal loading is dynamic and difficult to predict, due to a heterogeneous mix of components within the center, workload fluctuations of up to an order of magnitude [2], cooling unit failures, and incremental deployment or upgrading of components, possibly beyond the initial design parameters of the center. The resulting changes in thermal load motivate dynamic, continuous thermal load balancing and optimization of cooling resources.

Many of today's servers employ a simple form of feedback-controlled dynamic thermal management (DTM): they automatically shut down if high inlet temperature causes them to redline. Researchers are investigating more sophisticated DTM approaches to degrade gracefully by modulating power-intensive chip functions [5] [7]. In contrast, we propose to modulate temperature by migrating workload between servers to achieve a thermally balanced distribution of load across a PDC. Control-theoretic techniques can improve the precision and reduce the need for built-in margins of error in feedback-controlled DTM [11]; we expect that similar techniques would apply to our approach as well.

Our approach complements the existing work [12] on server performance by addressing efficiency of the cooling system rather than the servers themselves. In fact, these dynamic power management approaches strengthen natural variations in thermal load, further motivating dynamic thermal management.

This paper contends that software-driven dynamic thermal management is an important element of data center automation. A growing body of research uses related approaches to address other aspects of the data center automation problem [3][1][12].

3 Thermal Modeling Methodology

In this study we use a commercially available tool to create computational fluid dynamics models of a conventional data center (11.7mx8.5mx3.1m), with 0.6m deep, raised floor plenum, similar to Figure 2. The data center is populated with standard racks, each containing 20, 2U (90mm) high, Hewlett-Packard A-Class servers. Four rows, each consisting of seven racks, were arranged back to back to create cold and hot aisles. Four Computer Room Air Conditioning (CRAC) units, located along each wall, provide the cooling to the data center. The domain is discretized into 434000 cells; grid sensitivity studies verify the integrity of the numerical solution.

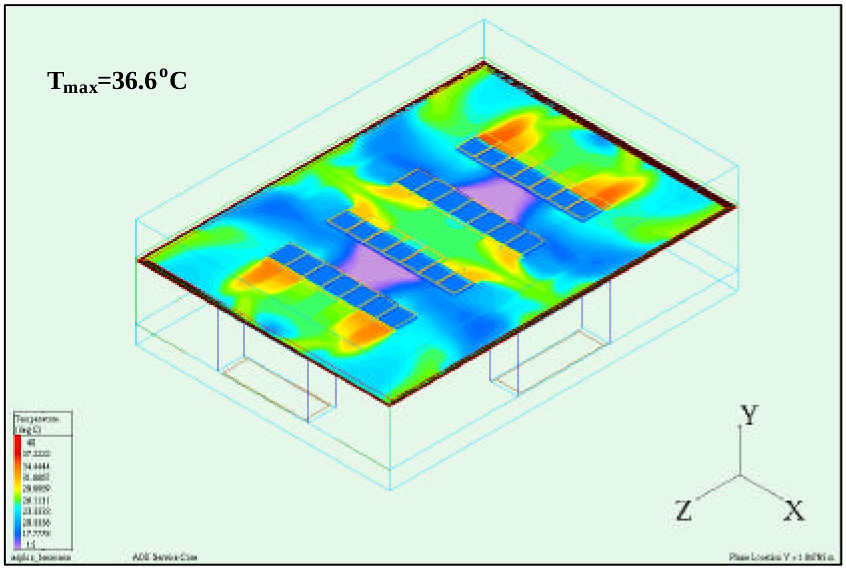

Figure 3: Temperature contour plot at 1.85m above floor for uniform workload distribution

Hewlett-Packard A-Class servers are modeled based on nameplate information. CRAC units were modeled as heat extraction devices with a characteristic outlet temperature and a limiting cooling capacity of 90kW. To evaluate the energy efficiency of the cooling system, we modeled the thermodynamic performance of the CRAC units with a vapor compression refrigeration cycle model using thermodynamic constraints and refrigerant properties. The model yields a coefficient of performance (COP) curve for each CRAC unit for applicable range of chilled air discharge temperatures. The COP of the CRAC unit is defined as the ratio of heat extracted by the CRAC unit to the electrical power consumed by the CRAC compressor.

$$COP = Q/W \tag{1}$$

where Q and W are the heat extracted by the CRAC unit and the electrical power consumed by the CRAC unit, respectively. The compressor efficiency and the condensing temperature were assumed to be constant over the evaporating temperature at constant heat load. Power consumed by air handling units was calculated from measured fan curve data at speeds specified in the CFD model.

Base case simulation results were obtained with servers uniformly loaded and dissipating heat at 75% of the rated capacity (252kW). The heat load per row was maintained at 63kW during all runs with or without workload redistribution. The discharge air from the CRAC was fixed at 15^{o}C for all units. The output of the simulation is a distribution of temperature and airflow through time within the PDC. Figure 3 shows the contour plot of temperature 1.85m above the floor in steady state. The figure illustrates that uniform workload distribution does not necessarily yield a uniform temperature distribution. While temperature is balanced near the middle of the rows, hot spots of varying magnitude and severity are apparent at the ends of the rows. Regions of moderately high temperature also exist along the hot aisles. The maximum temperature in the data center was 36.5^{o}C.

4 Thermal Load Balancing

This section demonstrates that variance in temperature leads to reduced energy efficiency, and quantifies the magnitude of the effect. It also discusses the risk of redlining as a function of thermal imbalance. From this viewpoint, the goal of thermal load balancing is to:

- Avoid cooling inlet air for each rack below the rated redline temperature of 25^{o}C., for improved energy efficiency.
- Maintain uniformity in inlet air temperature distribution to improve thermal management and reduce the risk of redlining in localized hot spots.
- Dynamically respond to thermal emergencies that result in reduced energy efficiencies and uneven temperature.

We propose two approaches to using workload redistribution to achieve these thermal management goals, and show how they contribute to improved robustness and energy efficiency of the cooling infrastructure. The *row-wise* approach focuses on thermal control locally on a row-to-row basis without affecting the entire PDC.. This is straightforward to implement and can resolve local thermal imbalances caused by uneven mixing of hot and cold air, enabling safe operation at higher average temperatures for better energy efficiency. A second *regional* approach generalizes dynamic thermal management to larger regions of the data center, and can handle larger imbalances, e.g., hot spots resulting from CRAC failure

4.1 Row-wise Thermal Management

Results from the base case simulation show that a typical deployment of computer equipment and corresponding cooling can produce large temperature variations localized within each exhaust aisle. In Figure 3, exhaust air temperatures of racks located at the end of the row are ~10^{o}C higher than that of the rack at the center of the row, even though the computational load on the servers is uniformly distributed. The complex airflow patterns that cause this imbalance are a function of thermo-mechanical architecture and are difficult to modify on a case-to-case basis.

Row-wise thermal management can resolve these imbalances by redistributing workload locally within a row based on measured temperature, equalizing temperature variation within the hot aisles. Although this concept can be implemented among servers in a rack or among rows in a PDC, we consider load distribution among racks within a row, due to ease of representation. To obtain uniform exhaust temperatures for a given total load, we propose to scale the compute load based on the excess temperature rise in each rack. Excess temperature rise of the i^{th} rack of a row is defined as the difference between its exhaust temperature (T_i) and the temperature of the cold air entering the room (T_{ref}). Higher excess temperature rise indicates a lower effectiveness in dissipating heat. Since center racks have lowest exhaust temperatures and hence, lowest excess temperature rise they are capable of dissipating the most power among all the racks in that row.

The power dissipated from i^{th} rack in the row after load redistribution (denoted by P_i) is inversely proportional to the excess temperature rise and can be represented as

$$P_i \propto \frac{1}{T_i - T_{ref}} \text{ where } i = -n \text{ to } +n \text{ for a row with } 2n+1 \text{ racks} \quad (2)$$

Since power dissipation, and consequently airflow, are equal through all the racks, it is reasonable to assume that the proportionality in Eq. (2) is identical for all racks in the row. Therefore, P_i can be represented in terms of P_0, as:

$$P_i = \left(\frac{T_i - T_{ref}}{T_{i=0} - T_{ref}} \right)^{-1} P_{i=0} = q_i P_{i=0} \quad (3)$$

where $i=0$ denotes the central rack in the row and q_i represents the thermal multiplier for the i^{th} rack. Summing up the power dissipation in a row

$$P_{row} = \sum_i P_i = \sum_{i=-n}^{+n} \left(\frac{T_i - T_{ref}}{T_{i=0} - T_{ref}} \right)^{-1} P_{i=0} \quad (4)$$

where P_{row} is the total heat dissipation from all the racks in the row. Based on P_{row}, individual rack power (P_i) can be calculated from Eq. (4). In the present case, each row had equal power dissipation of 63kW. The set of P_i, thus determined, represents the power profile for the row. To evaluate the effectiveness of row-wise redistribution, we conducted CFD simulations of the policy for the same data center and workload used in the base case. . Figure 4 shows the exhaust temperatures before and after row-wise thermal load balancing in this scenario. There is a marked reduction in hot aisle temperature variation from 10°C to 2°C. The absolute temperature values have also been reduced by several degrees. In this experiment, a reduction of maximum temperature by three degrees to 33.5°C is achieved without increasing refrigeration or air handling capacity. A significant reduction in air temperature was also obtained in other hot aisle regions of the data center. A selective implementation of the workload redistribution among the worst affected outer rows yielded an even lower maximum temperature of 32.4°C without affecting cooling demand on the CRAC unit.

Thermal load balancing can promote energy efficiency by allowing safe operation of CRAC units at a higher discharge temperature for a given computing load. Increase in CRAC unit air discharge temperature improves the COP of the refrigeration system. For example, based on our refrigeration model described earlier, a change of air discharge temperature from 15°C to 20°C, increases the COP by ~20% from 5 to 6. This results in 20% less power consumption by the refrigerant compressor for delivering an identical cooling load. This raises the average air temperature within the PDC, but the more uniform temperature distribution ensures rack inlet temperature safely below the specified redline level.

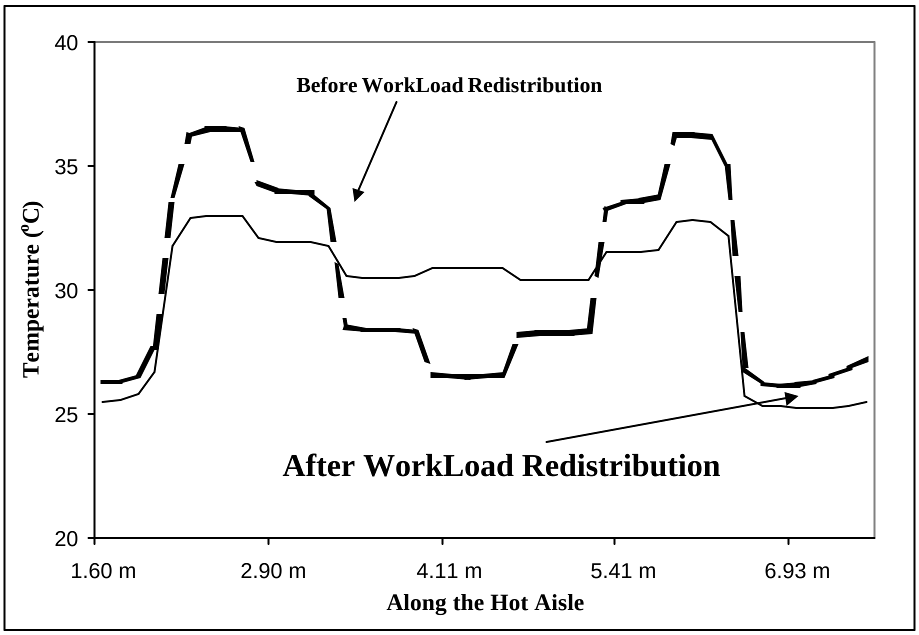

Figure 4: Rack exhaust temperatures in hot aisle before and after load balancing with row-wise thermal management approach.

A thermally optimized PDC would imply that the inlet temperature to each system would be at the redline condition of 25°C. Further if this matched the CRAC discharge temperature, implying the elimination of mixing, we would achieve an energy saving of 25% over the case shown in figure 3. This value represents the limit to which energy can be saved in the PDC and is achievable only with workload redistribution coupled with careful manipulation of the cooling infrastructure.

In this example, additional simulations indicate that row-wise thermal load balancing allows for an increase in the CRAC discharge temperature by 3°C without increasing the maximum temperature above the 36.5 °C level in the base case, yielding a cooling energy savings of 6.6 kW (9.6%). Further reductions are possible in this example without redlining. However, row-wise thermal balancing, alone, does not assure uniform temperature distribution at the server inlet. Development of an on-demand cooling infrastructure based on dynamic allocation of cooling resources can provide this feature and increase the cooling energy savings closer to the 25% limit for this PDC.

4.2 Regional Thermal Management

The regional thermal management approach depends on the ability of a workload redistribution mechanism to move large compute loads around the PDC in the event of infrastructural problems (like cooling failure, power delivery failure etc.) or a major increase in computational load. Regions can span across sections of rows with a shared chilled air supply. This technique generalizes the workload redistribution policy in the row-wise approach.

To evaluate the effectiveness of regional rebalancing in responding to a thermal emergency, we conducted fault injection experiments for a CRAC unit failure using the CFD model. As before, all the servers are working with a uniform load distribution at 75% capacity utilization dissipating 450W each when the CRAC B stops functioning. Figure 5 shows the temperature contour plot 1.2m above the floor shortly after the failure. As expected, the temperature around the unit rises, increasing the maximum temperature to 49.4^{o}C, with rack inlet temperatures above redline. Transient simulations indicate that this condition can be reached within 90 seconds of failure of CRAC unit B. The first three racks in each row closest to the unit are among the worst affected, with severe hot spots covering a significant area of the PDC.

Based on the number and arrangement of worst affected racks, we obtained a minimum region size for effective workload redistribution. For the regional rebalancing experiment, we used regions consisting of pairs of rows of three racks each. Workload was reduced in the affected region and shifted to matching regions farthest from the affected CRAC unit. Figure 5 shows the identified regions and the magnitude of the workload assigned for rebalancing.

After rebalancing, the racks closest to the unit functioned at 50% of rated power while those at the far end of the room away from the failed CRAC functioned at full capacity (100%) (as marked in Figure 5). As a result of this rebalancing, the PDC is free of high temperature zones and steep temperature gradients, barring a few minor hot spots. The maximum temperature in the room dropped by 8^{o}C. In addition, shifting of computational load away from the failed CRAC unit significantly reduces the temperature variation in the hot aisles.

To evaluate the energy savings obtained by workload distribution it is necessary to obtain identical maximum temperatures by modifying only the airflow and temperature rather than workload to avoid a thermal disaster. The airflows in the operating CRAC units are increased from 4.72 m^3/s to 7.1 m^3/s. The discharge air temperatures for CRAC units A and D are reduced from 15^{o}C to 10^{o}C.

Comparison of these PDC case results, one with cooling modification and one with workload distribution, show that there is an energy saving of 14.4% (11kW) with workload distribution. The analysis clearly shows that workload distribution is an energy efficient method for thermal management and disaster recovery of PDCs.

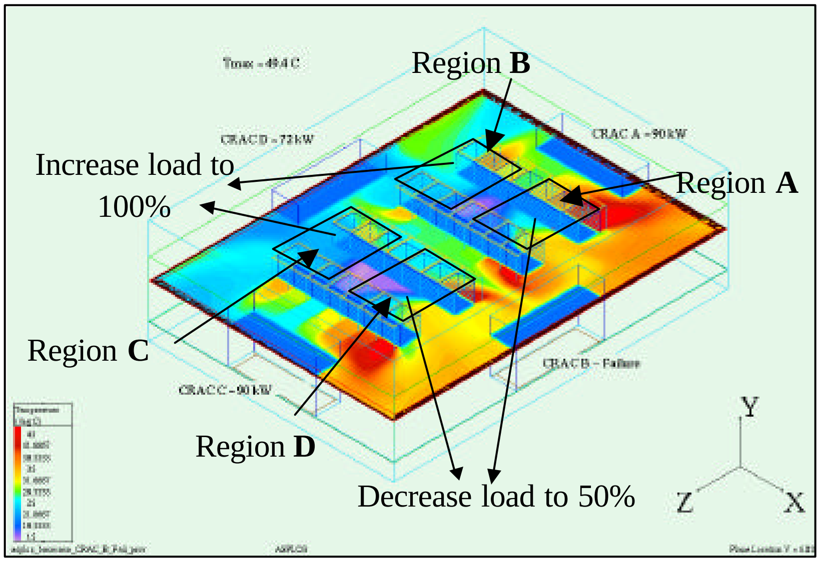

Figure 5: Temperature contour plot at 1.2m above floor for uniform workload distribution with a failed CRAC unit

5 *Implications for data center OS*

The data center OS must provide the resource manager for dynamic thermal management. Though CFD models have been used in this paper to simulate the thermo-fluid airflow behavior, our results show that simple heuristics can be derived and implemented in the data center OS to minimize the cost of computing workload placement decisions. From the results stated in this paper, we propose to use the measured server exhaust temperature, based on a uniform workload, to calculate the new load distribution. Row-wise thermal management can be implemented using the following workload placement algorithm:

- Initiate a uniform workload placement.
- Measure the exhaust temperature (T_i) for all racks
- Use the exhaust temperature (T_i) to calculate the thermal multiplier (q_i) for all racks (Eq. (3))
- Create the power profile (P_i) using Eq. (4)
- Place the workload according to the power profile, thus, yielding the optimum workload placement.
- New workloads entering the data center are placed on racks with higher thermal multipliers prior to racks with lower thermal multiplier, until all the racks are fully

utilized. The reverse course of action is taken when workload is completed and exits the data center.

Regional thermal management can be accomplished by workload placement into predefined regions outside the zone of influence of failed CRAC units. A key requirement for the data center OS is the ability to obtain *in situ* measurements of temperatures within the data center, either through server self-measurement or external measurement.

Implementation of the resource manager for dynamic thermal management assumes that workload placement can be achieved at an arbitrary granularity in data centers operating below saturation. In practice, QoS constraints, storage placement, and workload rebalancing granularity may constrain the thermal rebalancing policies. The paper has demonstrated the potential of workload placement.

6 Contributions and future work

Data center power consumption is expected to grow ten fold in the coming decade, placing great demands on data center cooling infrastructure and electricity consumption. This paper proposes a novel approach for thermal management and energy efficiency in programmable data centers by utilizing the techniques of workload placement. Two different approaches have been proposed. The approaches, based on regions and rows of data center equipment, can be implemented, separately or jointly, according to the operational conditions in the programmable data center. One of the key contributions emerging from this paper is the reaction to catastrophic thermal emergencies such as partial air conditioning failure through dynamic workload placement. Additionally, we have also demonstrated that energy consumption can be reduced by more than 14% by workload placement.

While our proposition in this paper was based on a given infrastructure problem definition of a data center, future work is currently underway to obtain geometry-independent correlations that can be implemented in the data center OS. Moreover, future research related to a combination of distributed measurements and the use of variable capacity air conditioning resources [9] is also required.

It is clear, given the increase in power density in data centers and the need for always-on infrastructures, that techniques such as dynamic thermal management in the data center OS will be required to efficiently manage large-scale programmable data centers of the future.

7 References

1. K. Appleby, S. Fakhouri, L. Fong, G. Goldszmidt, M. Kalantar and S.Krishnakumar, D. Pazel, J. Pershing and B. Rochwerger, *Oceano - SLA Based Management of a Computing Utility*, In Proceedings of the 7th IFIP/IEEE International Symposium on Integrated Network Management, May, 2001
2. M. Arlitt, D. Krishnamurthy, J. Rolia, "Characterizing the Scalability of a Large Web-based Shopping System," ACM Transactions on Internet Technology, Summer 2001.

3. Nina Bhatti and Rich Friedrich, "Web Server Support for Tiered Services" IEEE Network, 13(5):64-71, September 1999.
4. Pat Bohrer, Elmootazbellah N. Elnozahy, Tom Keller, Michael Kistler, Charles Lefurgy, and Ray Rajamony. *The Case for Power Management in Web Servers. In Power-Aware Computing* (Robert Graybill and Rami Melhem, editors). Kluwer/Plenum series in Computer Science, to appear, January 2002
5. D. Brooks and M. Martonosi, *Dynamic Thermal Management for High-Performance Microprocessors*, In Proceedings of the Seventh International Symposium on High-Performance Computer Architecture (HPCA-7), January 2001
6. Friedrich, R., Patel, C.D., *Towards planetary scale computing - technical challenges for next generation Internet computing*, THERMES 2002, January 2002, Santa Fe, New Mexico
7. M. Huang, J. Renau, S.-M. Yoo, and J. Torrellas. *A Framework for Dynamic Energy Efficiency and Temperature Management*. in Proc. of the 33rd Annual Intl. Symp. on Microarchitecture, 2000
8. Patel, C.D., Bash, C.E., Belady, C., Stahl, L., Sullivan, D., July 2001, *Computational fluid dynamics modeling of high compute density data centers to assure system inlet air specifications*, Proceedings of IPACK'01 – The PacificRim/ASME International Electronics Packaging Technical Conference and Exhibition, Kauai, Hawaii.
9. Patel, C.D., Sharma, R.K, Bash, C.E., Beitelmal, A, *Thermal Considerations in Cooling Large Scale High Compute Density Data Center*s, ITherm 2002 - Eighth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems" May 2002, San Diego, California
10. J. Rolia, S. Singhal, R. Friedrich, "Adaptive Internet Data Centers," SSGRR 2000, European Computer and eBusiness Conference, L'Aquila, Italy, July 2000.
11. K. Skadron, T. Abdelzaher and M. R. Stan, *Control-Theoretic Techniques and Thermal-RC Modeling for Accurate and Localized Dynamic Thermal Management*, In Proceedings of the Eighth International Symposium on High-Performance Computer Architecture, Feb. 2–6, 2002
12. J. Chase, D. Anderson, P. Thakar, and A. Vahdat, *Managing energy and server resources in hosting centers*, In 18th ACM Symposium on Operating System Principles (SOSP'01), pages 103--116, Chateau Lake Louise, Banff, Canada, October 2001.