DYNAMIC THERMAL RESPONSE OF THE DATA CENTER TO COOLING LOSS DURING FACILITY POWER FAILURE

A Thesis
Presented to
The Academic Faculty

by

Shawn Shields

In Partial Fulfillment
of the Requirements for the Degree
Master of Science in the
School of Mechanical Engineering

Georgia Institute of Technology
August 2009

/

DYNAMIC THERMAL RESPONSE OF THE DATA CENTER TO COOLING LOSS DURING FACILITY POWER FAILURE

Approved by:

Dr. Yogendra Joshi, Advisor
School of Mechanical Engineering
Georgia Institute of Technology

Dr. S. Mostafa Ghiaasiaan
School of Mechanical Engineering
Georgia Institute of Technology

Dr. Sheldon Jeter
School of Mechanical Engineering
Georgia Institute of Technology

Date Approved:

ACKNOWLEDGEMENT

I would like to thank Dr. Yogendra Joshi, Dr. Michael Patterson, and Michael Meakins: together they gave me the opportunity to be involved in this research and advised me through the various challenges, redirecting me when I needed redirecting and struggling with me when I struggled. I would also like to thank Dr. Sheldon Jeter, who served not only as a committee member, but discussed and aided in challenging experimental setups. Dr. Ghiaasiaan, was an excellent professor and very approachable when my challenges overlapped his experience.

Eric Burgett provided computing expertise by building a compute cluster capable of performing in these experiments out of salvaged, finicky machines; his help was indispensable. In addition, Hasan Abbasi and Matthew Wolf from the GA Tech College of Computing took time to help transfer the donated servers to extend their useful life in our lab.

My colleagues Ankit Somani and Qihong Nie discussed with me the science behind the observations. Ankit also helped me build the control volume that we both used and calibrate the sensors. It is difficult to include everyone who gave useful input and discussion due to the large numbers, but at least my colleagues in the METL and CEETHERM research groups should be mentioned. Sometimes we disagreed about the cause of certain physical phenomena, but those “arguments” were a great catalyst for understanding the problems under consideration.

I thank my wife for putting up with my graduate school salary and late nights at the lab. I thank those who brought me to this point with their training and investment in me as a person: parents, family, teachers, employers, mentors, business colleagues, and personal friends. Most importantly I thank God for life and strength and blessings to make it through all life’s challenges.

/

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF SYMBOLS

c_p	Specific heat, heat capacity
q	Quantity of energy transferred
$\frac{dq}{dt}$	Rate of energy transfer with respect to time (Power)
t	Time
m	Mass
T	Temperature
ΔT	Change in temperature over a period of time
$\frac{d\Delta T}{dt}$	Rate of change of temperature with respect to time
i	Index representing an element of thermal mass that may be considered
τ	Time constant: assumed to be first order
max	Highest (temperature) from which the CRAC response begins
min	Lowest (temperature) that the CRAC response achieves

LIST OF ABREVIATIONS

AH	Air Handler
BC	Boundary Condition
CFD/HT	Computational Fluid Dynamics and Heat Transfer
CHW	Chilled Water
CPU	Central Processing Unit
CRAC	Computer Room Air Conditioning
HVAC	Heating, Ventilation, and Air Conditioning
HX	Heat Exchanger
IT	Information Technology
LMTD	Log Mean Temperature Difference
PIV	Particle Image Velocimetry
RFP	Raised-Floor Plenum
UPS	Uninterrupted Power Supply
VFD	Variable Frequency Drive

SUMMARY

Provisioning for power failure is an important element of data center design. It is important to assess both tangible and intangible costs of unplanned data center downtime. These costs must be compared with the capital cost of providing various levels of backup power infrastructure to compute and cooling equipment. Various levels of backup power infrastructure each lead to a most probable transient scenario after utility power failure. Because of differences between facilities, the level of risk that unacceptable compute equipment inlet temperature associated with each level of backup power infrastructure is not standardized; in particular, facilities with differing compute equipment power densities may require different levels of backup power infrastructure to maintain safe operation. Choosing one level of backup power infrastructure above another is not necessarily obvious for every facility, as there may be large gaps in costs and unknown levels of risk for lower levels of provisioning.

A first order model is also used to compare inclusion of various thermal capacitance values with experimental results. Room level experiments also illustrate the relative level of risk associated with various levels of provisioning for the same control volume and compute equipment. Although provisioning to back up as much equipment as possible remains the “safest” solution, cost will continue to play a factor in facility design decisions. This work offers a step toward appropriate modeling of data center power failure events and suggests further steps to continue the process.

1 Introduction

1.1 Typical Elements of a Data Center

1.1.1 Variations on Compute Equipment within the Data Center

Data centers can vary substantially between facilities. The basic function is, of course, to house compute equipment and provide necessary power infrastructure and cooling infrastructure to compute equipment. Compute equipment can vary from clustered PCs housed on baker's racks (see Figure 1.1-1) to large one-piece mainframes (see Figure 1.1-2), but the typical data center is designed around clusters of rackable servers (see Figure 1.1-3). This thesis assumes the use of rackable servers, although much of the work could be easily extended to less standard compute equipment.

Figure 1.1-1: Perseus Beowulf compute cluster constructed from PCs and housed on Shelves. [15]

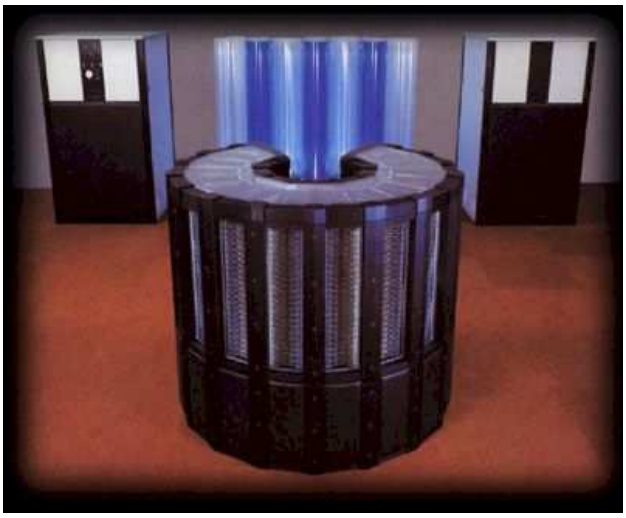

Figure 1.1-2: Cray-2 Mainframe. [17]

Figure 1.1-3: Server racks with clusters of rackable servers. [18]

1.1.2 Examples of Cooling Air Distribution Schemes

Because of cost considerations and due to concerns caused by pressurized water piping in close proximity to the electronics, compute equipment within the data center is typically cooled by forced convection of air that is drawn through the compute equipment via integral fans. This cooling air requires a distribution scheme. Server racks (SR) are

typically arranged in the hot aisle/cold aisle distribution scheme shown in Figure 1.1-4; alternatives and improvements have also been explored (See Figure 1.1-5). [15] Although overhead cooling is possible and exists at times (See Figure 1.1-6), particularly for office buildings with only a few server racks, the raised floor plenum (RFP) air distribution scheme has become dominant within the industry. This thesis assumes air cooling with raised floor plenum distribution to server racks in hot-aisle/cold-aisle arrangement. However, the work could easily be extended for use with other cooling schemes, especially those using air cooling.

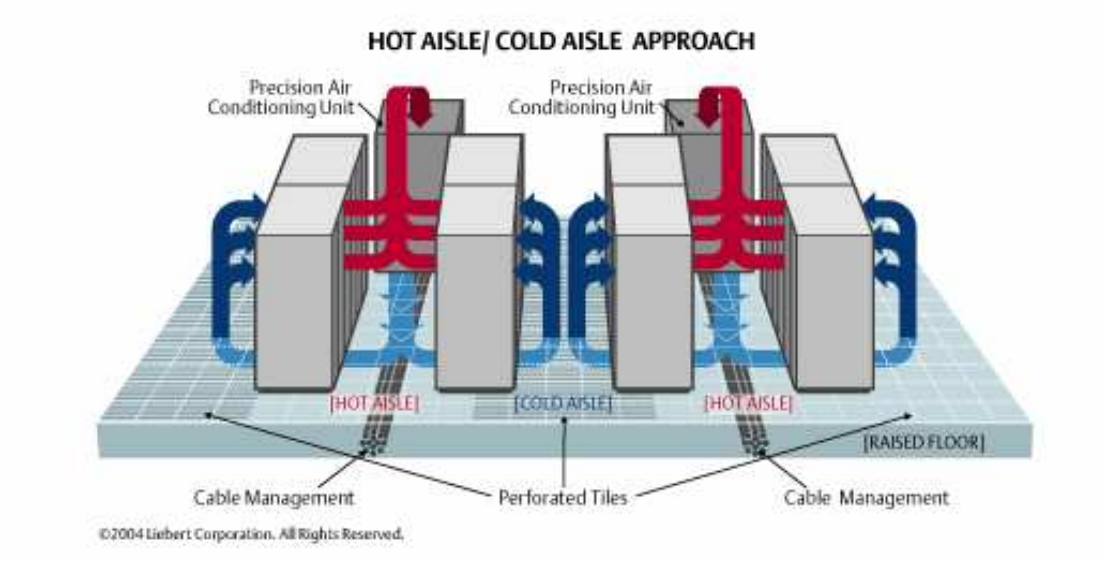

Figure 1.1-4: Hot aisle cold aisle configuration. [18]

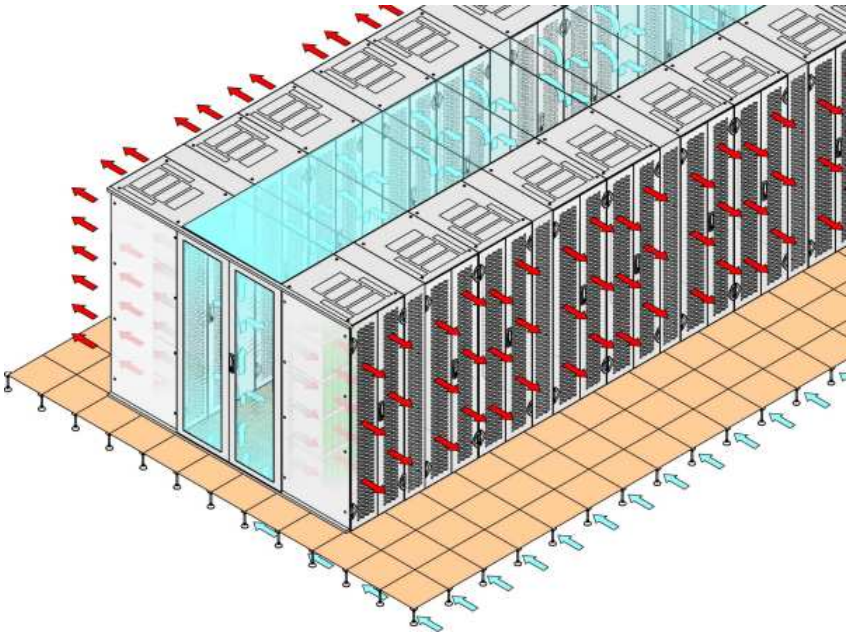

Figure 1.1-5: Cold aisle containment. [20]

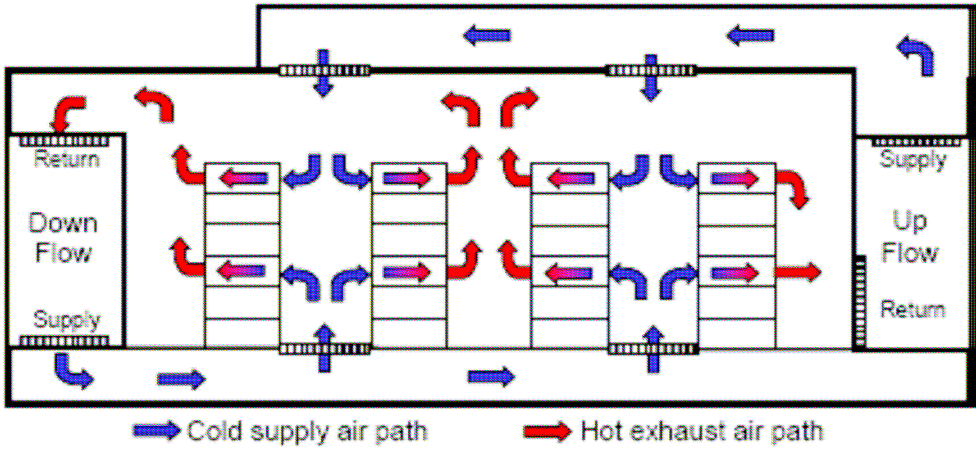

Figure 1.1-6: Two alternate air distribution schemes overhead ducts vs. raised floor plenum. [6]

1.1.3 Computer room air conditioning unit

Air cooling of a data center with a raised floor plenum is accomplished with some form of air handler (AH). This AH has traditionally been a computer room air conditioning (CRAC) unit, but built up AH and other permutations are also available. Much of this thesis could be applied to any AH, but experimental and computational

works were performed with a CRAC unit—specifically a CRAC unit employing an air to water heat exchanger. Visual representations of CRAC unit type AH are shown in Figures 7 and 8.

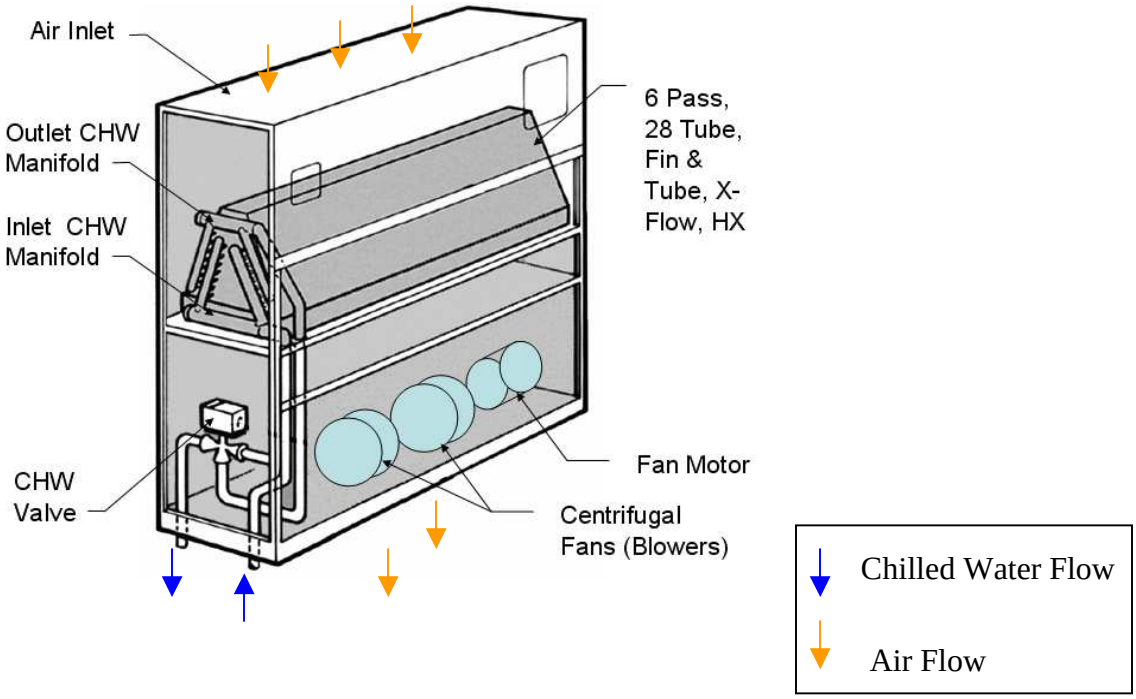

Figure 1.1-7: Schematic of CRAC AH unit used in experiments. [14]

1.2 Increases in Compute Equipment Power Density

1.2.1 Historical Developments in Compute Equipment and Cooling Systems

Compute equipment has gone through several generations of development. Initially, computers were room sized, built with light bulb sized bipolar junction transistors and generally air-cooled. The next generation were mainframes; early mainframes were air-cooled, but liquid cooling became necessary as the density of transistors increased. Early liquid cooling schemes employed such fluids as fluorocarbons and liquid nitrogen. With the invention of CMOS chips compute equipment began to accomplish more data processing per unit of power required and heat dissipated; as the power density (PD) decreased, so did the heat fluxes. As physical understanding and manufacturing processes have developed, the feature sizes of electronics on CMOS chips

have been reduced and more transistors are included per unit area. Add to this the development of dual and quad core architecture that places more chips per unit volume. Though there have been gains in efficiency that lead to higher level of data processing per unit of heat dissipated, the power density still continues to increase. Some compute equipment has begun to again employ liquid cooling schemes, but, due to the higher cost and frequent upgrades, most compute equipment remains air-cooled. Compute equipment is often designed to be as highly dense as possible without making cooling impractical; the goal is not to raise power density, but that is a by-product of raising the density of data processing. Facility floor space is expensive and it is desirable to accomplish more computing per unit area. Thus, data center design becomes more challenging as the limits of air-cooling are explored.

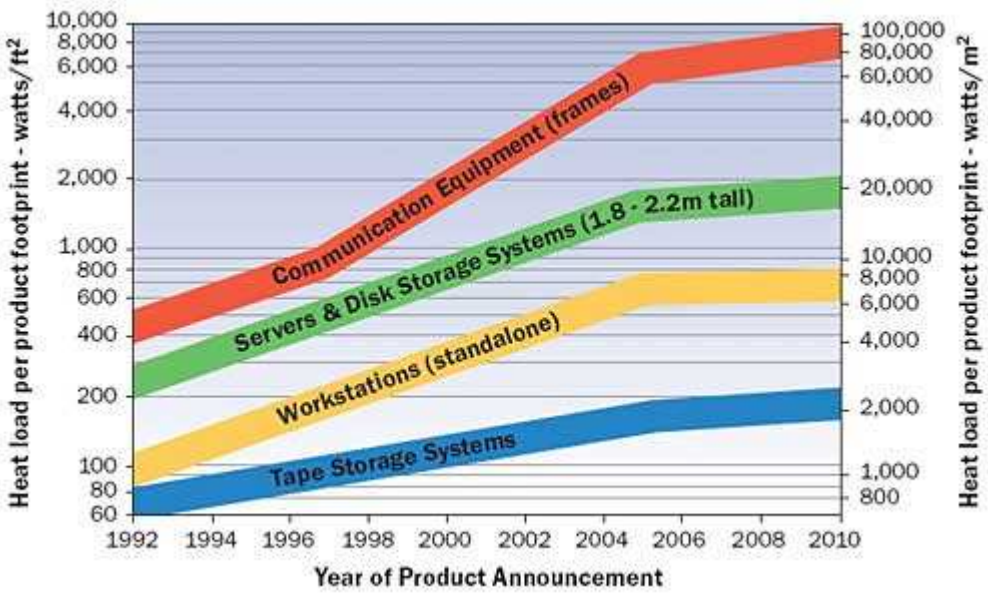

Figure 1.2-1: Increases in compute equipment power density. [29]

1.2.2 Design Considerations for Increased Power Density

Increased power density in data centers introduces new considerations for facility designers. Recirculation of air from the outlet of compute equipment back around the server rack can occur, with hot air passing either over the top of the row of racks or

around the sides at the end of the row; this can cause compute equipment to experience inlet temperatures above the ASHRAE acceptable limit (32°C, dry bulb [34]). To counter this effect designers and facility managers have increased CRAC flow rates to force more air towards the top of the racks, decreased temperature set points (SP) to lower the temperature of recirculated air, and/or installed barriers to separate the hot and cold air streams. Such steady state problems are an area of concern and ongoing research, but this thesis is concerned primarily with predicting and exploring the transient scenarios introduced during a data center power failure (PF) event.

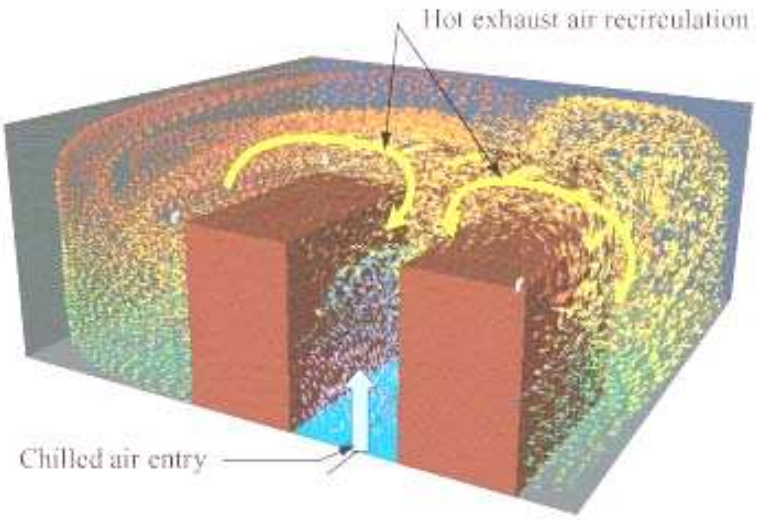

Figure 12: Recirculation of air from hot aisle to cold aisle. [30]

1.3 Importance of Provisioning for Power Failure

1.3.1 Cost of Data Center Downtime

The US EPA estimates that 90% of data centers will be affected by a power failure within a 1-year period. [3] While it must be acknowledged that some of these power failures are too short to affect facility operation, the percentage is still very high. Not all power failures in data centers are published or receive public attention beyond mention in internet blogs, but a recent failure at a Dallas Rackspace data center illustrated that even the best prepared facility can stop operating given a series of events that repeatedly causes infrastructure to fail. [23] The detailed description of the power failure

given by Rackspace was meant to show how improbable it was that such a failure should take place again. It illustrates how seriously such events are taken by data center users.

Data centers are operated for a variety of purposes, which are not limited to the following: hosting of websites, scientific computing, rendering of animation, storage and processing of company information, control of industrial processes, and sale of compute space under service level agreements (SLA). Downtime of all of these processes incurs expense to the user and typically the equipment owner if these are not the same party. Under the SLA, the owner may have to pay, or forfeit being paid by, the user. A company computer may lose information currently being processed if it goes down. A data center controlling an industrial process could extend the process downtime caused by a power failure as the data center is being restarted, unavailability of websites may cause missed opportunity for sales. In general, a compute cluster that loses power unexpectedly may have an awkward restart that requires expert IT support.

- Impact on stock price
- Cost of fixing / replacing equipment
- Cost of fixing / replacing software
- Salaries paid to staff unable to undertake productive work
- Salaries paid to staff to recover work backlog and maintain deadlines
- Cost of re-creation and recovery of lost data
- Loss of customers (lifetime value of each) and market share
- Loss of product
- Product recall costs
- Loss of cash flow from debtors
- Interest value on deferred billings
- Penalty clauses invoked for late delivery and failure to meet Service Levels
- Loss of profits
- Additional cost of credit through reduced credit rating
- Fines and penalties for non-compliance
- Liability claims
- Additional cost of advertising, PR and marketing to reassure customers and prospects to retain market share
- Additional cost of working; administrative costs; travel and subsistence etc.

Figure 1.3-1: Potential causes of loss due to downtime. [31]

Industry Sector Revenue Per Hour	Lost Revenue Per Hour (Millions of US Dollars)
Energy	2.8
Telecommunications	2.0
Manufacturing	1.6
Financial Institutions	1.4
Information Technology	1.3
Insurance	1.2
Retail	1.1
Pharmaceuticals	1.0
Banking	0.97

Figure 1.3-2: Lost revenue per hour of data center downtime in various industries. [31]

1.4 Determining Required Level of Backup Power Infrastructure

1.4.1 Complexity of the Task and Need for Modeling during Facility Design

All of the above negatives represent expenses caused by compute equipment downtime; these expenses must be compared with the initial cost required to provide infrastructure that allows operation of the data center during a power failure event on a case-by-case basis. The extra expense incurred by raising the reliability of the data center can include not only the capitol cost of added infrastructure, but also the increased energy consumption caused by inherent inefficiencies in uninterrupted power supply (UPS) systems—as UPS load capacity increases, so does its associated power loss.

Proper design of power infrastructure and cooling infrastructure can delay, or even virtually eliminate, temperature rises during power failure. However, such increased infrastructure requires increased capital expenditures and raises the first cost of the facility. Facility owners (owners) who place reliability as a first priority may be willing to make such expenditures based on broad recommendations, but many could benefit from

quantification of the risk associated with various power infrastructure/cooling infrastructure combinations. It is not possible to give a blanket recommendation of the infrastructure that will achieve an acceptable cost/risk combination; the optimum combination depends upon factors such as equipment power density and cost of equipment downtime. A model that predicts the time available for safe equipment operation during a power failure under various power infrastructure/cooling infrastructure combinations could help quantify the risk of each combination. Thus, owners and designers would have a basis for choosing among various options. One of the goals of this thesis is to make progress toward such a model.

1.4.2 Basic Cooling Scheme and Equipment Considered

A variety of designs are possible for HVAC applications and it is outside the scope of this thesis to discuss them in detail. Instead, a simplified, typical system will be used to illustrate the power failure scenarios under consideration. Consider the cooling system shown in Figure 14. Heat flows from the compute equipment to the data center air, the energy stored in the air is released to the CHW via an air to water HX within the AH, the chiller then removes the heat from the CHW through a thermodynamic cycle. For the purposes of this thesis, it is not relevant to question whether the chiller rejects heat directly to ambient air or into cooling tower water—the answer does not change the power failure scenarios considered.

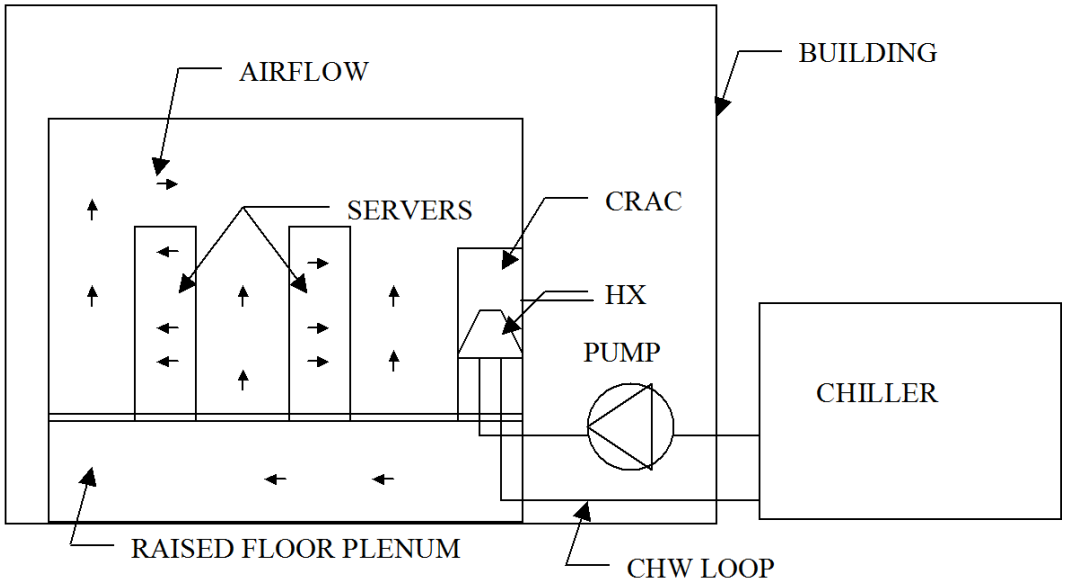

Figure 1.4-1: Heat flow from server equipment to chiller

1.4.3 Levels of Backup Power Infrastructure and Uptime Institute Tiers

There are various permutations of backup power infrastructure that are designed to maintain data center operation during a power failure. Plans for backup power to data center equipment are included in a classification of mission critical facilities developed by the Uptime Institute [5]; facilities are classified into categories called Tiers. This classification requires evaluation of other factors besides power and cooling infrastructure. The choice of backup power infrastructure provided, or not provided, to each level of cooling infrastructure determines the most likely transient scenario during a power failure. Various scenarios are shown in Figure 1.4-3 and described below.

Figure 1.4-2: Emergency generator fueled by natural gas. [21]

The Uptime Institute recommends that some cooling infrastructure, namely air handler fans and chilled water pumps, be provided with UPS. However, there is some ambiguity over whether UPS backup power infrastructure is *required* for cooling equipment in order to achieve higher tier classifications. Emergency generator backup power infrastructure, on the other hand, is unambiguously required in order to achieve tier of Tier II or higher. There is good reason for the ambiguity over whether to provide UPS backup power infrastructure to cooling equipment. The main goal of raising the tier rating of a data center is to ensure continuous system operation; some data centers can easily maintain operation during the time it takes cooling infrastructure to come back online after a power failure, while others may develop localized areas with unacceptable compute equipment inlet temperatures within seconds after the AH stops circulating air. Thus, the tier rating attempts to quantify risk, rather than specifying a required level of backup power infrastructure for the cooling infrastructure of the data center.

1.4.4 No Backup Power Provided

Some data centers do not have any backup power, or provide backup power within the rack for the head nodes only, just enough to minimize the amount of labor necessary to bring the compute equipment back online after the power failure is over. In this case cooling of compute equipment is generally not an issue, since the power density

of data center drops drastically during the power failure. Such a data center falls in to the category of Tier 1.

t_0 = Power failure: servers, CRACs, and pumps on UPS. Chiller on generator.	t_1 = Emergency generator comes online	t_2 = Chiller controls initiate compressor restart sequence.
t_0 = Power failure: servers and CRACs on UPS. Pumps and chiller on generator.	t_1 = Emergency generator comes online	t_2 = Pumps restart, chiller controls initiate compressor restart sequence.
t_0 = Power failure, servers on UPS. CRACs, pumps, and chiller on generator.	t_1 = Emergency generator comes online	t_2 = CRACs restart, pump restarts, Chiller controls initiate compressor restart sequence.

Figure 1.4-3: Some possible power failure and cooling equipment restart scenarios.

1.4.5 Uninterrupted Power Supply for Compute Equipment Only

The lowest level of provisioning in designing backup power infrastructure (BPI) would be to provide UPS for compute equipment only. This means that compute equipment would continue to process data, but also to dissipate heat; running compute equipment without cooling infrastructure is cause for concern, particularly as the power density of compute equipment continues to increase. This design allows for the data center to continue smooth operation through very short interruption in power supply, but may cause compute equipment inlet temperatures to rise unacceptably if the power failure disrupts cooling for an extended period. The problem is compounded as power density increases. It will likely be necessary for this type of data center to shutdown its compute equipment during an extended power failure. Shutdown of compute equipment can take

place automatically or manually, depending on the control scheme employed by the manufacturer. Manual shutdown runs the risk of bad user judgment, or slow reaction. Automatic shutdown runs may cause problems if the machine is shutdown during an important task. However, in some cases, a data center with lower power density may be able to continue operation without downtime if an emergency generator is also installed.

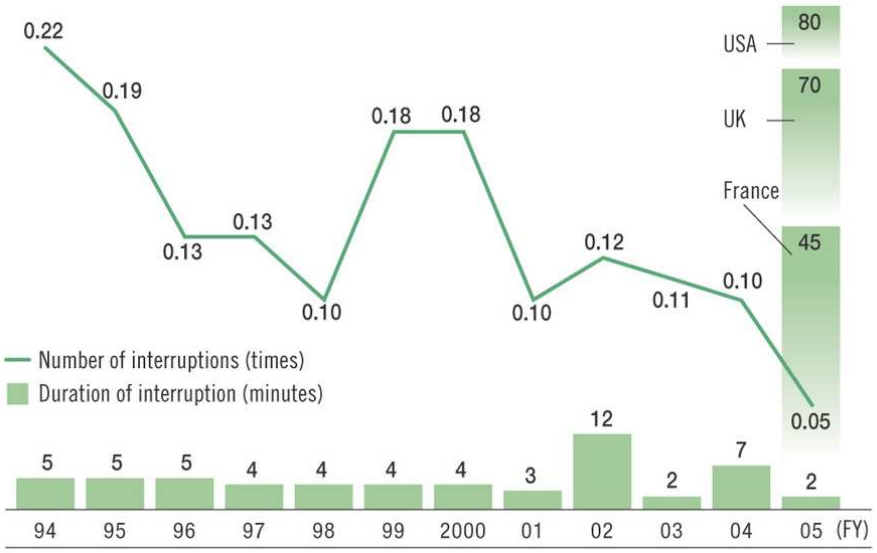

Figure 1.4-4: Example utility data showing average downtime per power outage. [22]

1.4.6 Emergency Generator provisioning for Cooling Infrastructure

The addition of an emergency generator to the backup power infrastructure shortens the maximum theoretical time span of a power failure for any equipment that it supplies. Adding power infrastructure to the emergency generator tends to be substantially less expensive than providing it with UPS. Thus, chillers—the cooling infrastructure that tends to draw the largest amount of power—generally are not provided with UPS, but may receive emergency generator backup power infrastructure. Typically, an emergency generator can restore power within about 30 seconds or less after a power failure. At this point the AH and CHW pump with emergency generator backup power infrastructure would start almost immediately. The facility can begin recovery from a power failure event even before the chillers restart. When the AH and CHW pump begin

working, heat is transferred from the data center air into the CHW; this CHW is still at approximately the chiller evaporator setpoint until it passes through the heat exchanger (HX) in the AH. The CHW will eventually make a complete circuit around the CHW loop—the time it takes to make a complete circuit depends on pipe diameter, pipe length, and CHW flow rate—once the CHW makes a complete circuit the AH will no longer have as much cooling capacity (This decrease in capacity is due to reduced log mean temperature difference (LMTD) across the HX). However, the AH can still continue to transfer heat to the CHW and limit the temperature rise in the data center. With CHW storage—discussed in a later paragraph—the temperature rise in the CHW loop can be extended proportionally to the amount of CHW stored. Chillers typically require a restart sequence for each compressor. The restart sequence, if successful, may take up to 180 seconds for the first compressor and 110 seconds for each additional compressor. Provisioning for a failed compressor may also be a wise step. Thus, a facility provided with emergency generator backup power infrastructure to for its compute equipment and cooling infrastructure needs to determine what additional steps are needed to ensure the facility does reach unacceptable temperatures before the chiller has a chance to come back online.

1.4.7 Uninterrupted Power Supply for Air Handler Fans

It may seem as though adding an emergency generator as backup power infrastructure for the cooling infrastructure would be sufficient in every case, but there are times when more is required. Without any UPS provided for the AH, there is necessarily a delay in air circulation within the data center. With higher power density compute equipment, this delay can lead to recirculation of air from compute equipment outlets to inlets. In certain facilities, this may lead to unacceptable compute equipment inlet temperatures as soon as airflow patterns change. An improvement to the backup power infrastructure, after providing UPS to compute equipment and emergency generator backup power infrastructure to all cooling infrastructure, is to provide UPS for the AH. The first expected benefit from running the AH is circulation of the air within the data center during the power failure. Running the AH will retain the air flow patterns that pressurize the cold aisle and minimize recirculation and hot spots. Stratification of the air

within the data center could also contribute to temperatures that are substantially above that predicted by a uniform, well mixed assumption; running the AH mixes colder air from the plenum with hot air from higher in the room to reduce this effect. Another benefit of running the AH is the increase in heat transferred to other media besides the room air. The cooling coil within the AH is likely the most important thermal storage medium that becomes available when the AH keeps running. Although its temperature rises, the coil keeps storing energy as long as the air coming into the AH keeps rising. Other media that may become available for energy storage due to continued airflow are solid surfaces in the facility. Building materials, such as the concrete floor, raised floor plenum tiles, walls, and ceiling may begin storing energy, or even conducting it away, as the time scale of the power failure event increases and the air temperature rises. Any extension of the time window within acceptable operating temperatures gives the facility more time for power to be restored, either by emergency generator or by the utility service, and cooling infrastructure to come back online.

1.4.8 UPS for Chilled Water Pumps and Air Handler Fans

In many cases, providing backup power infrastructure—especially UPS—to the AH may be enough to keep temperatures within acceptable ranges long enough for the emergency generator and chillers to come back online. However, some owners may desire a greater degree of reliability. If power density is considered high enough within the data center that there is concern that server inlet air temperatures may rise too rapidly, or if a greater degree of reliability is desired, the CHW pump can also be placed on UPS backup power infrastructure. Under this scenario, the data center will continue operating at steady state, with no temperature rise, until all the CHW in the piping has made a complete circuit through the AH. With sufficient chilled water storage, the facility can operate with no rise in air temperature for an extended period. This would allow for false starts of both the generator and chiller.

1.4.9 Discussion of Redundancy Tier Requirements Excluded

The highest tier ratings for data centers are based not only on backup power infrastructure, or CHW storage, but also on redundancy. It is not the purpose of this thesis

to discuss the redundancy required by the rating system, but suffice it to say that redundancy allows pieces of equipment or systems to fail without disrupting operation of the data center. This thesis is concerned with utility power failure and planned infrastructure response. It does not consider what happens when facility equipment fails.

1.5 Modeling of Transient Temperature Response of Data Center air during Power Failure

1.5.1 The Need for a Model

Attempts to choose the level of backup power infrastructure needed for cooling infrastructure in order to maintain acceptable compute equipment inlet temperature during power failure expose the need for development of improved modeling techniques. Namely, it is difficult to accurately predict how much time is available for the compute equipment to operate, given a particular choice of backup power infrastructure that allows part of the cooling infrastructure to stop operating. Previous models generally are either too conservative, or worse yet, too optimistic. Some models even combine overly optimistic and overly conservative elements. The absence of sufficiently accurate or descriptive models speaks to the difficulty of the subject and lack of previous exploration.

In addition to helping determine what level of backup power infrastructure is needed for each element of the cooling infrastructure, an accurate transient model could be used to predict the amount of time gained by lowering the temperature setpoint in the data center. Within a certain range, raising the setpoint can increase cooling efficiency in the data center. However, above a certain temperature, the server fans will begin drawing enough power to mitigate further increases in efficiency. [24] Some managers have ignored the recent trend toward increasing data center setpoint due to a belief that cooler air at the beginning of a power failure even may buy them more time to operate the data center with acceptable compute equipment inlet temperaures. A reasonable model could shed light on how much time is actually gained.

1.5.2 Assumptions in Existing Models

Early models were performed with building modeling software such as Energy Plus. [13] The results of these models were offered, not as a design tool, but as a general warning that temperatures can increase rapidly in a data center during power failure. They also demonstrated that increasing power density exacerbates this effect. [12] These were useful first steps. However, such modeling software assumes a uniform, well-mixed air temperature within the facility and considers the thermal capacitance of building materials only. These assumptions are correct for thermal events that take place over the course of hours or days and have relatively uniform air temperatures within each volume considered. Data center power failures do not fall into this category. Therefore, these models should be considered a step forward rather than a final word. Comparison with more detailed models and experiments would provide needed validation or improvements.

1.5.3 Models Assuming Well Mixed Air and Steady State Servers

Figure 1.5-1 and

Figure 1.5-2 show data and predictions taken from a data center design manual. These were intended as a warning that higher power density data centers will experience a faster temperature rise in a power failure. Very little information is given regarding model geometry, or input parameters.

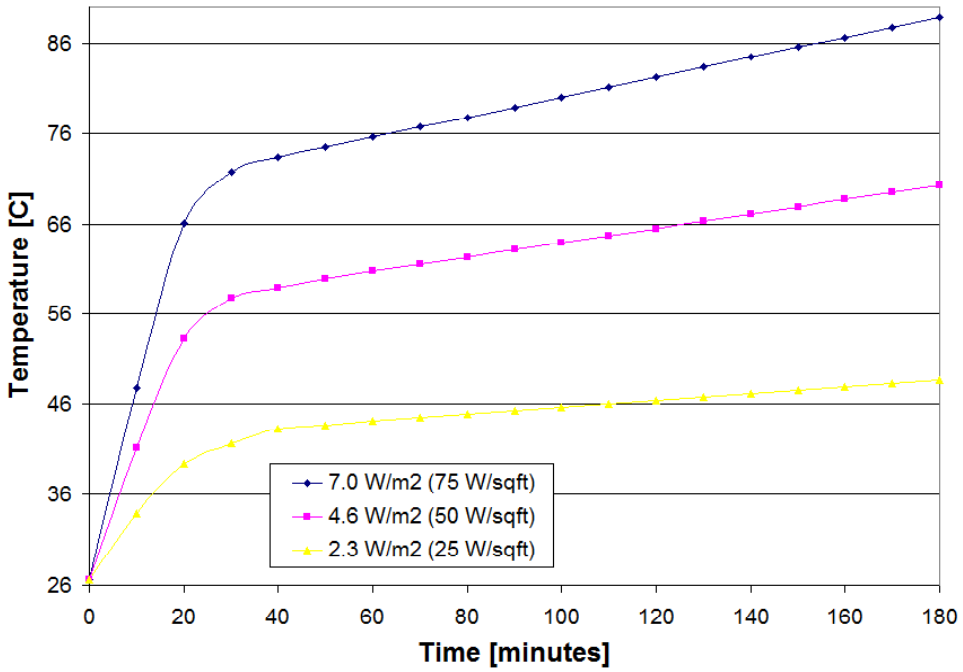

Figure 1.5-1: Early Energy Plus modeling average air temperature rise predictions. [13]

Power Dissipation	Rate of Rise-Model	Rate of Rise-Experiment
[W/m2 (W/SQFT)]	[℃/min (℉/min)]	[℃/min (℉/min)]
810 (75)	2.1 (3.8)	2.4 (4.3)
540 (50)	1.4 (2.5)	1.2 (2.2)
270 (25)	0.72 (1.3)	0.56 (1)

Figure 1.5-2: Comparison of Energy Plus predictions with experiment (setup details unavailable). [12]

On the one hand, uniform, well-mixed models are too optimistic, because they do not account for the substantial variations in temperature that are often present in a data center. Recirculation from hot aisle to cold aisle and stratification can give rise to substantial difference in compute equipment inlet temperature within the same facility. [15, 24] If the AH stops circulating air, this deviation from the average compute

equipment inlet temperature increases. [4] On the other hand, these models are too conservative, because the thermal mass of the compute equipment itself is not considered. Although a welcome first step, such models may not prove useful for prediction of safe time intervals for mission critical design.

1.5.4 Modeling of CRAC Failure with CFD/HT and Steady State Servers

A more recent modeling development uses computational fluid dynamics and heat transfer (CFD/HT) analysis to model transient scenarios within the data center. These models dispense with the popular uniform, well-mixed assumption. Therefore, they have the potential to predict the location of hot spots caused by recirculation and stratification. However, they assume the compute equipment to release a constant amount of heat and yield a constant temperature rise, as would be the case in a steady state analysis. In reality, some heat will be stored within the compute equipment before it comes to steady state. Thus, the models predict a result that may be too conservative, yet allow for some prediction of hot spots.

CFD/HT have been used to model the failure of a single CRAC unit in a facility that relies on multiple CRAC units to support its cooling load. This study explored the temperature rise in various parts of the data center upon failure of one of the six CRAC units. [4] There were three power density regions low, moderate, and high, which dissipated 1.08 kW/m^2 (100 W/SQFT), 1.40 kW/m^2 (130 W/SQFT), and 1.88 kW/m^2 (175 W/SQFT) respectively. The results predicted a rise in maximum server inlet temperature from 29°C (84°F) to 40°C (105°F) in less than 150 seconds. In this case the predicted maximum temperatures are likely more accurate than the predicted time from one steady state condition till the next because of the steady state server assumption.

Another CFD/HT model was used to predict the time available before unacceptable temperatures were reached in a high density data center after complete cooling system failure. Servers were modeled as steady state heat inputs. The model predicted how long it would take for a temperature sensor in the hot aisle for the fire suppression system to reach 35°C (95°F); it predicted 11 seconds. The model also predicted that without CRACs running, colder air in the raised floor plenum will not mix

with warmer air above floor. [26] Because servers are modeled as steady state, the model may under predict the time available while still providing insight into air flow patterns.

The use of CFD/HT may at times be necessary, particularly if the AH unit ceases to mix the air or significant hot spots are present. However, the steady state model of compute equipment within the server rack leaves room for improvement.

1.5.5 Inherent Challenges in Modeling and Experimental Validation in Data Centers

A challenge in evaluating computational models of power failure events in data centers is that so little data is available for comparison of computational results with experiment. It is understandable that performing cooling failure experiments in a data center would be generally unacceptable because of the cost associated with compute equipment and the demand for reliability by their users. While computational models may use well-established equations to describe physical phenomena, every model requires assumptions and boundary conditions BC. BC must be provided on a case-by-case basis and generally require some assumptions. One common discrepancy between models and the physical situation under consideration is caused by the large variation in length scales involved in physical processes. In a data center, Heat transfer takes place from interconnects, to chips, to heat sinks, to the air streams in individual servers, to the room air, to the cooling system. This process includes length scales ranging from less than 100 nm to multiple meters in length; it is not possible to model every detail of this of this system exactly because the mesh size required would be too large for currently available computers. In addition, the time required to define every detail of such a model would be infeasible. Even if human labor and compute memory were not issues, not every detail of a system is available. If one could measure every detail of the situation and model them, then the model might still be said to work only for the exact situation modeled.

1.6 Proposition to Model Server Thermal Capacitance

1.6.1 Event Timescales and the Need for Modeling of Thermal Capacitance

It is reasonable to question whether it is necessary to take the thermal capacitance of the compute server into account. Therefore the following criteria are suggested: if the time constant (TCON) for temperature rise at the outlet of the compute server is comparable to or greater than the time scale of the temperature rise at the server inlets, then the compact model of the compute server should account for thermal capacitance. Otherwise, if the time constant for the compute server is much less than the time scale of the room level dynamic event, then the steady state compact model may be used. One CFD/HT model was developed to determine when the temperature of a fire suppression system sensor, which was located in the hot aisle, would become unacceptable. This model predicted the temperature to be unacceptable in 11 seconds. [26] Normally, the compute server inlet temperatures are of greater interest, still this study provides a baseline for comparing experimental determined compute server time constants and the time scale of the room level dynamic event.

2 Thesis Objectives

This thesis has three primary objectives and one secondary objective, see Figure 2.1. The primary objectives focus on experimentally mimicking power failure of a facility with various levels of backup power infrastructure. The first primary objective intends to show that there is significant thermal storage in compute equipment and that their outlet temperature response is not instantaneous. The second primary objective is to explore the rate of temperature rise within the data center when all cooling infrastructure fails. The third primary objective is to investigate the thermal storage within cooling infrastructure when various cooling equipment remains online. Various scenarios are possible, but only some are considered. It is necessary to consider the situations in which the compute equipment and AH retain power, as well that in which the compute equipment and AH and CHW pump retain power without CHW storage. The secondary goal of this thesis is to explore the temperature response at the outlet of the CRAC unit when power is restored to the pump.

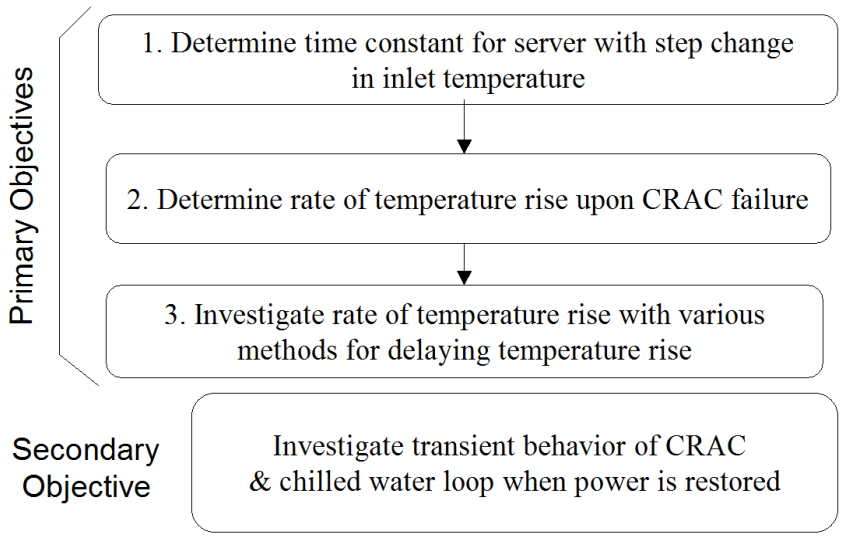

Figure 2.1: Thesis Objectives.

Another goal of this thesis is to show that servers have significant thermal capacitance compared with the short windows of safe operation predicted by previous computational modeling. This is accomplished by measuring the temperature reaction at

the outlet of two generations of servers given a step change in inlet temperature. A control volume representing the larger data center space is constructed and used to quantify the effects of thermal capacitance from various elements of the system.

3 Server Time Constant Experiment

3.1 Objectives

Transient models form a relatively small and slowly developing segment in data center modeling; the existing CFD/HT models tend to assume zero thermal capacitance for the servers. [4, 12, 26] That is to say, they assume the server reacts so fast that it can be modeled with a steady state compact model. The experiment that follows was designed to test this assumption. As power density increases and models predict shorter time windows for safe data center operation during power failure, the question of validity becomes more acute—one model predicted unacceptable temperatures within 11 seconds after AH failure. However, it was not possible to test every model of server on the market. Therefore, two generations of compute equipment with different power densities were compared. If a significantly long time constant is found, a compact model of a server that includes thermal capacitance becomes advisable.

3.2 Experimental Setup

In order to determine the time constant for the temperature rise at the compute equipment outlets, it was necessary to approximate a step change at the compute equipment inlets. The experimental setup for this experiment is shown in Figure 3.2-1 and Figure 3.2-2. The location of the sensors is shown in Figure 3.2-3. The data center laboratory was separated into two control volumes (CV); polyethylene sheet was used to separate air between the two control volumes. On cold side of the polyethylene sheet, the CRAC unit maintained a temperature between 12.3 and 14.8°C. On the other side of the polyethylene sheet, the hot side, a heater was operated, resulting in a compute server inlet temperature of 43 to 53°C. A server rack was filled with 5 Dell Poweredge 8450 7U (legacy) servers and one Intel 2U SR2500 configured with dual Xeon processors (2U Intel) server. These were chosen for two reasons: they were readily available and they illustrate the difference in compute equipment time constants due to advancements in compute equipment. Processors in the servers were loaded as close as possible to

maximum power draw with the program Prime95, except that data was also collected from a server with idle processors. The servers were instrumented with multiple type-T thermocouples at the inlets and the outlets and, in some cases, CPU and fan speed data were collected from the motherboard. When fan speed data was not available, a thermal anemometer was used to collect air speed data. Details of data collection equipment, programs, and associated uncertainty analysis are included in Appendix A.3. The uncertainty of the temperature measurements is estimated to be within ±0.3°C. Data were recorded every 10 seconds.

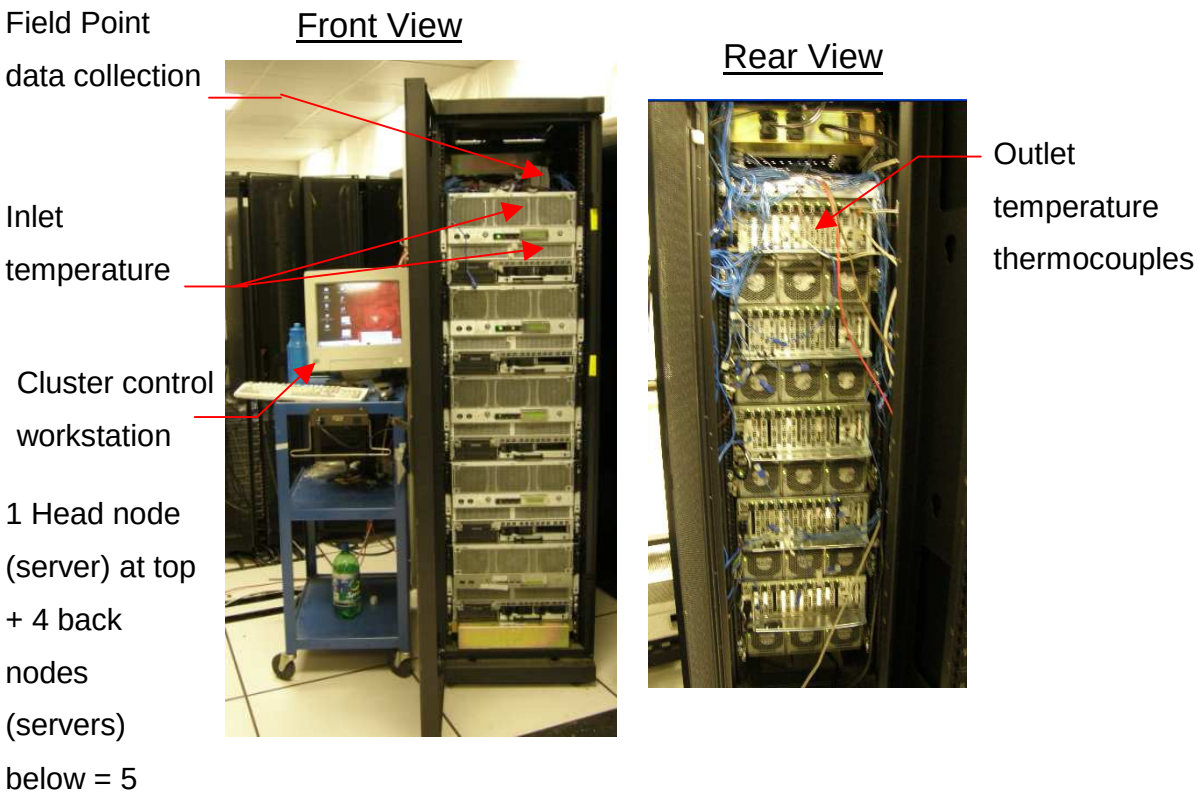

Figure 3.2-1: Servers in rack instrumented for time constant experiment (Room division also visible in background).

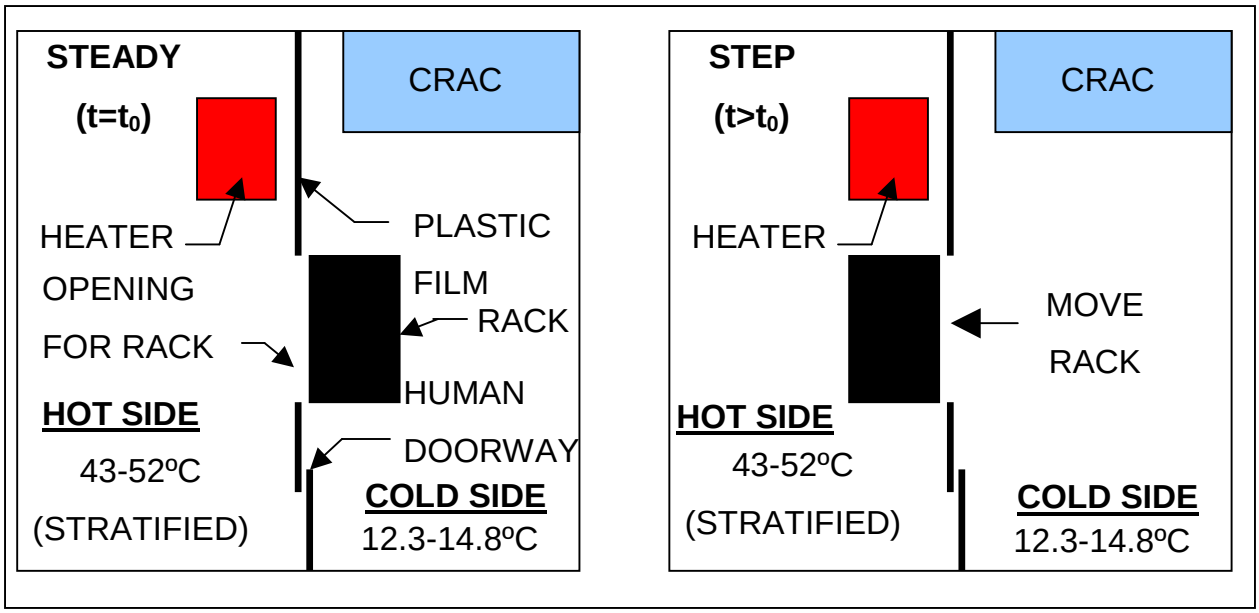

Figure 3.2-2: Division of room, conditions, and sequence of events for server time constant experiment.

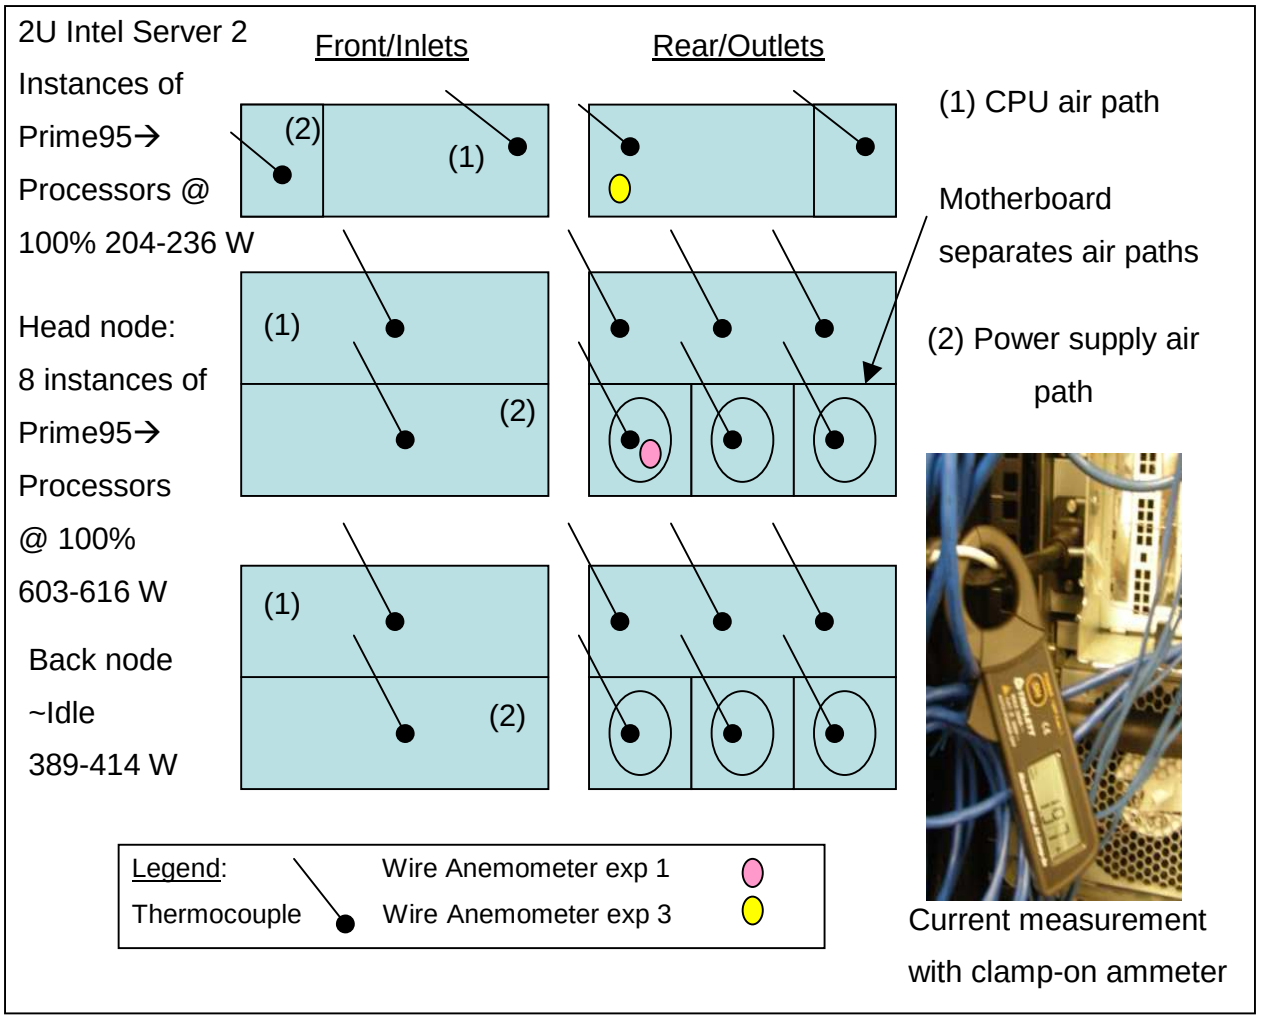

Figure 3.2-3: Location of sensors for server time constant experiment.

3.3 Procedure

To perform the experiment, the hot side and cold side were allowed to reach steady state. Then the rack of servers was pushed from the cold side to the hot side. It took approximately 1-2 seconds to transfer the rack from the cold side to the hot side. Figure 3.4-1 shows sample data from a legacy server. As shown, the inlet temperature to the server changes from the low temperature to the high temperature within 30 seconds and is therefore a close approximation to a step function. The legacy servers are constructed with two distinct air paths. The upper air path flows over the eight processors, the lower air path flows over the memory modules and power supplies, and

the motherboard separates the two air paths. Pictures of the server components are included in Appendix A.4.

3.4 Results and Discussion

3.4.1 Legacy Server with Full Processor Load

Figure 3.4-1 shows the response of the Legacy server with processors fully loaded. The time constant for the processors is 340 seconds, whereas the time constant for the power supplies is 380 seconds. These time constants are found from least squares analysis of the first 500 seconds of the response. The corresponding analysis is shown in Appendix 11. The corresponding analytical curves yield errors in predicted temperature of about 3% and 5% of the total temperature rise respectively over this 500 second region. The initial part of the curve was chosen because it was considered to predict a more conservative (lower) time constant. There is a minor departure from first order behavior in both experimental curves; that is to say that the data begin to depart from the analytical curve as time progresses. This phenomenon is best explained by considering the data shown in Figure 3.4-5. Figure 3.4-5 shows that the fan speed increases with time after the compute servers are moved to the hot side. A change in fan speed means that compute server inlet temperature is not the only boundary condition (BC) that is changing in the experiment. This is supported by Figure 3.4-2, which shows that the velocity of the air exiting the legacy power supply increases when the server is moved to the hot side. The departure is much smaller for the legacy servers because the change in fan speed is smaller by comparison and because the legacy servers do not incorporate throttling of electric current into their control scheme. Figure 3.4-3 shows components with temperature or fan speed sensors and their locations within the legacy server chassis.

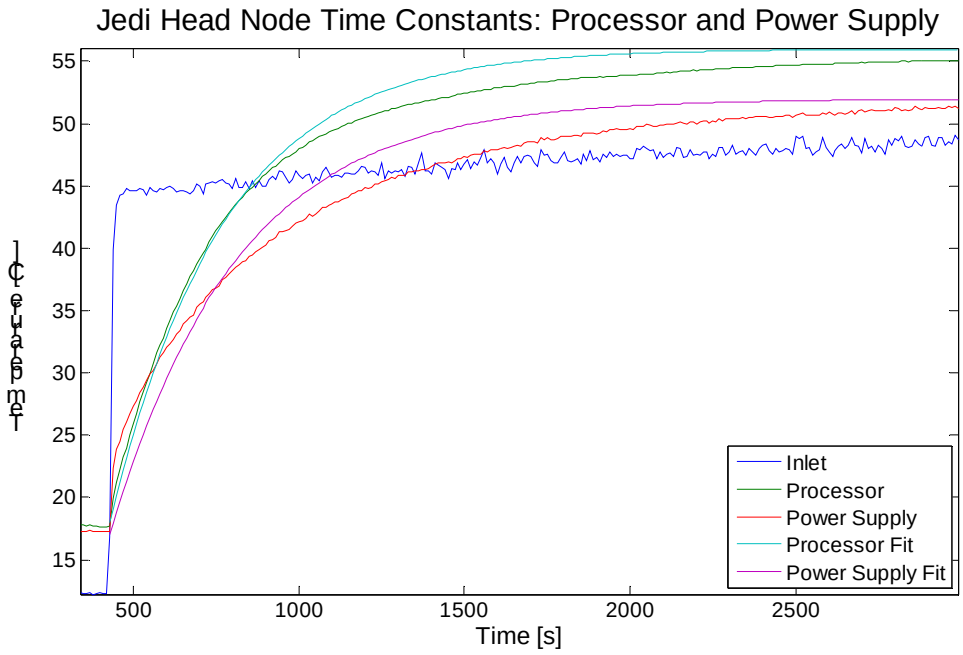

Figure 3.4-1: Time constant experimental results from legacy server with processors fully loaded.

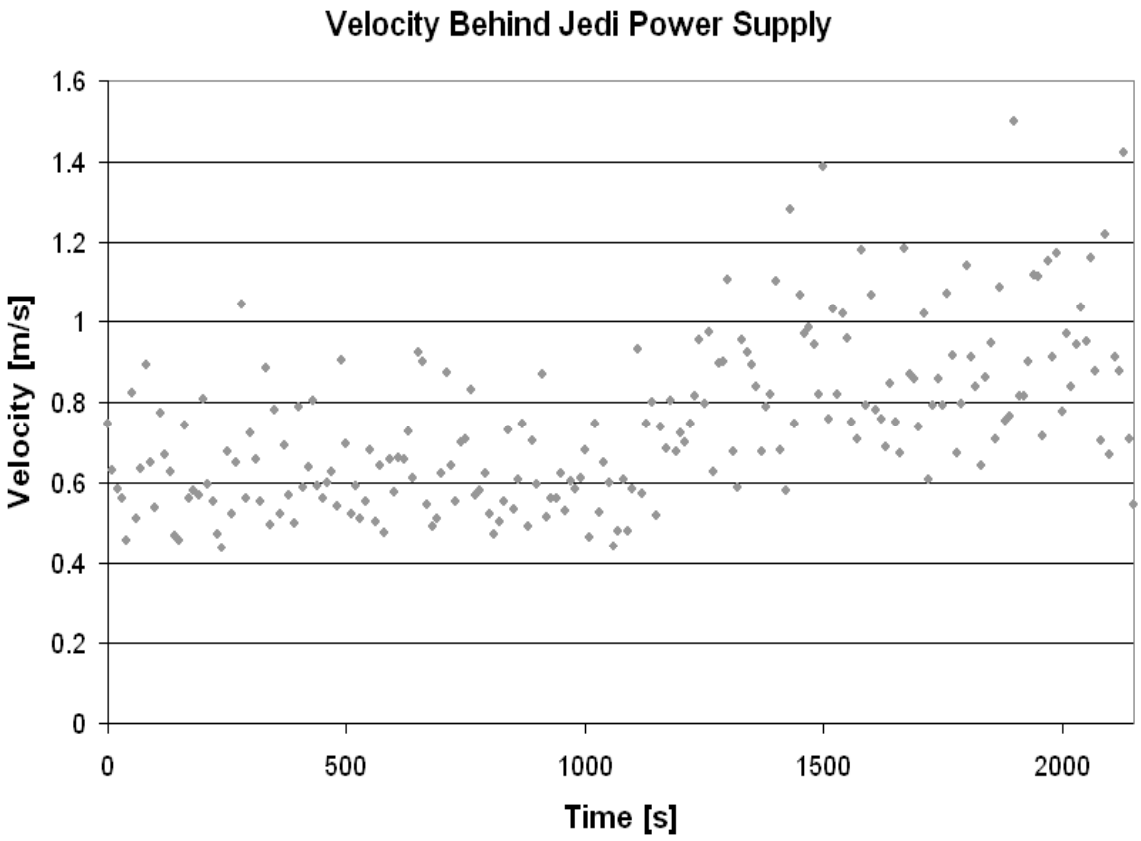

Figure 3.4-2: Velocity of air leaving legacy server power supply.

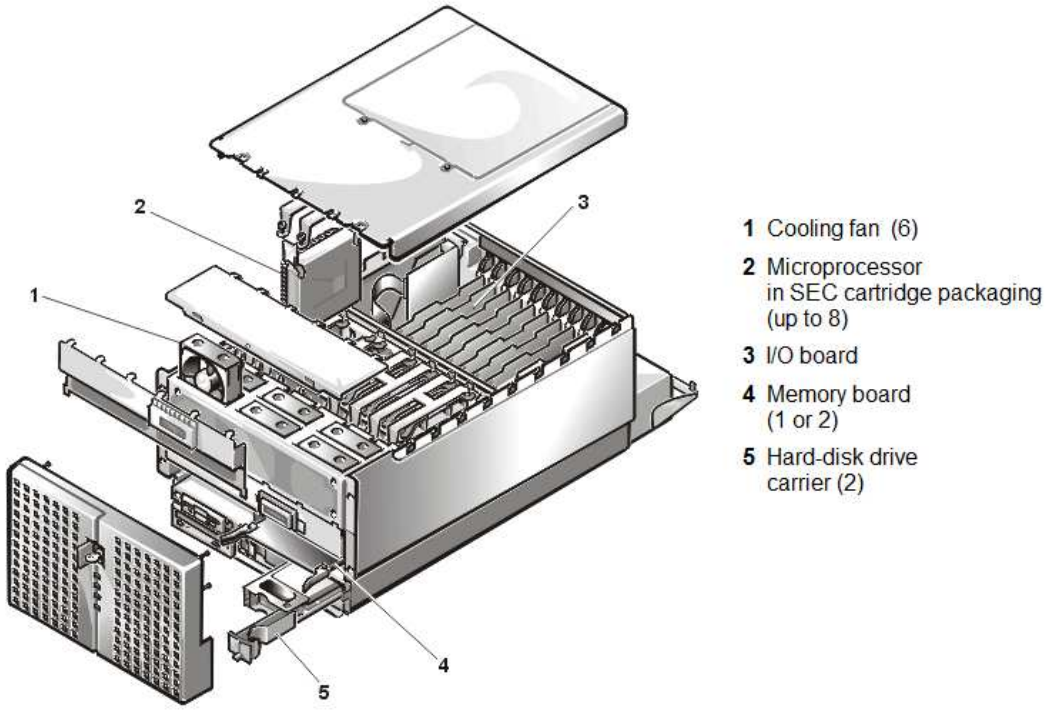

Figure 3.4-3: Layout of components with temperature and fanspeed sensors within legacy server chassis. [32]

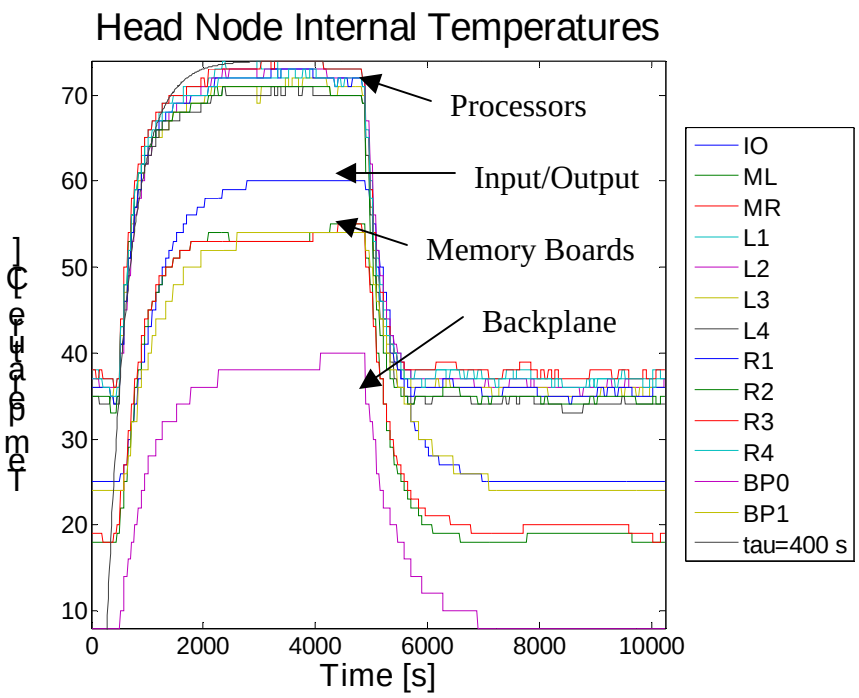

Figure 3.4-4: Temperature data from legacy motherboard sensors and first order curve.

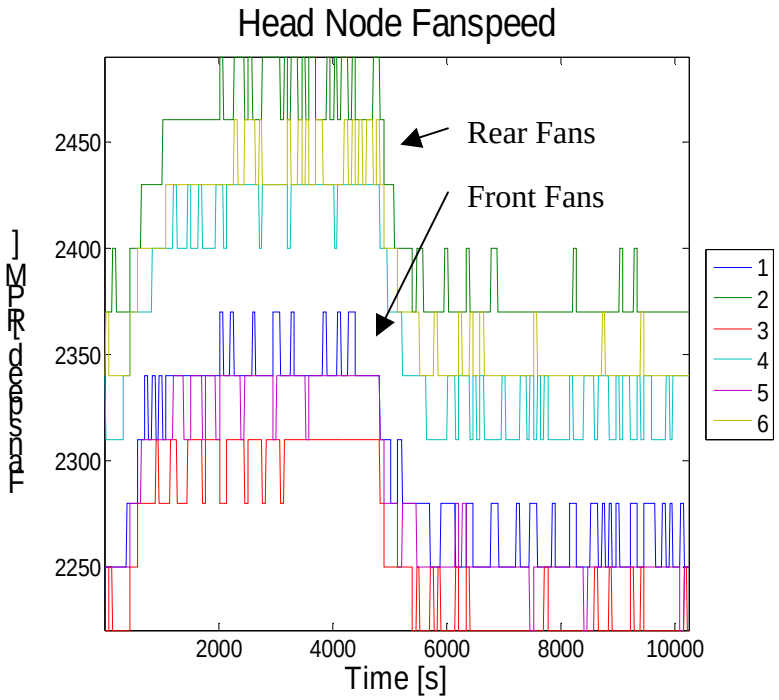

Figure 3.4-5: Processor fan speeds measured from legacy motherboard sensors.

3.4.2 Legacy Server with Idle Processors

Figure 3.4-6 shows the response of a legacy server that has idle processors and draws less power. The time constant of 370 seconds at the processor outlets and 300 seconds at the power supply outlets shows that reducing the heat dissipated also reduces the time constant, all other things being equal. As the figure suggests the power to the servers takes more than one value. These time constants are found from least squares analysis of the first 500 seconds of the response. The corresponding analytical curves yield errors in predicted temperature of about 3% and 4% of the total temperature rise respectively over this 500 second region. The corresponding analysis is shown in Appendix 11. The initial part of the curve was chosen because it was considered to predict a more conservative (lower) time constant. The departure from first order behavior is least pronounced with the idle Legacy servers.

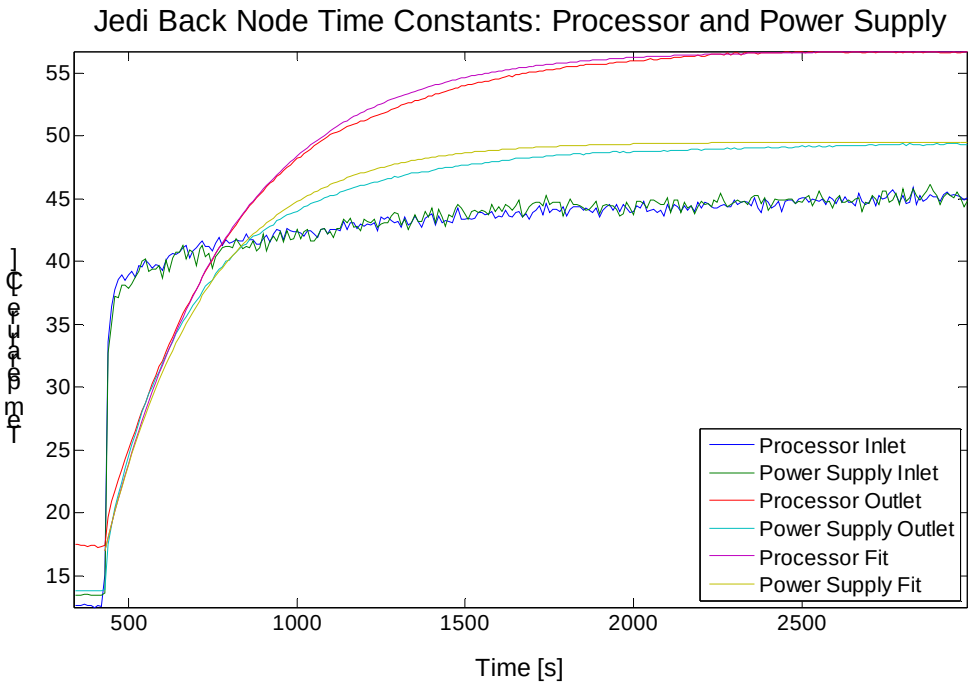

Figure 3.4-6: Time constants for legacy server with idle processors.

3.4.3 2U Intel Server with Higher Power Density and Fully Loaded Processors

Consider Figure 3.4-7, which shows the response of the 2U Intel server to the same step change experiment. The response at the outlet of the processor airflow path departs much more strongly from first order behavior; it rises quickly at first and then flattens. This response can be partially explained by Figure 3.4-8, which shows sample thermal anemometer data of the increase in air speed behind the modern server processor outlet. This change in air speed is much more substantial than that of the legacy servers, therefore, there is a substantial change in this second BC. However, this server also incorporates a processor power throttling sequence during periods of high temperature; this means that a third BC is also changing substantially and first order behavior cannot be expected. Nevertheless, it is possible to fit a first order time constant to the initial temperature rise at the processor outlet with a value of 130 seconds. This time constant will be conservative and will represent the initial response of the modern server to changes in inlet temperature at lower temperatures. The outlet temperature of the modern CS power supply does display first order behavior, as it does not employ these control schemes; its time constant is 990 seconds. Due to the faster response and stronger departure from first order behavior, the time constant was fit to the first 200 seconds of the response for the processor. The corresponding analytical curves yield errors in predicted temperature of about 3% and 4% of the total temperature rise respectively during the portion of the response that was analyzed.

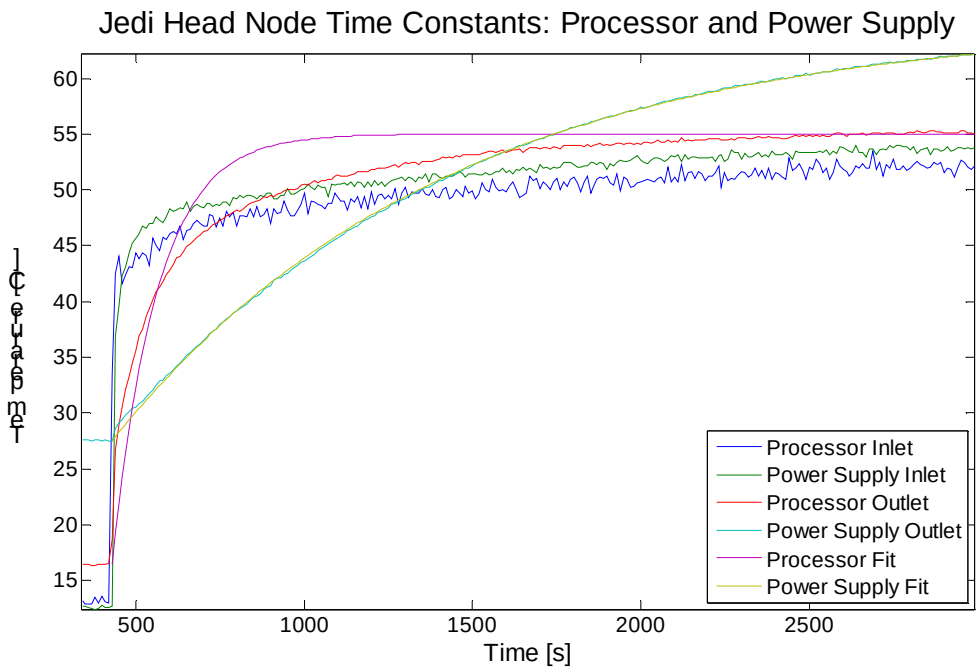

Figure 3.4-7: Response of 2U Intel server to step change in inlet temperature.

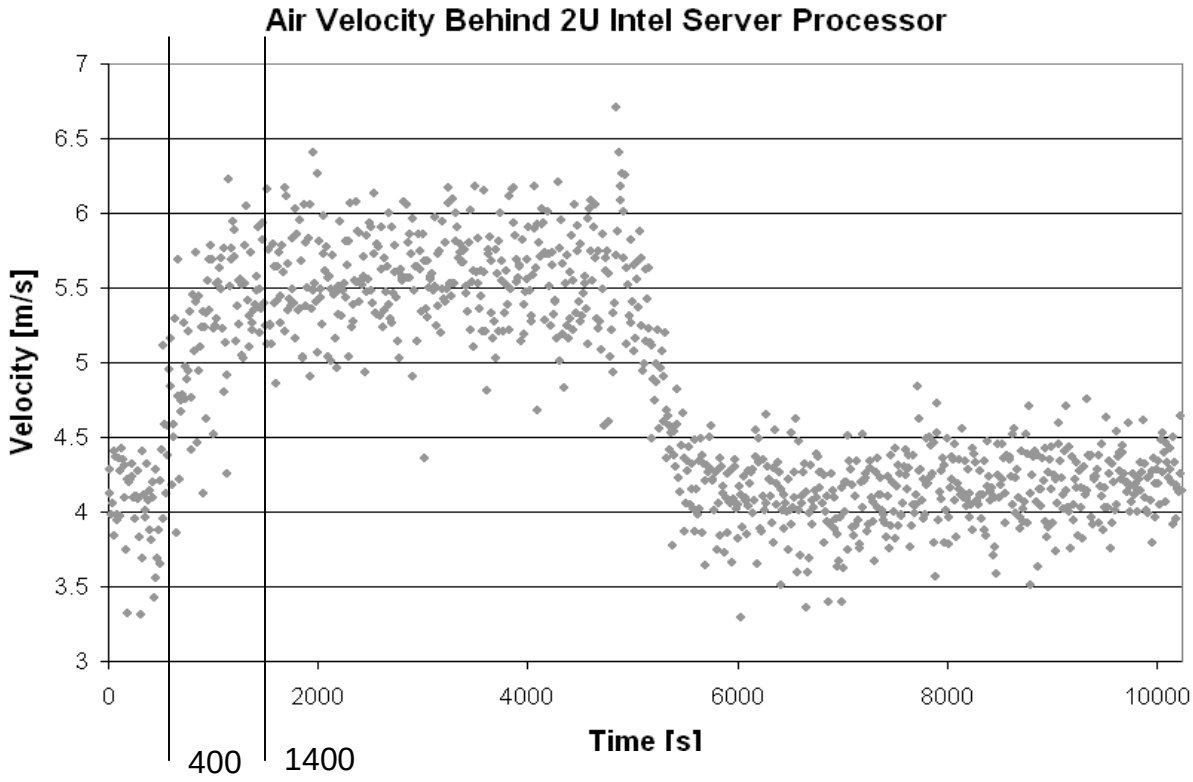

Figure 3.4-8: Air velocity at outlet of Intel server processor air path.

3.4.4 Comparison with Heating Element Time Constant as Lower Limit

Figure 3.4-9 shows the air temperature response at the server simulator outlets to a step change in inlet air temperature. The time constant for the server simulator is about an order of magnitude lower than that of the servers at around 50 seconds. The composition of the server simulator heating element is shown in Figure 3.4-10; the element contains nichrome wires surrounded by ceramic insulation and a steel heat sink.

A diagram of a typical CPU package is included as Figure 3.4-13 and pictures of the CPU package from the legacy machines are include as Figure 3.4-11and Figure 3.4-12. Because of the materials and geometry involved, it is not likely that any servers will have a time constant as low as the server simulator. Thus its time constant might be seen as a lower limit on server time constants. Even this short time constant is long enough to have an effect on temperature rise during a power failure with events taking

place in about 30 seconds (Emergency generator startup will cause AH fans and CHW pumps to operate in approximately 30 seconds after power failure.).

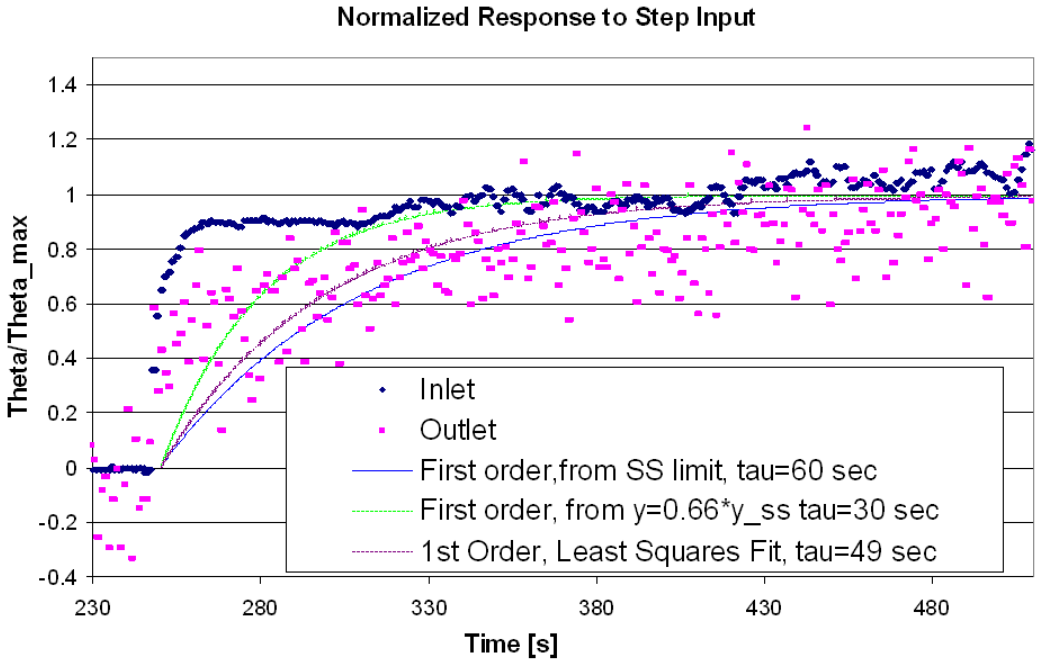

Figure 3.4-9: Response of server simulator to step change in inlet temperature.

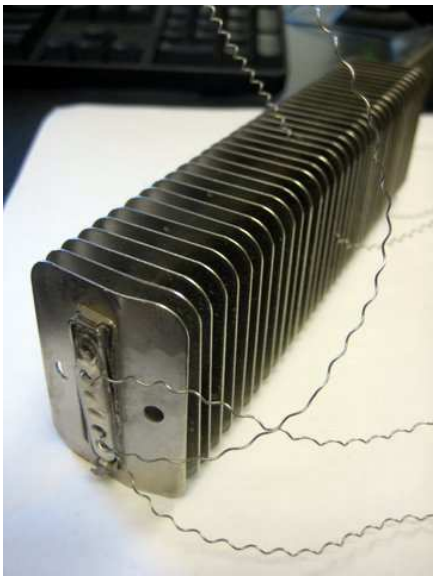

Figure 3.4-10: Server simulator heating element.

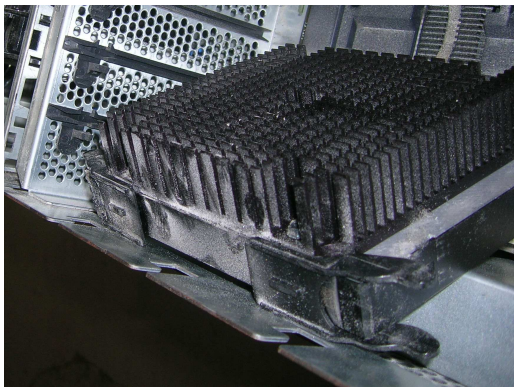

Figure 3.4-11: Legacy server processor package.

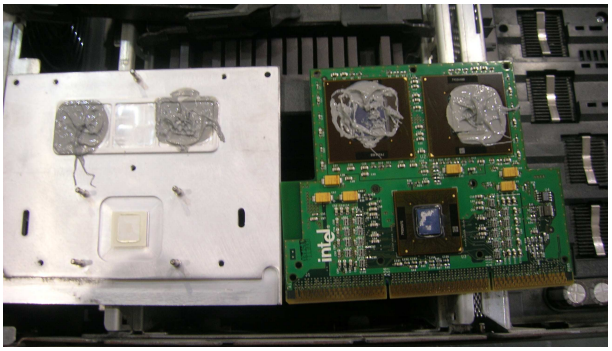

Figure 3.4-12: Legacy server processor circuit board.

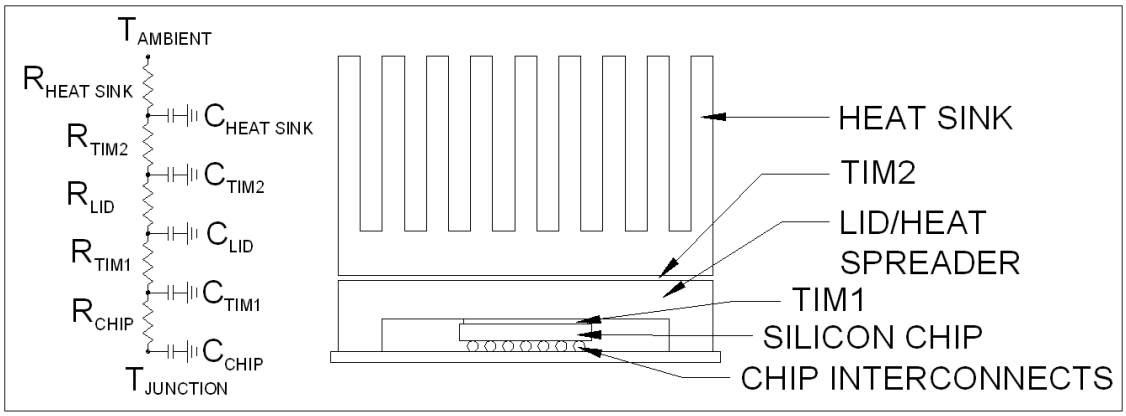

Figure 3.4-13: Thermal capacitance of materials in heat transfer path from CPU to air.

3.4.5 Implications of Relatively Long Server Time Constant

When compared with transient events predicted to occur in 11 seconds, or even 1 minute, the thermal capacitance of a server with a time constant of 2.5-6.7 minutes is substantial. These experiments show that some steady state based CFD/HT models may be predicting failure before the server has a chance to approach steady state. By storing energy within the materials of the server, the actual compute equipment delays the temperature rise of the outlet air. Thus, the room air temperature and the server inlet air temperature will rise more slowly. This amendment to the model will be explored more fully in the control volume experiments in Chapter 5.

Unless appropriate modeling methods for predicting server time constants become available, it would be necessary to repeat the above experiment for each model of server. This would be a simple test for a server manufacturer, but is not likely to take place in a data center environment. However, if enough data becomes available it may become possible to make reasonable guesses at server time constants based on the characteristics of the server.

4 Control Volume Experiment

4.1 Experimental Setup

Data center heat loads are often measured in heat dissipation per unit of floor area. In order to maximize the heat dissipated per unit of floor space, a smaller volume of air was separated from the rest of the laboratory. In addition, it was unacceptable to subject the other compute equipment housed within the laboratory to the conditions experienced by the legacy servers. Although airflow patterns were brought closer to those that would be experienced under the hot-aisle/cold-aisle configuration, predicting flow patterns for a typical data center is not the purpose of the model.

Although complete thermal isolation of the control volume was impossible, as much isolation as practical was attempted. The control volume was built from 12 mm (nominal ½") thick expanded polystyrene foam board, which had an R-value of R-3. However, where the control volume coincided with the walls of the data center, they were incorporated into the control volume. Above the raised floor, the foam board was supported by framing made from 38 mm x 89 mm (nominal 2" x 4") studs set at 610 mm (24 in) on center. Within the raised floor plenum, the foam board was supported by the stanchions that also support the raised floor; the stanchions were also set at 610 mm (24 in) on center. Diagrams of the control volume are shown in Figure 4.1-1, Figure 4.1-2, and Figure 4.1-3, and photos are shown in Appendix A.5. The server position designations shown in Figure 4.1-3 are used to show where temperature measurements were acquired in the data sets in the following sections.

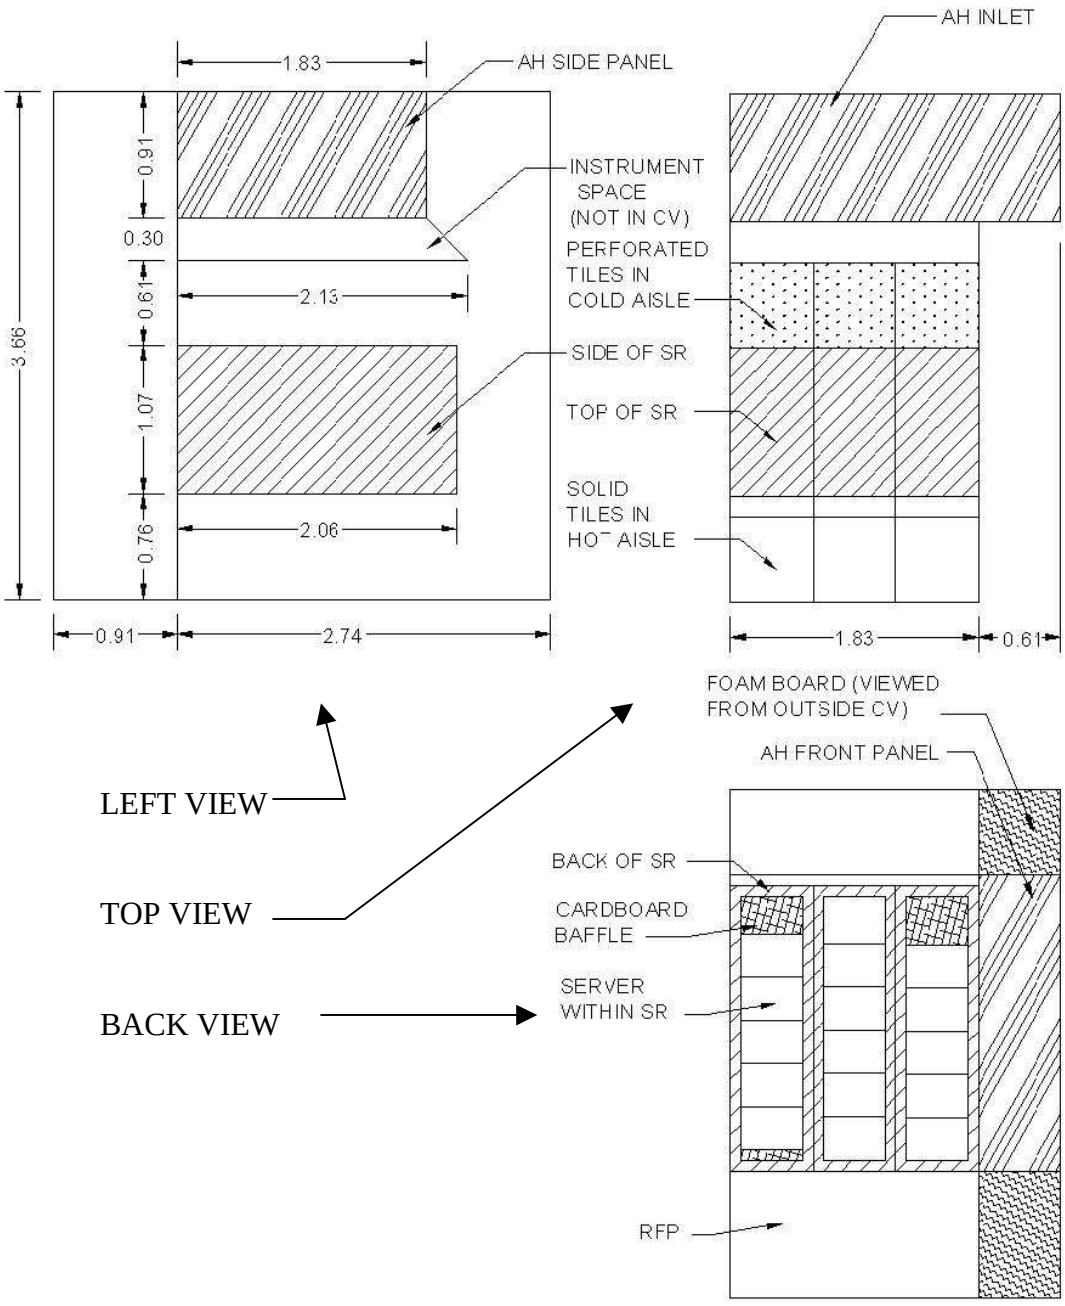

Figure 4.1-1: CV layout and dimensions.

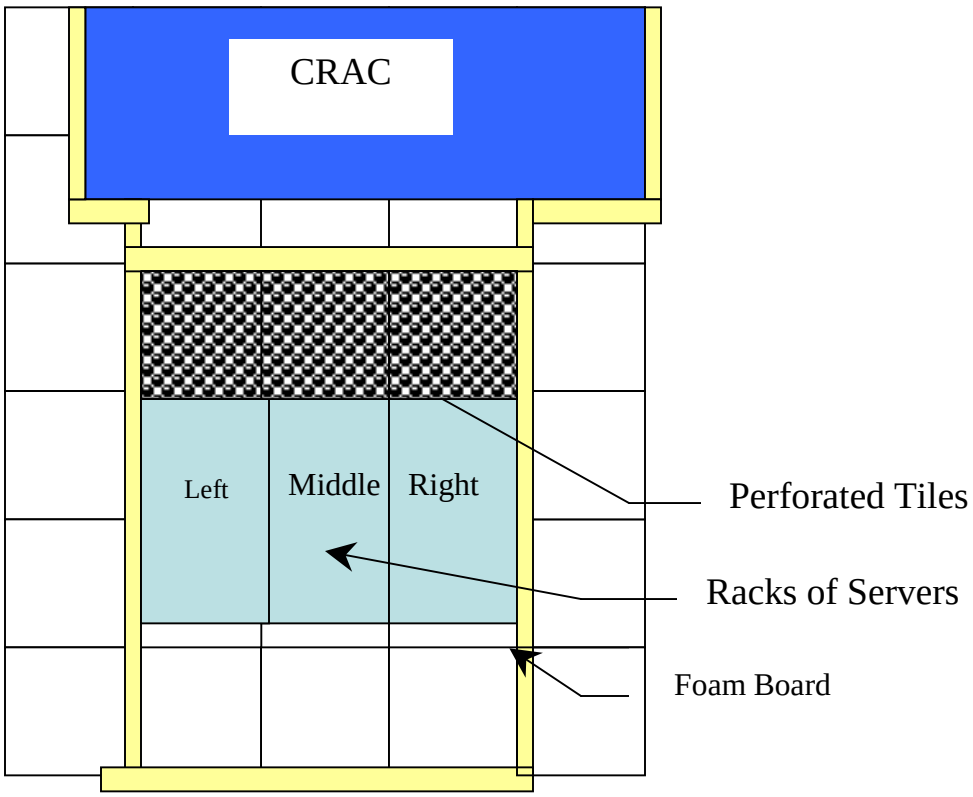

Figure 4.1-2: Schematic depiction of control volume floor plan.

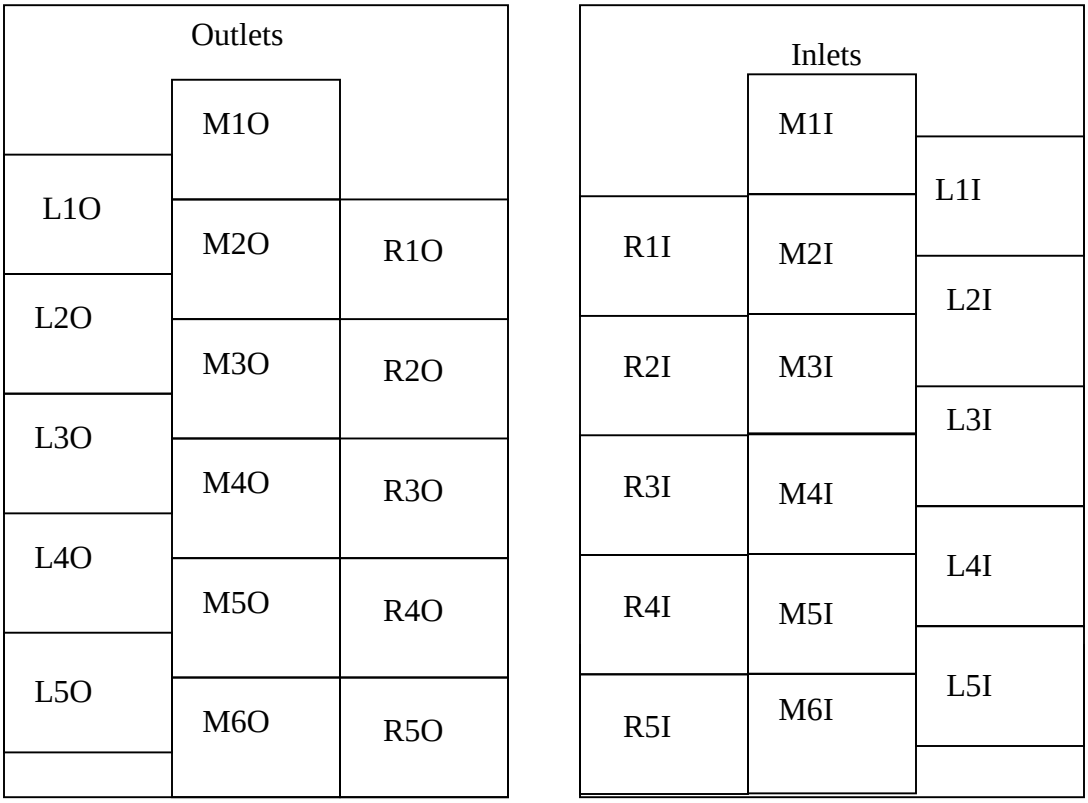

Figure 4.1-3: Designation of server outlets and inlets by rack and vertical location.

The control volume allowed experimentation with conditions that would not have been possible in the data center at large, but it also suffered from some inherent difficulties. The most important of these was that the AH was designed for a much larger load and air volume. The AH had a variable frequency drive (VFD) to vary the fan speed, but this fan speed was not independently controllable by the user. Instead the AH had its own control system that was designed for a larger volume of air. Therefore, when used for such a small control volume, the AH control system modulated the fan to a speed that was toward the low end of the range at the room level, but very high for the control volume. The result was high pressure differences between the control volume and the outer data center air, leading to leakage and higher than desired velocities from the perforated tiles. In addition to leakage and infiltration of air, conduction across the foam board was a source of discrepancies between the model and the control volume. It should be noted that the chilled water temperature experiences oscillations due to intermittent operation of the chiller, making it impossible to achieve true steady state operation.

Prime95 was used to load the processors of all 16 legacy servers. The results of the control volume experiment are included in the following sections. Measurements were taken with the same type-T thermocouples used in the server time constant experiment; they were attached to server inlets and outlets, as well as other areas of interest. A handheld thermal anemometer (See Figure A 5) was stabilized via an instrument stand and used to estimate the face velocity at the AH fan outlet. This reading represents a point and should not be interpreted as the average velocity at the fan outlet.

The data show that the point velocity at the AH outlet typically remains around 20 m/s (4000 FPM), but can reach over 46 m/s (9000 FPM) when the temperature of the CHW rises. This is only possible because the AH fan motor is equipped with a variable frequency drive (VFD).

Figure 4.1-4: Gap in control volume enclosure caused by large negative pressure at air handler inlet requiring repair to obtain usable results.

4.2 Control Volume Experiment: Sources of Uncertainty

The CV experiments suffered from uncertainties due to measurement equipment, but more prominently from inability to completely isolate the control volume from the surrounding environment. The entire temperature measurement system was calibrated to yield a total uncertainty within ± 0.3°C.

Uncertainties originating from inability to completely isolate the control volume were more difficult to quantify and depend on the temperature difference across the control volume boundary. A MATLAB script was written to estimate these and is included in Appendix A.6. This script suffers from the common necessity of assuming steady state conduction across the boundary. At a temperature difference of 10°C, loss of heat through conduction is estimated at less than 1 kW. At around 30°C, this loss exceeds 2 kW. Both values are an order of magnitude lower than the heat being dissipated by the server CPUs. However, this estimate does not take infiltration into account.

An attempt has been made to quantify infiltration within the MATLAB script in Appenidx A.6. The positive pressures within the raised floor plenum and negative pressures within the control volume above floor are predicted by the CFD/HT model to be on the order of 5 kPa. With leakage, this value will be lower in practice. Nevertheless, the ASHRAE chart referenced does not provide data beyond 70 Pa at best. Although every attempt has been made to seal the control volume, it is difficult to determine what quality of construction should be chosen. In light of these, the highest infiltration value on the chart is used for the calculations and the results should be viewed as an order of magnitude estimate. The script predicts a heat loss of 2 kW with a 10°C temperature difference across the control volume boundary and 6 kW loss with a 30°C temperature difference.

The total heat loss from infiltration and conduction predicted by the script ranges from 3 kW at a temperature difference of 10°C to 8 kW at 30°C. The experiment begins with the interior of the control volume about 6°C lower than the surroundings and ends at about 30°C higher; the beginning of the experiment therefore does not see large errors produced by incomplete isolation, but the error increases as the temperature inside the control volume rises. This effect is apparent in the data, because the slope of the temperature rise flattens as the temperature rises.

4.3 Control Volume Experiment 1: Mimicking AH Power Loss

The first control volume experiment mimics the scenario in which UPS backup power infrastructure is provided for only the compute equipment. The experiment is performed by allowing the control volume to come as close as the chilled water oscillations will allow to steady state and then turning off the AH. This allows the servers to heat the room air. The experiment was stopped when one of the server inlet temperatures reached about 45°C. This sequence of events is illustrated by Figure 4.3-1.

Figure 4.3-1: Timeline of events during control volume experiment 1.

A few general comments are possible about general temperature trends during the experiment from Figure 4.3-2. First, most temperatures rise significantly after the AH is switched off; those that do not are measuring temperatures of the outer surface of the control volume wall and air outside the control volume. Even the temperature of the chilled water manifolds on the AH HX rises by more than 5°C despite the fact that the fans are turned off. Another general trend worth noting is that slope of the inlet and outlet temperatures decreases gradually over the course of the experiment. This effect corresponds to the inability of the control volume to completely isolate the air inside against conduction and infiltration; as described previously, the errors introduced by incomplete isolation increase as the temperature inside the control volume increases.

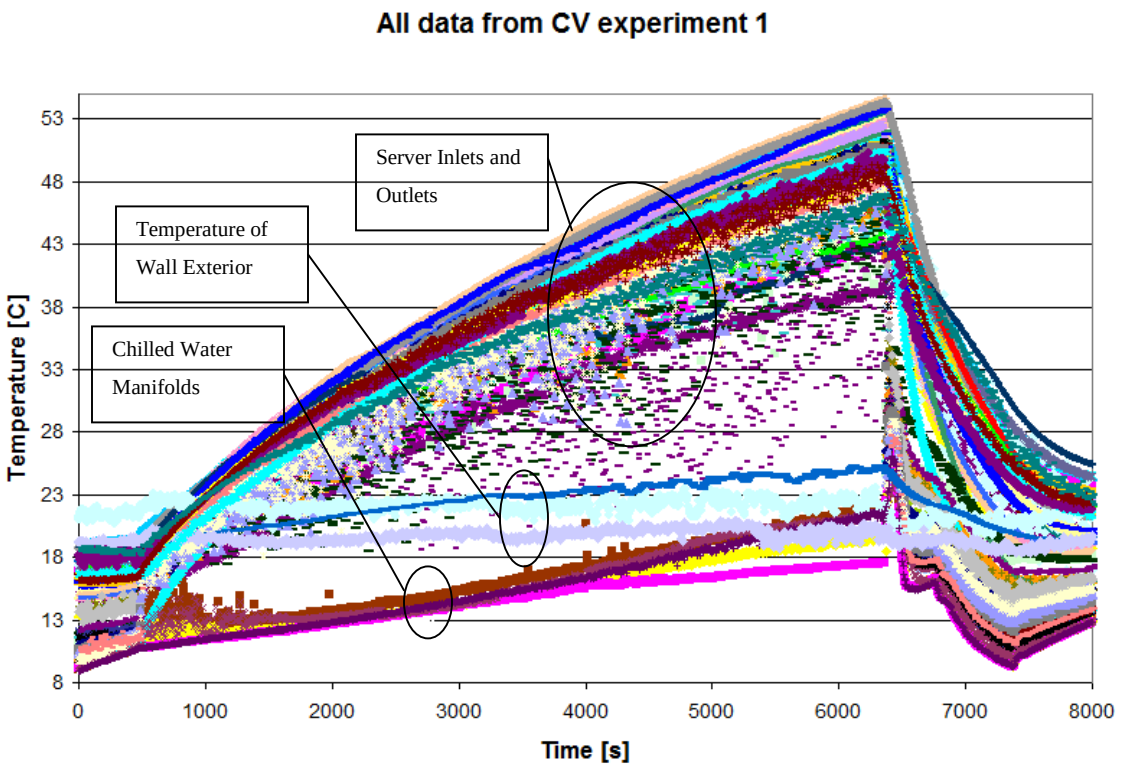

Figure 4.3-2: Overview of data from control volume experiment 1.

Figure 4.3-3 shows all server inlet temperatures from control volume experiment 1; the legend follows the indexing structure shown in Figure 4.1-3. A significant upward jump in inlet temperatures can be observed jut after the AH is switched off. This is caused by suction of warm air originating from the server outlets into the server inlets; the pressurization of the cold aisle by the AH had not allowed this before it was switched off. The figure shows that inlets near the top of the racks follow a smooth, steadily increasing trend. However, some of the servers at the bottom show temperature trends with large scatter. The airflow patterns were established by the AH fans until it was switched off; thus, the air was forced into the raised floor plenum, up through the perforated tiles, and was then drawn in by the server fans. After the AH is switched off the flow is driven by the server fans; thus, air is not forcibly circulated through the raised floor plenum. Instead it moves from server outlets, over the top of the server rack, and back to the server inlets. Since air from the raised floor plenum is not forcibly mixed with the air above, cold air remains just below the perforated tile. During

the course of the experiment, unsteady flows alternately cause the servers to draw warm air originating from the server outlets and cool air originating from the raised floor plenum. The poor mixing of air from the raised floor plenum with that above floor has been predicted in a previous model. [25]

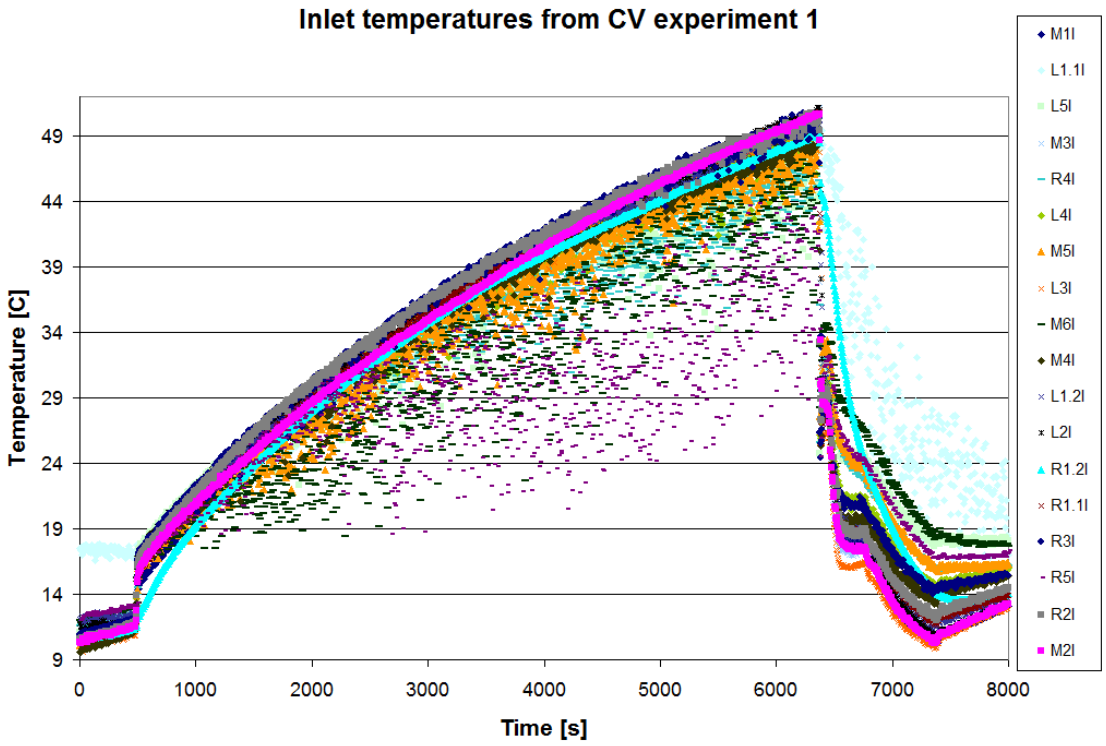

Figure 4.3-3: Inlet temperatures from control volume experiment 1.

Figure 4.3-3 shows inlet temperatures from the middle rack. Inspection of the data shows that temperatures near the top of the rack are generally higher than those near the bottom, but not always sequentially decreasing. Also somewhat unexpected is the narrow range of server inlet temperatures experienced at any given point in time. Save the very lowest servers, which receive air intermittently from the raised floor plenum, the

temperature responses remain within a few degrees of one another.

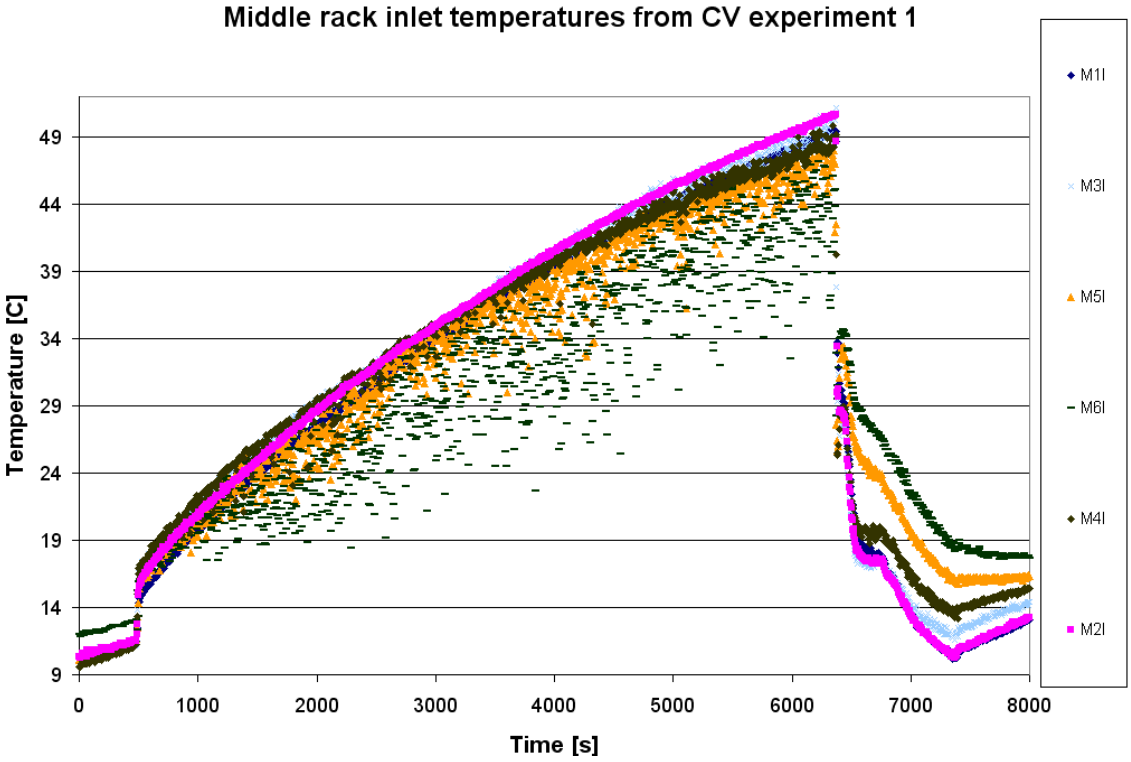

Figure 4.3-4: Inlet temperatures from middle rack of control volume experiment 1.

Figure 4.3-5 shows outlet temperatures from control volume experiment 1. Temperatures are generally lower toward the bottom of the rack than the top. The significant fluctuations in temperature observed at the inlets to the lower bottom servers are not observed at the outlets. Apparently, the air has been well mixed within the server and exits with a uniform temperature.

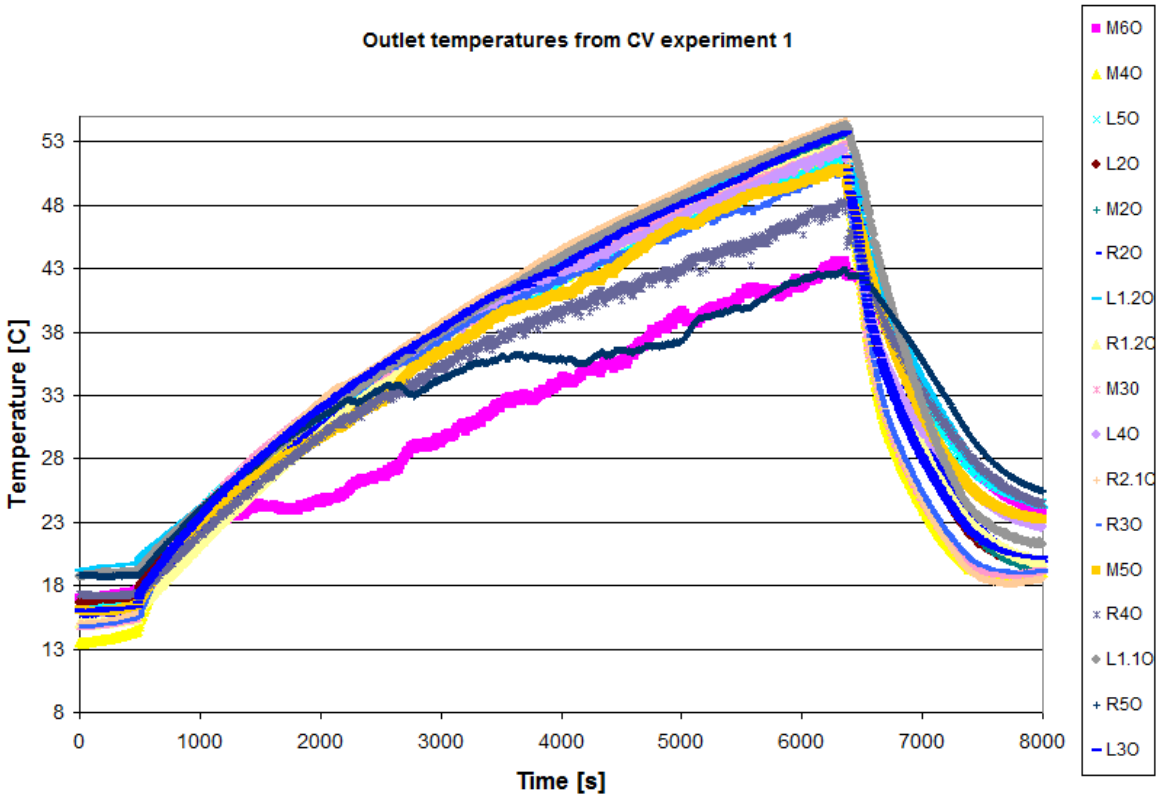

Figure 4.3-5: Outlet temperatures from control volume experiment 1.

Figure 4.3-6 shows outlet temperatures from the middle rack during control volume experiment 1. The lowest server outlet temperatures stand out as lower throughout the experiment. Other responses are not so simply described. The higher servers remain within a narrow range of a few degrees from one another, but their relative rankings vary. The uppermost server never has the highest temperature, as would be expected based on recirculation patterns.

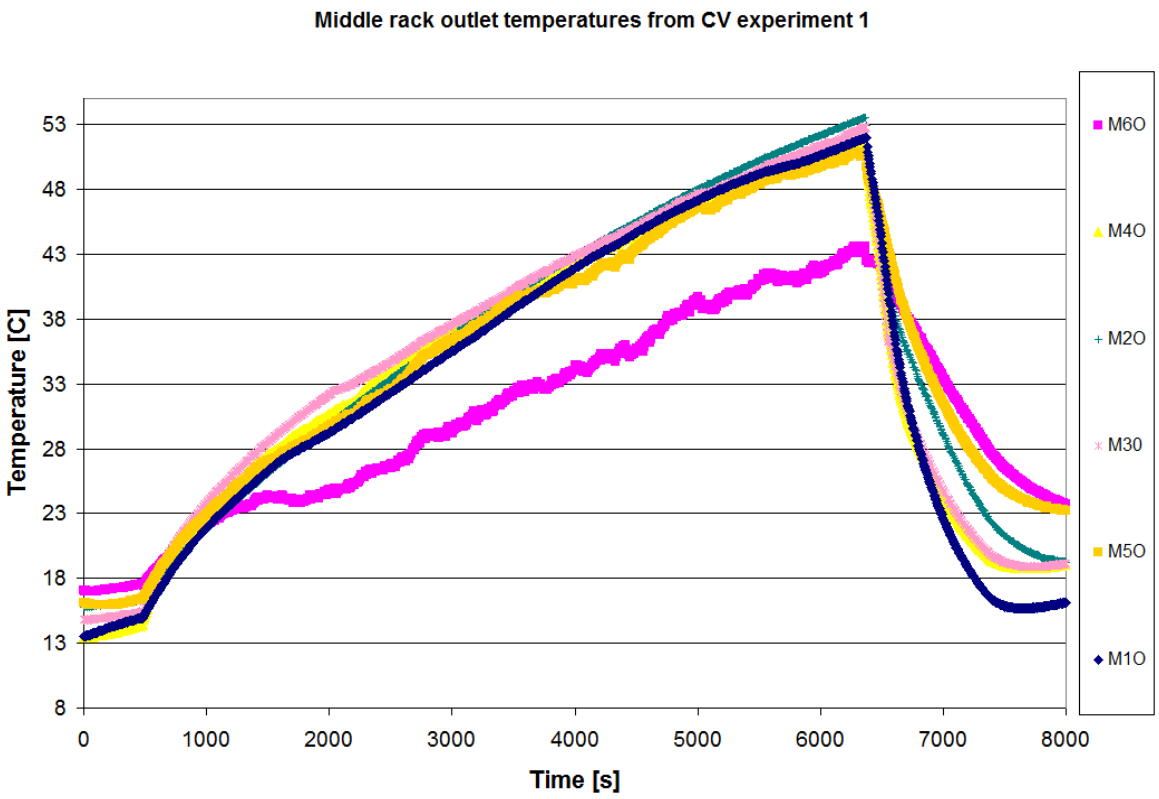

Figure 4.3-6: Middle rack outlet temperatures from CV experiment 1.

Figure 4.3-7 shows the average of all server inlet temperatures, the average of all outlets, and the average temperature difference computed from these curves. If all the servers have similar airflow rates, then this represents the mass averaged temperature difference across the servers. Just after the AH is switched off, the temperature difference across the servers drops from 3.8°C to nearly 0°C; it rises during the course of the experiment, but only reaches 2.8°C by the time it is switched back on over 2.5 hours later. Since the temperature difference drops, it can be inferred that energy is being stored within the materials of the server.

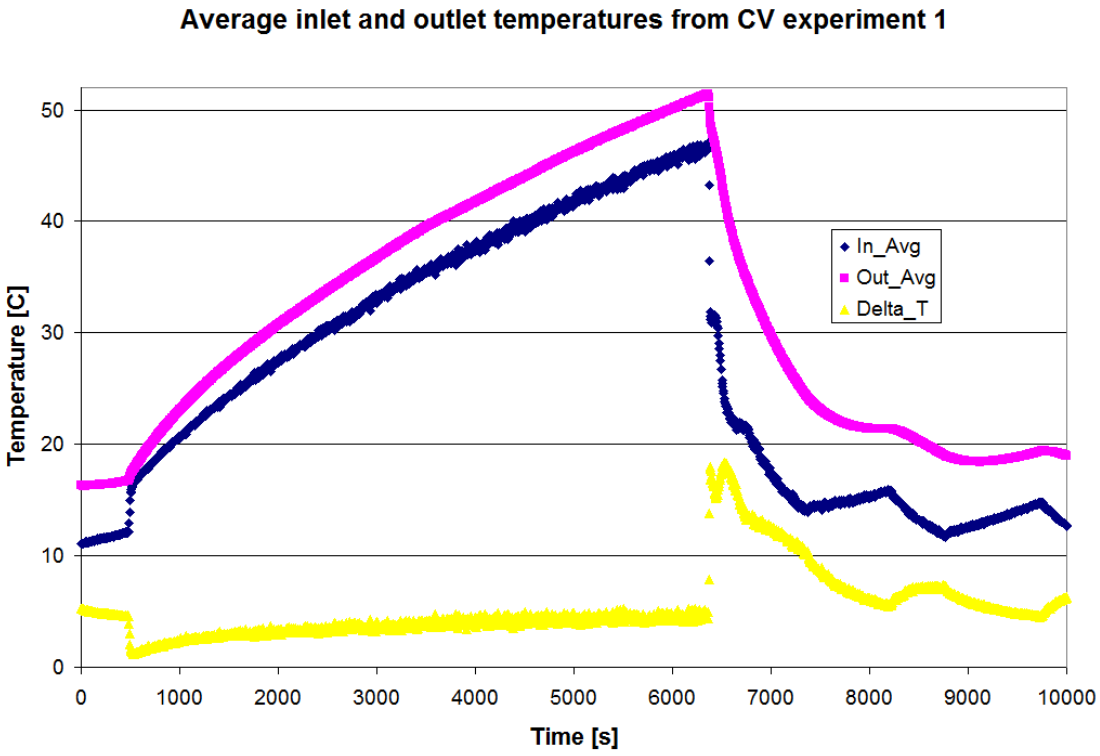

Figure 4.3-7: Average temperature rise across servers in CV experiment 1.

Figure 4.3-8 shows colormaps representing temperatures at various server inlets over the course of control volume experiment 1. The first pair of colormaps is taken at steady state operation; the greatest variations in temperature are among the server outlets. The top right server is the warmest. As the experiment progresses, the inlet temperatures experience a much larger degree of variation. The temperatures of the bottom servers may reflect fluctuations in temperature due to turbulent mixing rather than the general trend.

Inlets **Outlets**

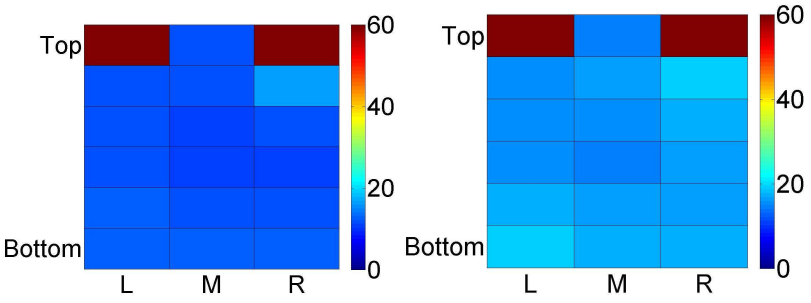

t_0 = 435 s

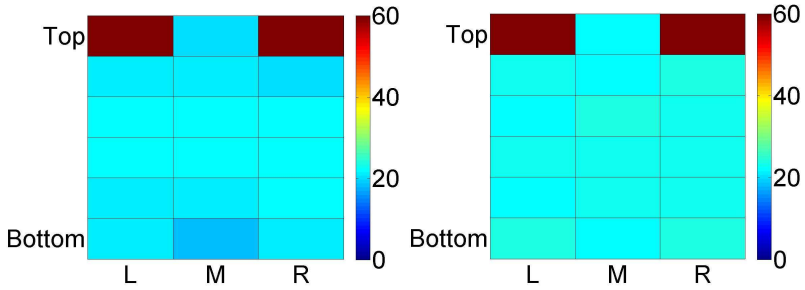

t = 1000 s

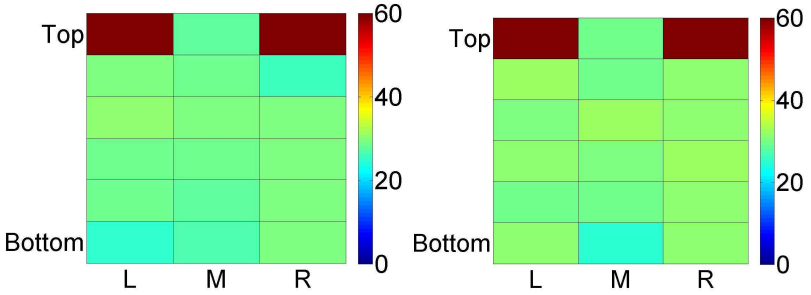

t = 2000 s

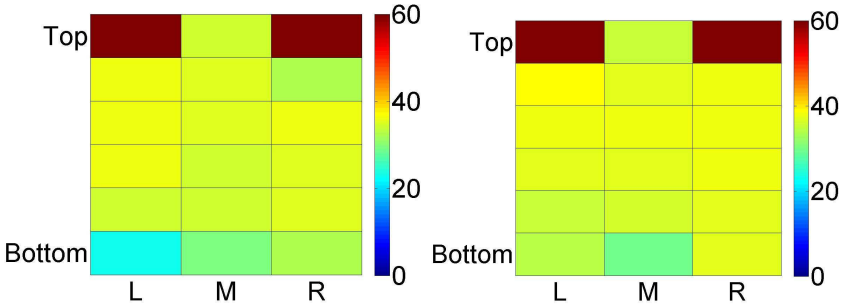

t = 3000 s

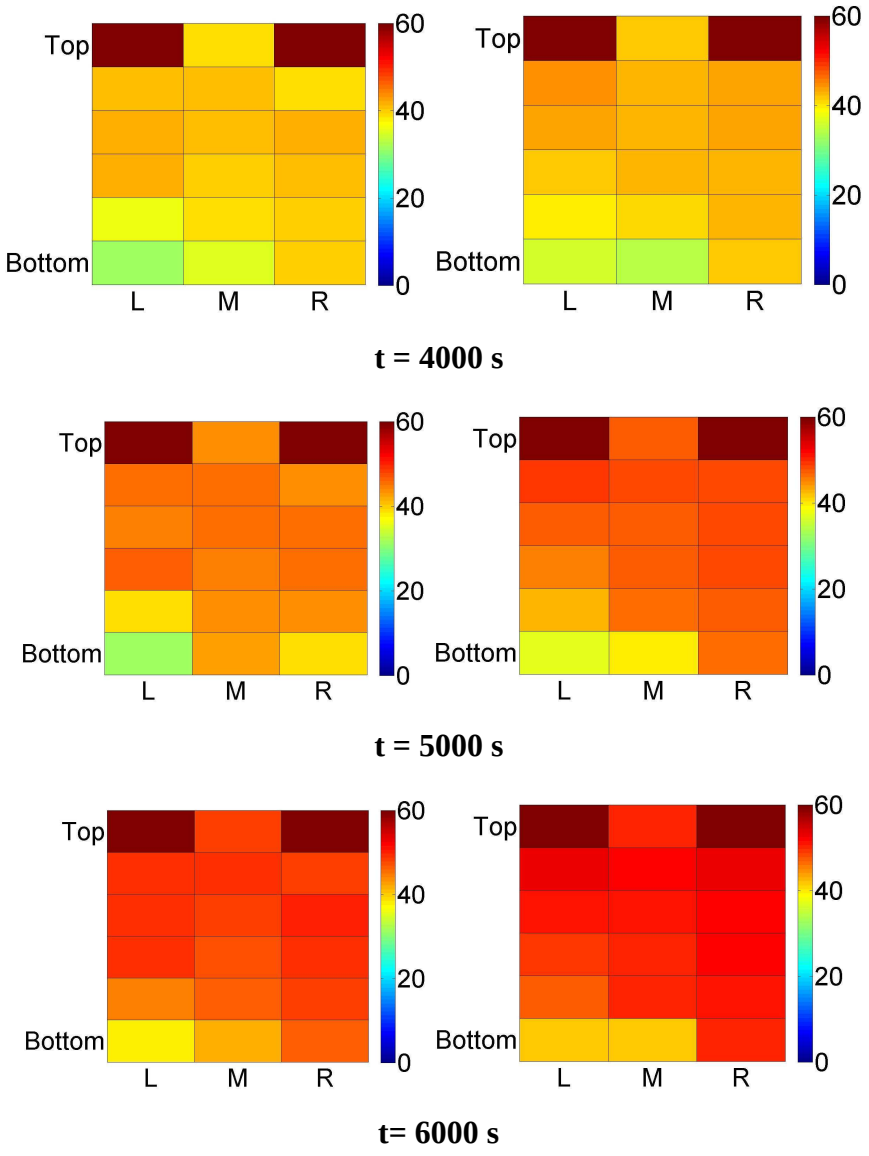

Figure 4.3-8: Colormap of inlet and outlet temperatures during CV experiment 1.

4.4 Control Volume Experiment 2: Mimicking Pump Power Loss

Figure 4.4-1 shows the timeline of events for control volume experiment 1. The second control volume experiment mimics the scenario in which UPS backup power infrastructure is provided for compute equipment and AH, but not CHW pumps. Once again, the control volume is allowed to come to an initial quasi-steady state. However, in this case, the AH is not switched off; instead, the CHW valve is closed. The servers still

heat the air within the control volume, but now the air flow patterns remain relatively the same and air is circulated through the plenum and control volume at high flow rates. Therefore, it takes much longer for the compute equipment inlet temperatures to reach unacceptable values. In fact, the temperature within the control volume rises more slowly as the experiment progresses. This phenomenon could be attributed to increased heat transfer across the walls of the control volume; both conduction and infiltration would be substantially increased by the high airflow rates. However, there could also be some leakage through the AH CHW shutoff valves if positive shutoff is not achieved.

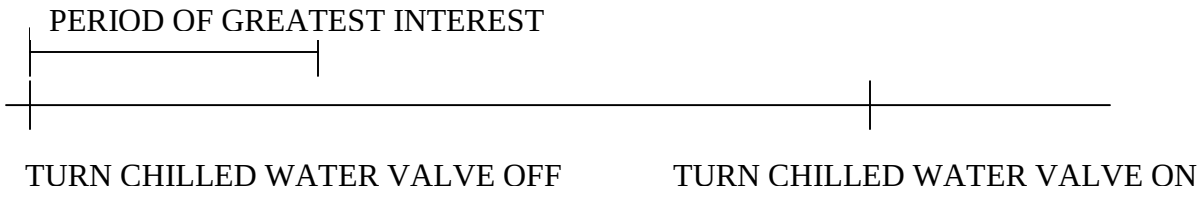

Figure 4.4-1: Timeline of events during CV experiment 2.

Figure 4.4-2 shows all the temperature data collected from control volume experiment 2. The experiment takes place over a larger timescale than the control volume experiment1 because it took much longer to approach the limiting temperatures at the server inlets. Most of the same general trends with the data are observed as in control volume experiment 1, but some effects are more or less exaggerated. In particular, the decrease in the slopes of the temperature rise curves is much more obvious. The wall temperature outside the control volume changes slightly over the course of the experiment—an effect not previously observable and likely due to increased convection to the foamboard. Since air is driven through the AH HX, the chilled water manifold temperatures increase much more dramatically in this experiment.

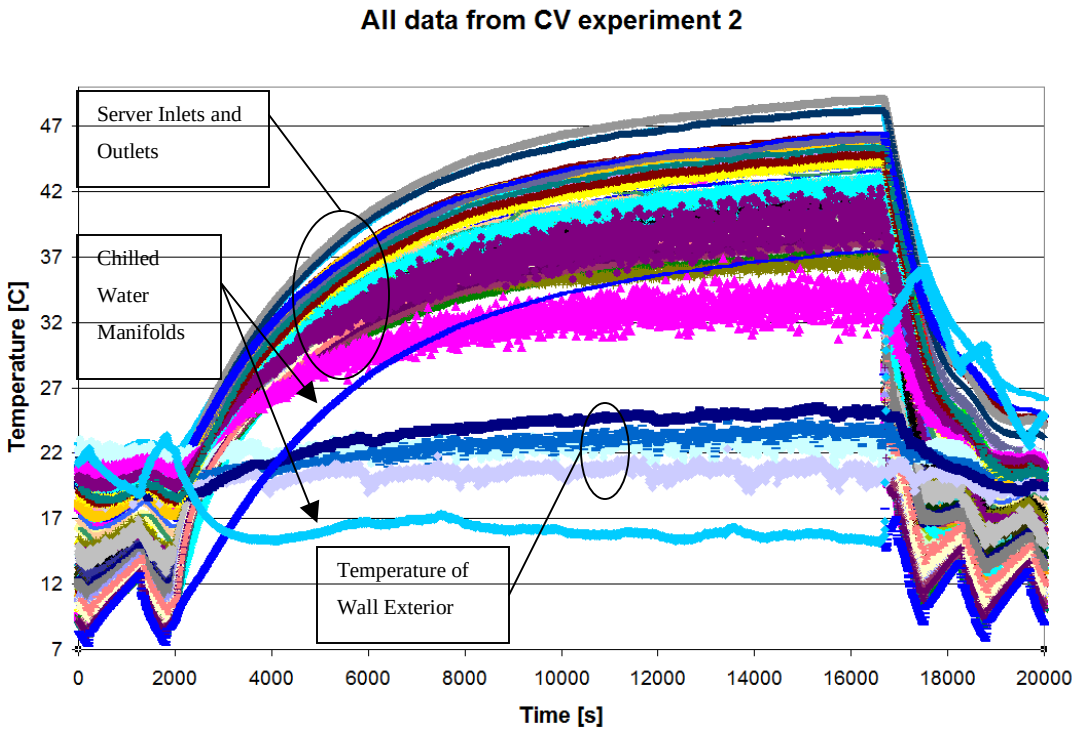

Figure 4.4-2: All data from CV experiment 2.

Figure 4.4-3 shows all server inlet temperatures from control volume experiment 2. There is a marked difference between the inlet temperature trends from control volume experiment 1 and 2; none of the server inlet temperature trends experience significant scatter. The lack of scatter is can be attributed to the fact that the AH fans keep blowing and maintain essentially constant airflow patterns throughout the course of the experiment.

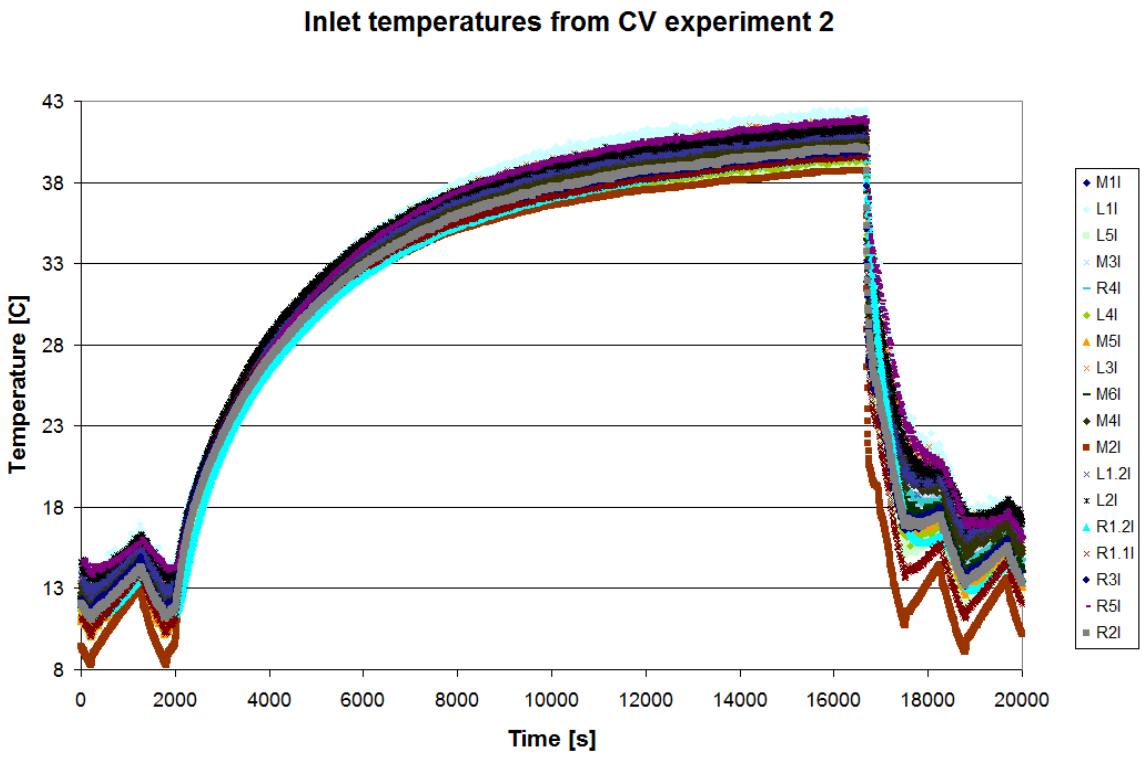

Figure 4.4-3: Inlet temperatures from CV experiment 2.

Figure 4.4-4 shows inlet temperatures from the middle rack during control volume experiment 2. The range of temperatures during the period of interest, where the slope is the greatest right after the valve is closed, are within about 1°C of one another. Therefore it is not relevant to talk about the temperature trends relative to one another during this time. Even after the period of interest, the temperature trends are all very close to one another.

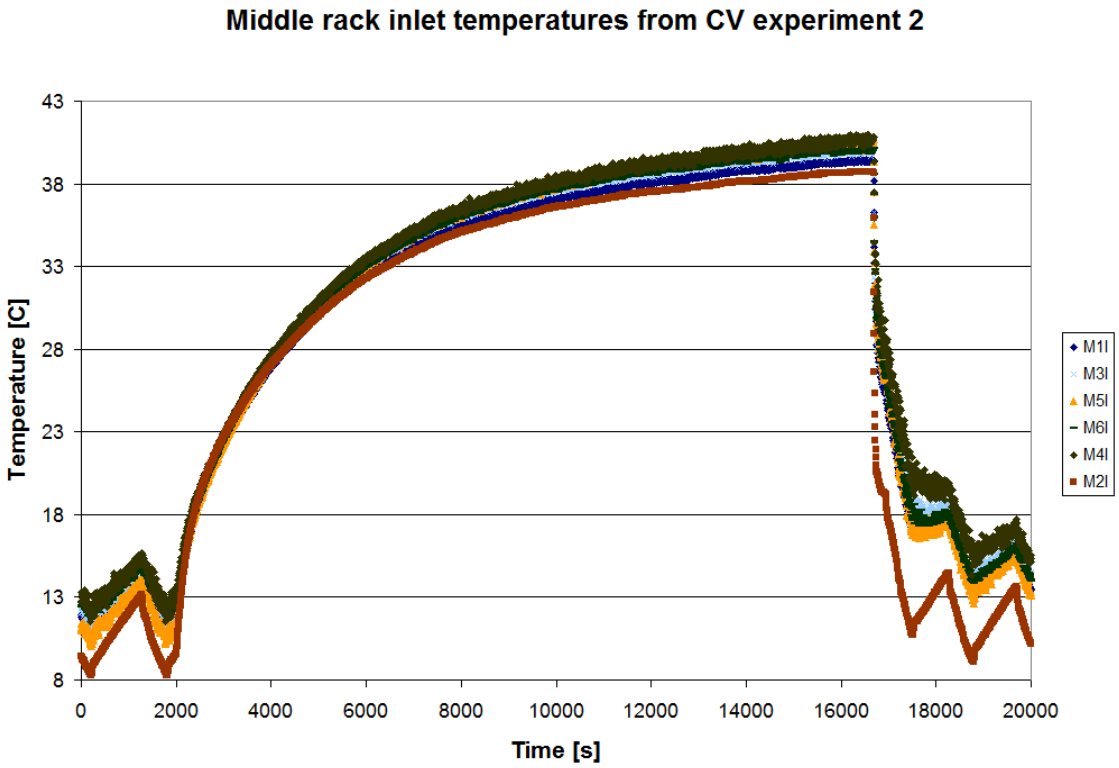

Figure 4.4-4: Middle rack inlet temperatures from CV experiment 2.

Figure 4.4-5 shows outlet temperatures from control volume experiment 2. After the valve is closed, the outlet temperatures move at first toward a narrow range, but the range of temperatures increases somewhat toward the end of the experiment. It is interesting to note that the highest temperatures are not necessarily at the highest part of the rack; the top racks have the highest temperature on the left and the bottom on the right. This shows that more than just recirculation patterns influence server cooling; pressure variations among the server inlets, when high velocity air by-passes the inlets, also inhibits cooling. Also worth noting is that the highest temperatures before the period of interest are not necessarily the highest afterwards.

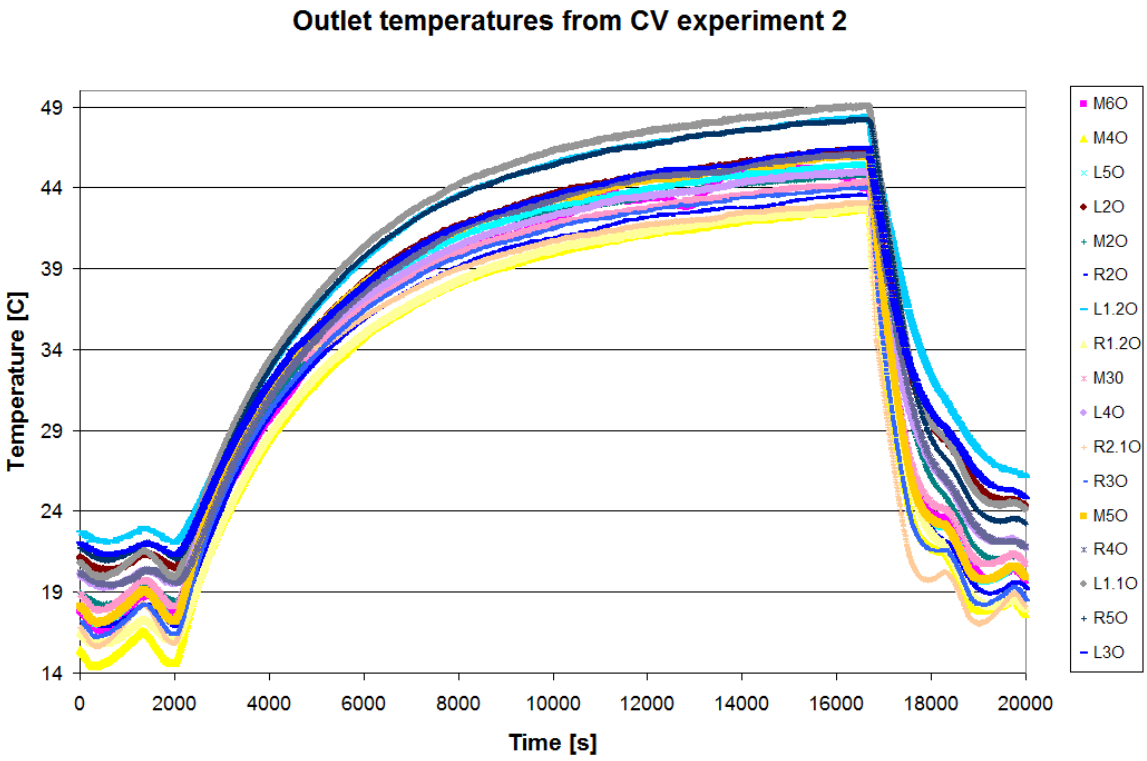

Figure 4.4-5: Outlet temperatures from CV experiment 2.

Figure 4.4-6 shows once again that the relative temperatures of servers within the rack do not necessarily decrease sequentially from top to bottom. Figure 4.4-7 shows the average temperature of server inlets, server outlets, and the average temperature rise across the servers during control volume experiment 2. The temperature difference curve does not drop as sharply during the period of interest as in the first control volume experiment, but this is likely due to the fact that the servers do no start from true steady state operation before the experiment. The temperature difference increased just before the valve was closed. Therefore the drop was not as prominent as in control volume experiment 1. Since the servers were releasing stored heat at the time when the valve was closed, the sharp dip in temperature rise from the first control volume experiment is obscured.

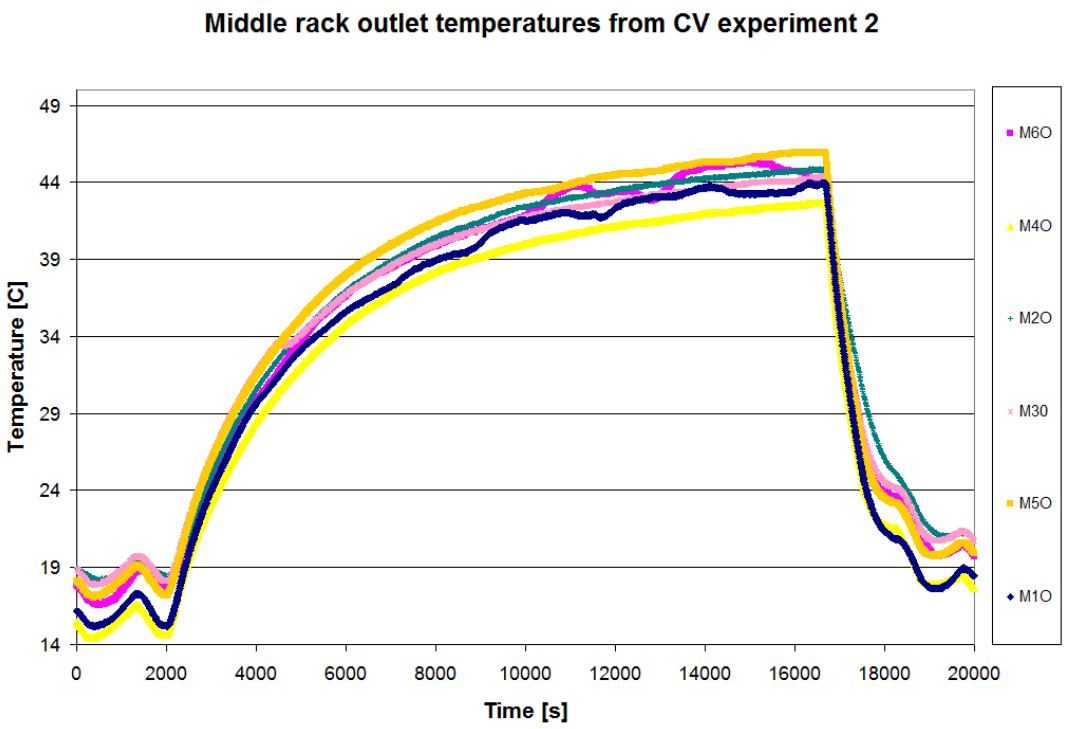

Figure 4.4-6: Middle rack outlet temperatures from CV experiment 2.

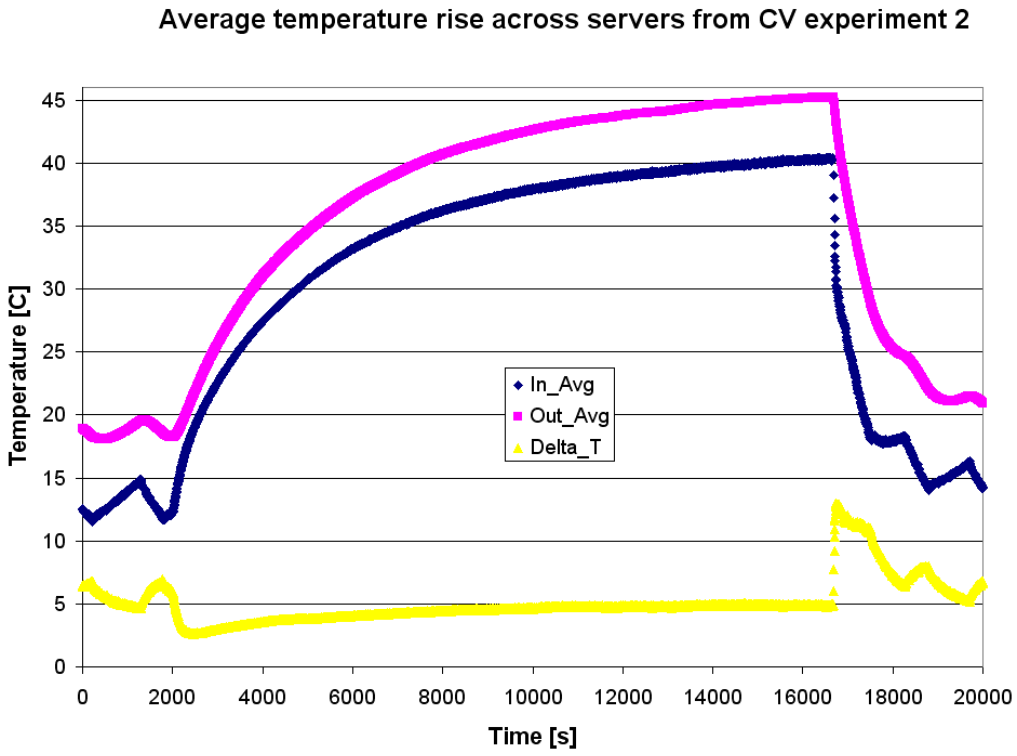

Figure 4.4-7: Average temperature rise across servers.

Figure 4.4-8 shows a series of colormaps representing server inlet and outlet temperatures at various servers. It can be observed that the variation in inlet temperatures is much less than the variation in outlet temperatures. Although the entire cold aisle is pressurized with cold air and blocks warm return air from being recirculated to the server inlets, this ensures only proper inlet temperatures and not proper airflow. Some of the servers show higher temperature trends; it appears that air with a high vertical velocity may have a low dynamic pressure and, thus, the fans cannot draw enough air.

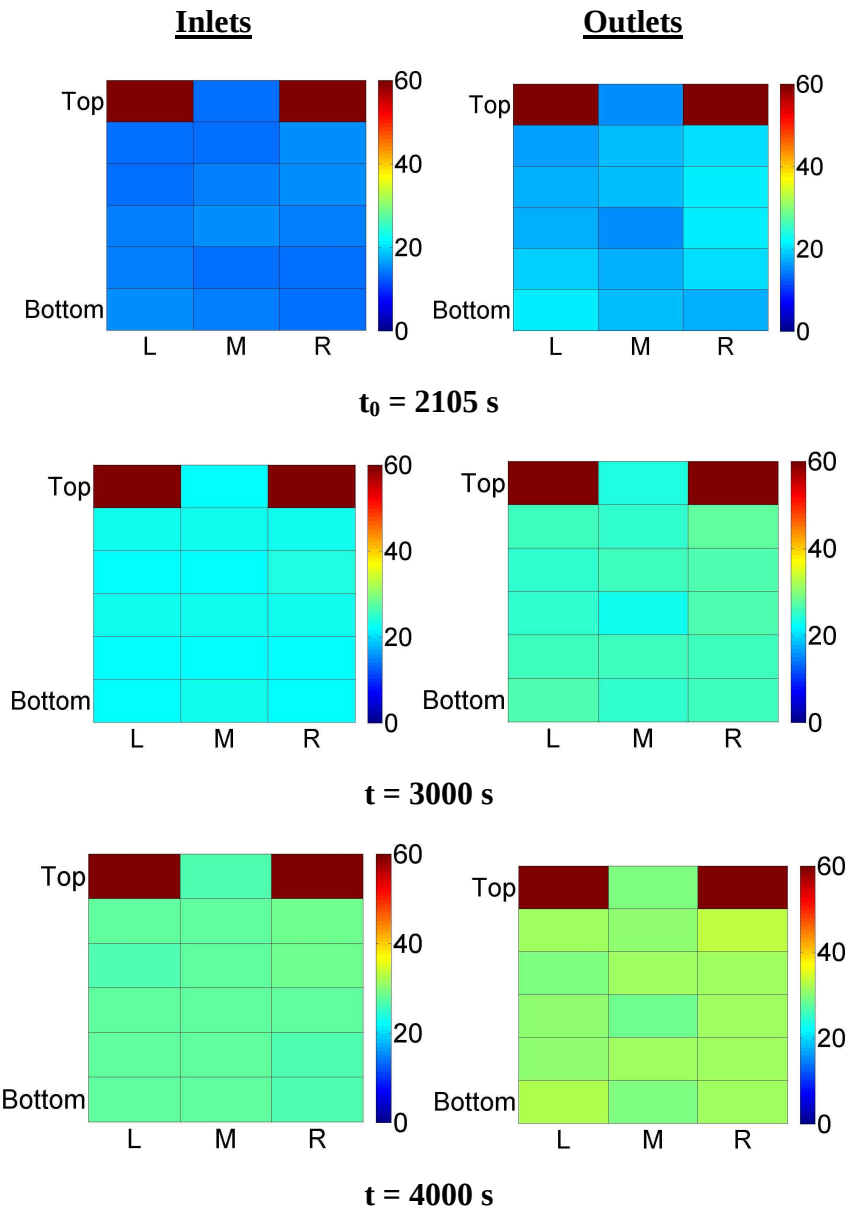

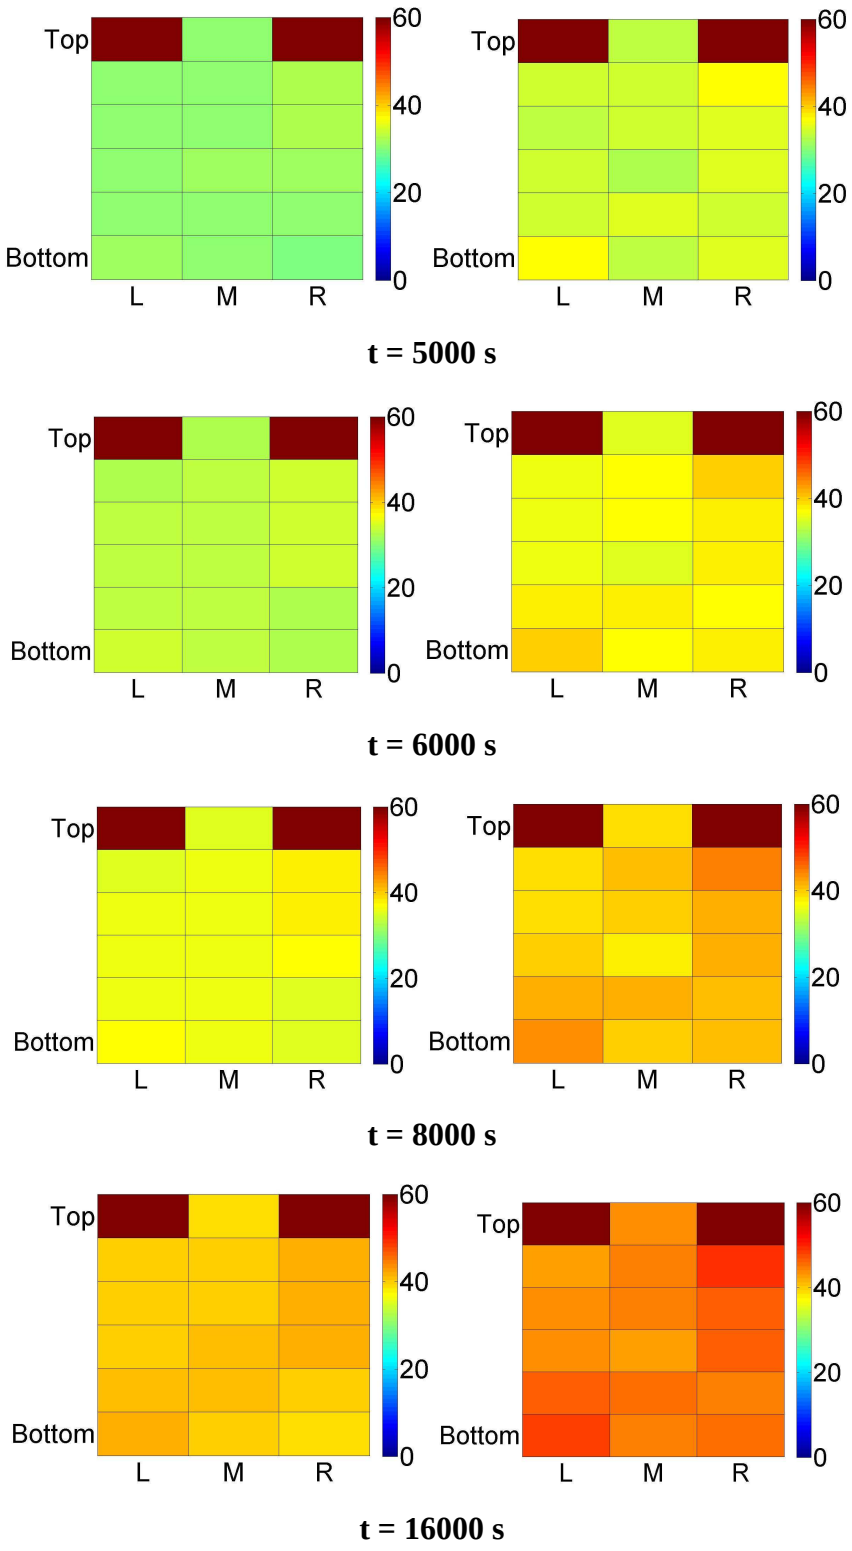

Figure 4.4-8: Colormap of inlet and outlet temperatures during CV experiment 2.

4.5 Control Volume Experiment 3: Mimicking Chiller Power Loss

Figure 4.5-1 shows the timeline of events during control volume experiment 3. The third control volume experiment mimics the scenario in which UPS backup power infrastructure is provided for compute equipment, AH, and CHW pumps. With sufficient supply of CHW storage, the data center should remain at steady state for as long as needed to restore power. However, The CHW loop in the CEETHERM laboratory is only about 37 m (120 ft) m long and has a diameter of 76 mm (3 in). There is also a 757 L (200 gal) storage tank in the CHW loop, yielding 840 Liters (220 Gallons) of CHW. This would be a relatively small amount of CHW for the entire data center, but experimental results show that the temperature rise is substantially slowed when cooling is supplied only to the control volume (The other compute equipment was shutdown for the experiment.).

Figure 4.5-1: Timeline of events during CV experiment 2.

Figure 4.4-2 shows the temperature data taken during control volume experiment 3. While the temperature of the room air outside the control volume remained relatively constant during the previous control volume experiments, here it rises significantly. This effect could have been caused by losses of energy from the control volume to the room, or by heating of the room air by the small fraction of CE within the room without cooling, or some combination of the two.

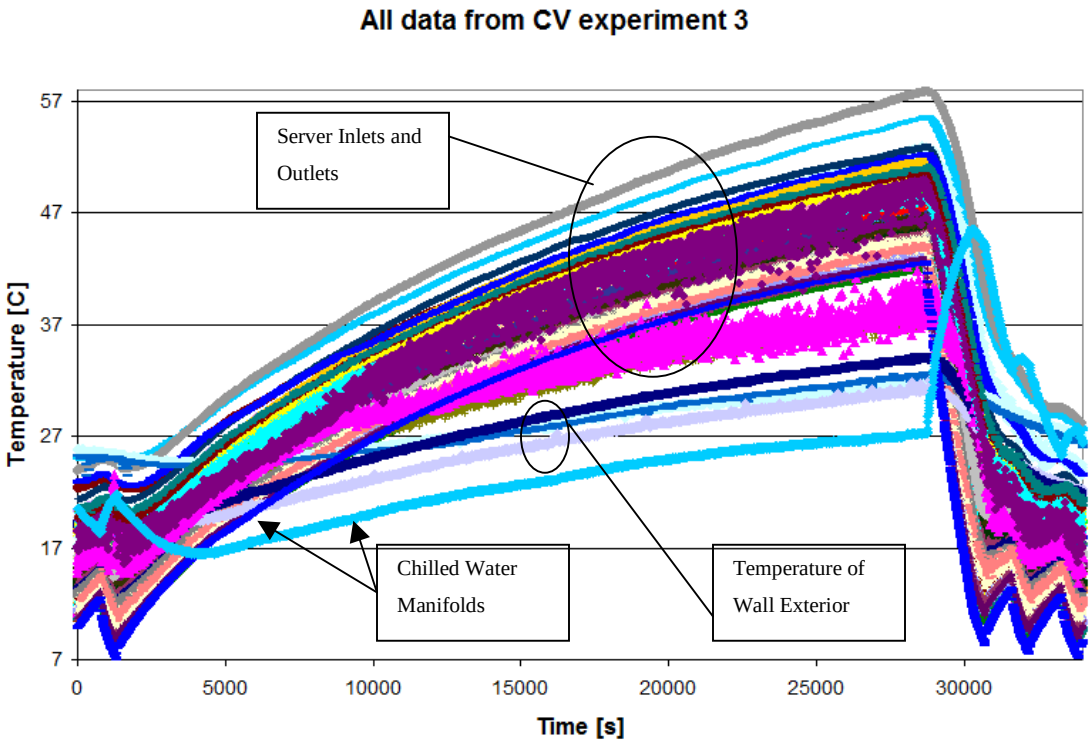

Figure 4.5-2: All data from CV experiment 3.

Figure 4.5-3 shows server inlet temperatures over the course of control volume experiment 3. The narrowing of the temperature trends observed during the period of interest in the first two control volume experiments is not readily observed here. This phenomenon is likely due to the much larger time scale of the experiment. Since changes in temperature take much longer, the server responses can be assumed quasi-steady.

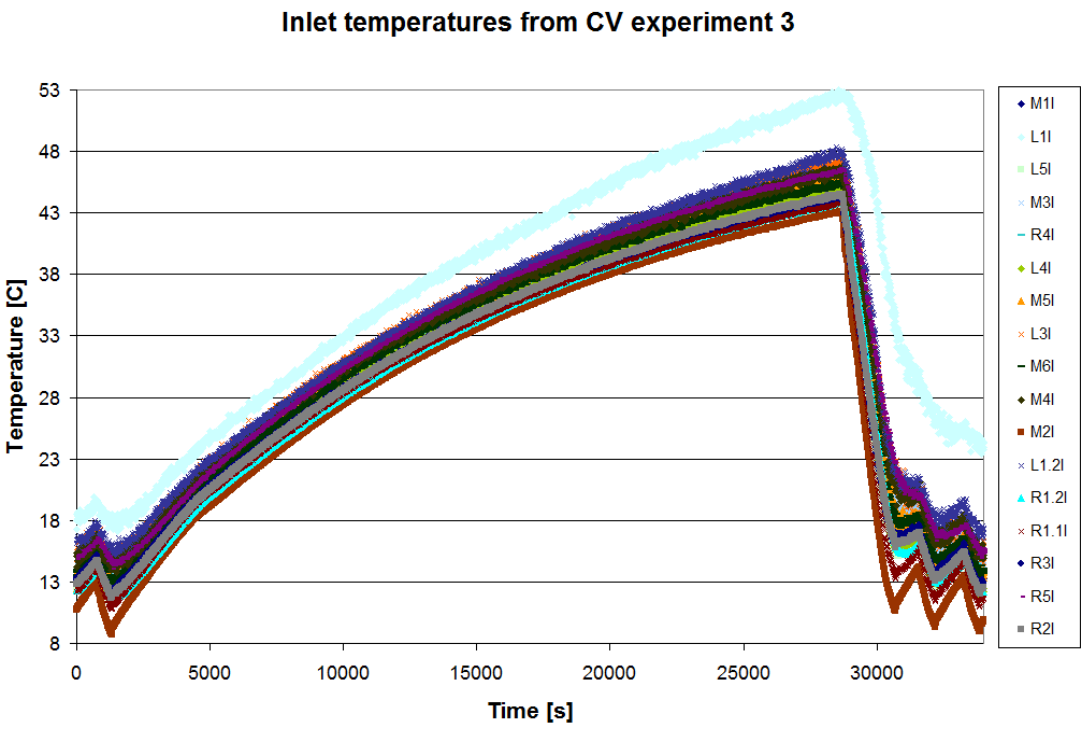

Figure 4.5-3: Inlet temperatures from CV experiment 3.

Figure 4.5-4 shows middle rack inlet temperatures from control volume experiment 3. The relative temperature distribution among the inlets remains relatively the same; there are a few minor exceptions as servers adjust fan speed to compensate for increasing CPU temperatures the temperature differences between the server inlets remains constant as the room warms.

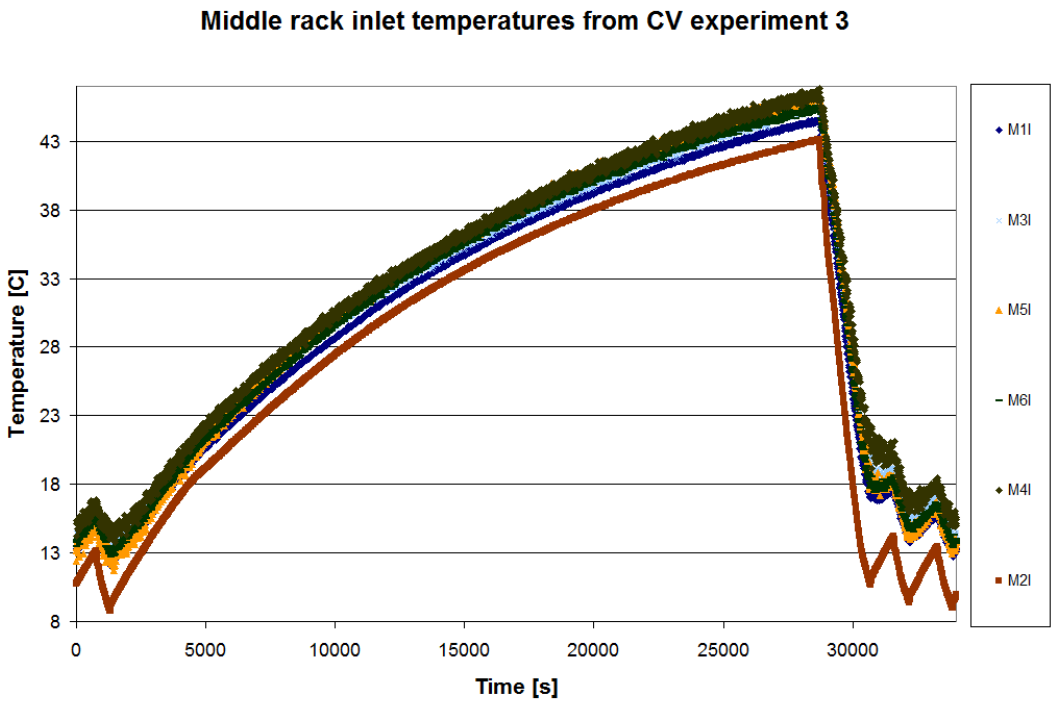

Figure 4.5-4: Middle rack inlet temperatures from CV experiment 3.

Figure 4.5-5 shows the temperature response at the server outlets during control volume experiment 3. Again it is apparent that the temperature differences between the server outlets remain relatively constant. However, Figure 4.5-6 shows that, for the middle rack, the top and bottom server experience fluctuations in outlet temperature; these fluctuations are more gradual than in control volume experiment 1, showing that the flow is well mixed to a nearly uniform which fluctuates at these outlets.

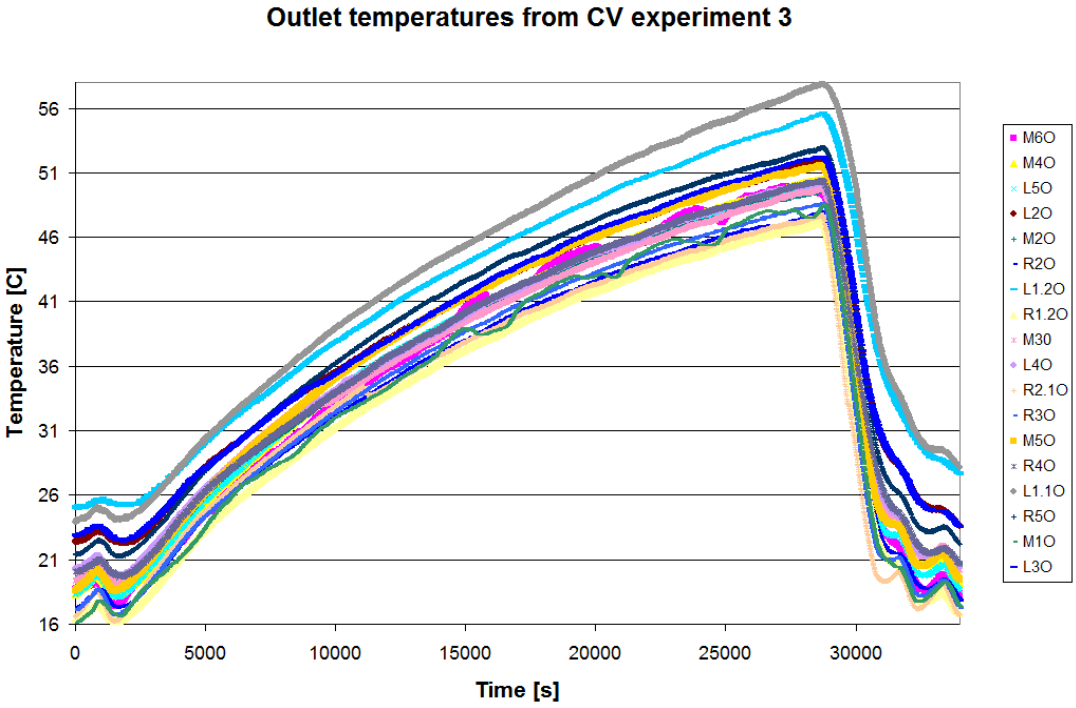

Figure 4.5-5: Outlet temperatures from CV experiment 3.

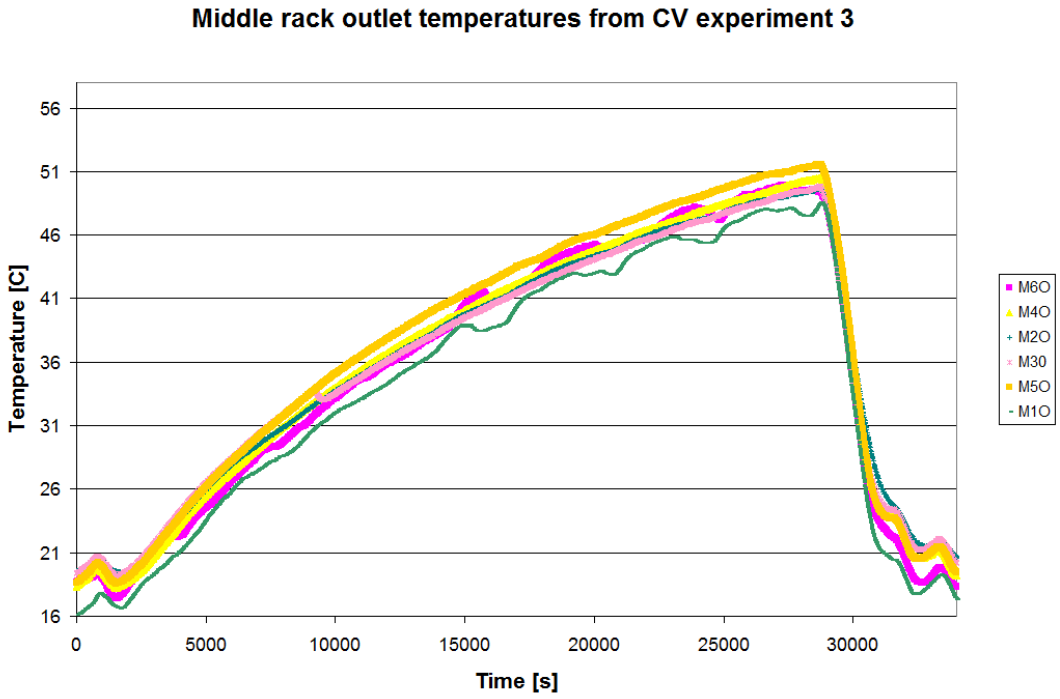

Figure 4.5-6: Middle rack outlet temperatures from CV experiment 3.

Figure 4.5-7 shows the average air temperature at server inlets and outlets during control volume experiment 3, as well as the average air temperature rise across the servers. The experiment was started a time when the servers were releasing heat and the rate of temperature rise in the control volume was gradual. Therefore, the initial drop in temperature difference across the servers observed in control volume experiment 1 is not present. However the effects of thermal capacitance within the servers is observed before the chiller is switched off and again after it is switched back on, since the temperature rise does not remain constant throughout the experiment.

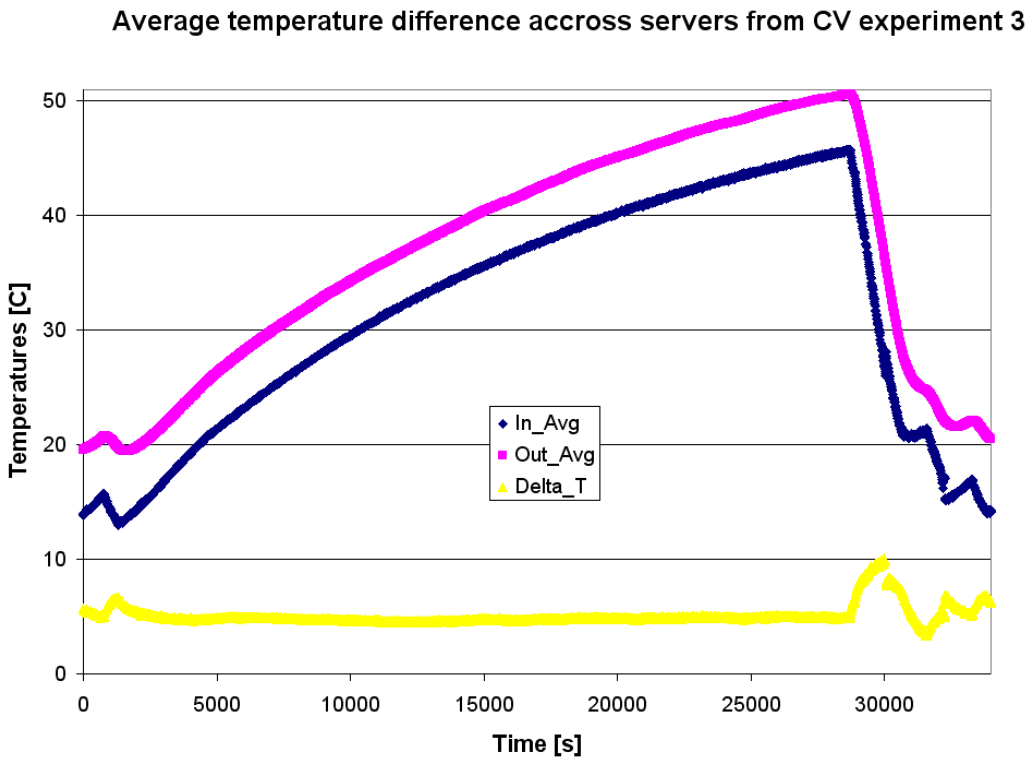

Figure 4.5-7: Average temperature rise across servers from CV experiment 3.

Figure 4.5-8 shows colormaps representing the temperature at server inlets and outlets at various times throughout control volume experiment 3. It can be observed that the relative temperature distribution among the inlets and outlets remains almost constant and all the temperatures increase at approximately the same rate.

Inlets **Outlets**

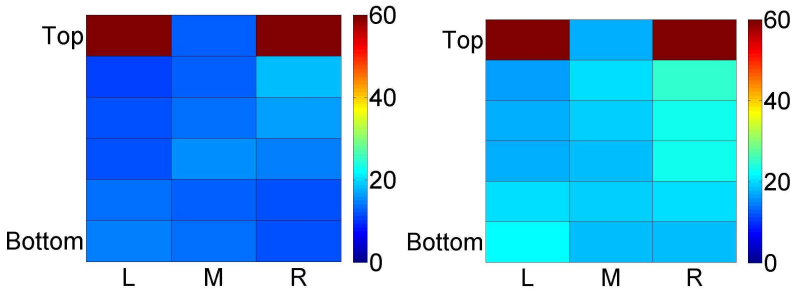

t_0 = 1320 s

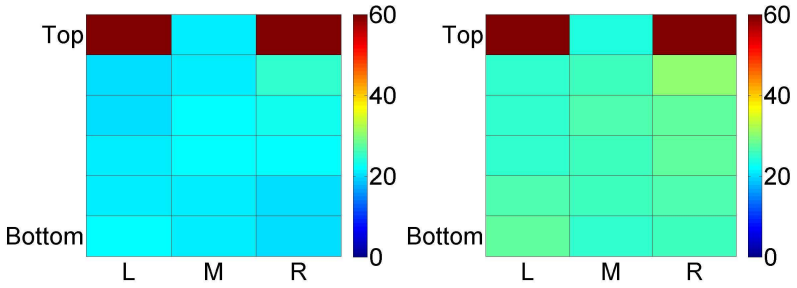

t = 5000 s

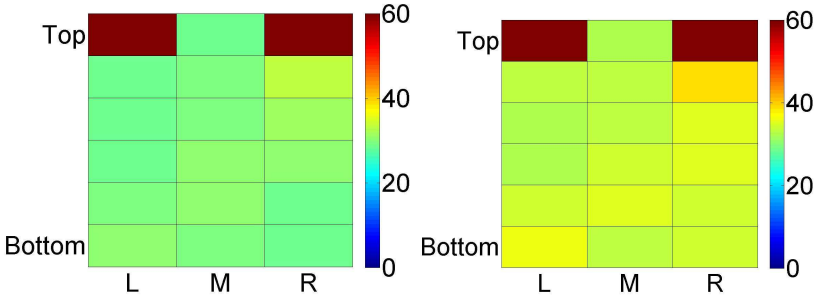

t = 10000 s

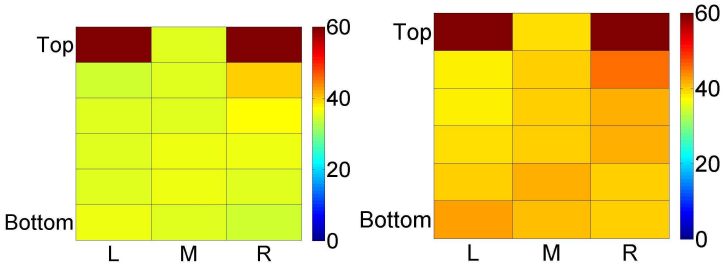

t = 15000 s

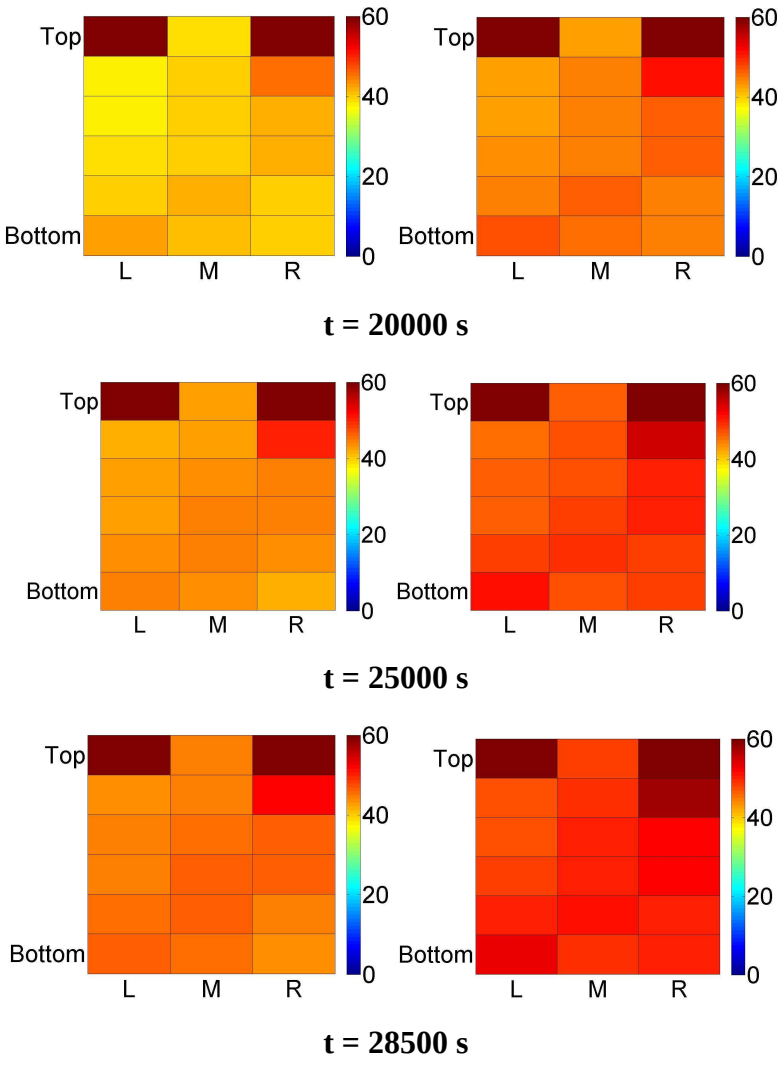

Figure 4.5-8: Colormap of inlet and outlet temperatures during CV experiment 3.

4.6 Comparison of CV Experimental Results with First Order Analytical Model

A simple analytical model has been developed based on a control volume energy balance for a closed system. Figure 4.6-1 shows the system considered for each variation on the analytical model. In every case, the walls of the control volume are assumed to be insulated and infiltration of air is neglected, making the control volume a closed system. Equation 4.6-1 shows the energy balance for a closed system with internal generation of energy, Equation 4.6-2 expresses the energy balance as a rate equation, and Equation 4.6-3 averages the thermal capacitance of the various materials within the

control volume. An underlying assumption in Equation 4.6-3 is that all materials within the control volume experience the same rate of temperature rise.

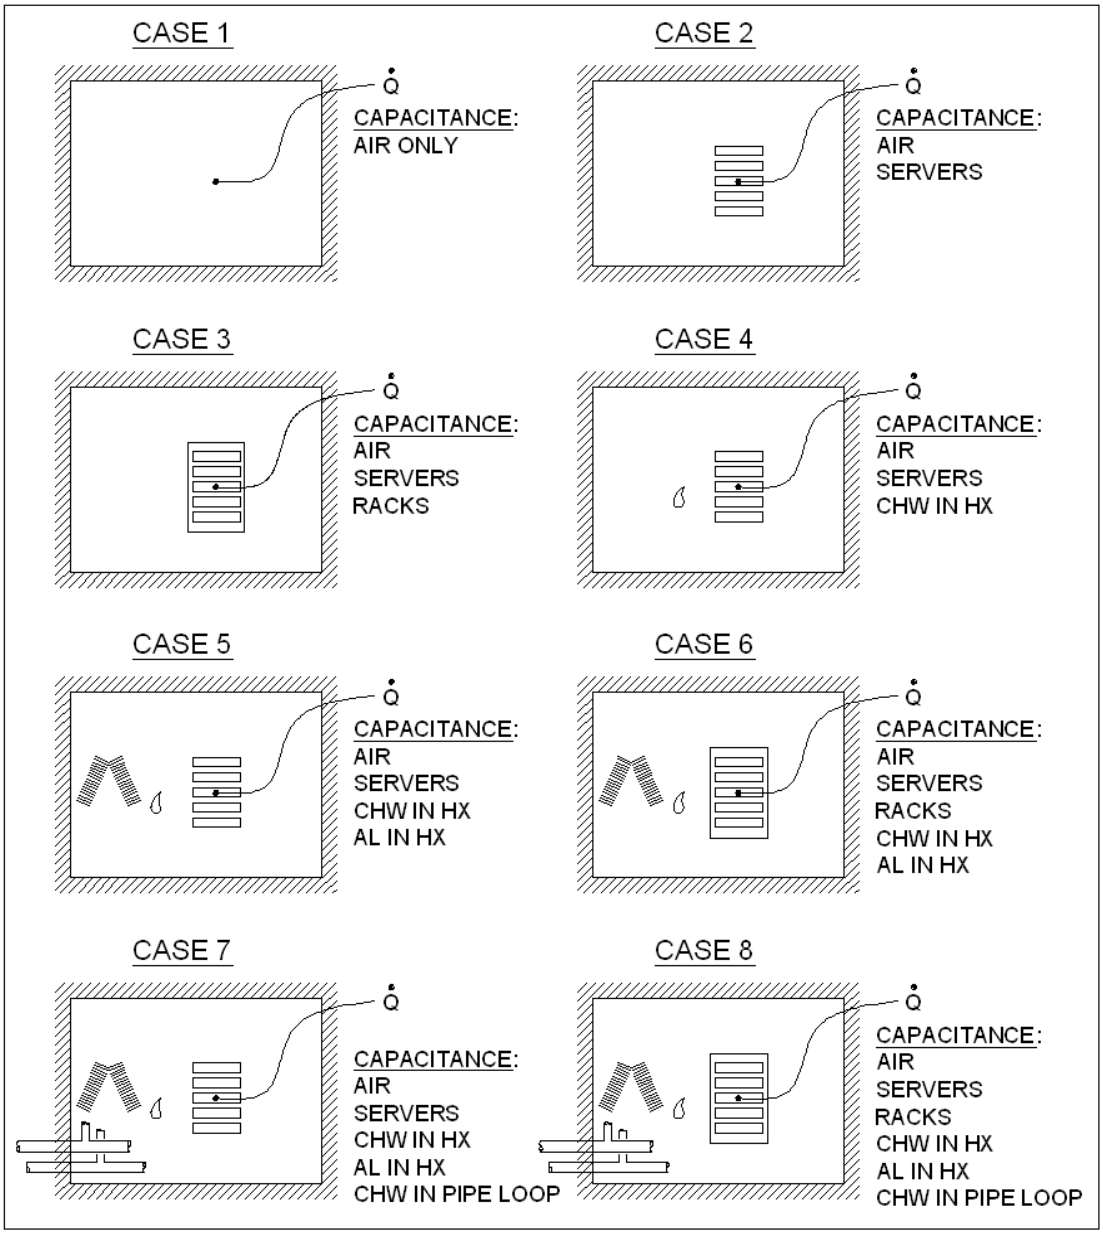

Figure 4.6-1: Thermal capacitance considered in various permutations of the analytical model.

$$E_{in} - E_{out} + E_{gen} = \Delta U = mc\Delta T = E_{gen} = Q \quad (4.6\text{-}1)$$

$$mc\Delta \dot{T} = \dot{Q} \quad (4.6\text{-}2)$$

$$\Delta \dot{T} \sum_i m_i c_i = \dot{Q} \quad (4.6\text{-}3)$$

Table 4.6-1 compares the three control experimental results with those obtained from variations of the first order analytical model. Experimental slope is based on the highest slope possible, at the beginning of each control volume experiment; during this period the temperature difference between the inside of the control volume and the outside room air should be small, and therefore, heat losses to conduction and infiltration should be low.

It is interesting to note that the model that only takes into account the thermal capacitance of the air predicts a rate of rise at least an order of magnitude higher than any experimental result. In fact, including the thermal mass of the servers still predicts a slightly higher rate than observed even in control volume experiment 1. However, assuming the entire mass of the racks to be available for energy storage underestimates the slope.

Including the thermal capacitance of the CHW within the HX of the AH has little effect on the analytical model, particularly as other thermal masses increase. However, including the thermal mass of the Al within the AH HX brings the result within an order of magnitude of the experimental results from control volume experiment 2. Adding the mass of the racks again underestimates the slope; nevertheless, this prediction is the most comparable to the experimental results. The larger timescale of the second control volume experiment may make the racks more available for storage of thermal energy.

Adding the thermal capacitance of the CHW in the piping loop brings the analytical results within an order of magnitude of the rate of rise found in control volume experiment 3, adding the thermal mass of the racks slightly lowers the prediction, but both analyses underestimate the experimental slope. An underestimate in the analytical model cannot be due to heat losses from the control volume, as losses would actually decrease the experimental slope. It is possible that the volume of CHW in the loop may be overestimated.

Table 4.6-1: Comparison of analytical and experimental results.

Result Source: Thermal Capacitance Considered (Slope Temp. vs. time in K/s)	Compare with CV Experiment Number			
	N/A	1	2	3
Analytical: Air Only	0.420			
Analytical: Air and Servers (Note-1)		0.0130		
Experimental: CV Experiment 1		0.0112		
Analytical: Air, Servers, and Racks (Note-2)		0.0063		
Analytical: Air, Servers, CHW in AH HX	0.0130			
Analytical: Air, Servers, CHW in HX, HX Aluminum (Note-3)			0.0082	
Experimental: CV Experiment 2			0.0058	
Analytical: Air, Servers, CHW in HX, Al in HX, Racks			0.0049	
Experimental: CV Experiment 3				0.0023
Analytical: Air, Servers, CHW in HX, Al in HX, CHW in Piping Loop				0.0019
Analytical: Air, Servers, CHW in HX, Al in HX, CHW in Piping Loop, Racks (Note-4)				0.0016

Notes:

1. **Server weight taken from [31] for maximum configuration and assumed to be constructed of 100% Al.**
2. **Rack weight taken from [27] for similar rack and assumed to be 100% steel.**
3. **HX assumed to comprise most of the weight of a similar AH [14].**
4. **Pipe lengths based on [33] to calculate volume of CHW.**

5 CRAC Heat Exchanger Response to Step Change in Chilled Water Flowrate

5.1 Experimental Setup

The CEETHERM data center laboratory employs the cooling scheme shown in Figure 1.1-1, Figure 5.1-1, and Figure 5.1-2. Various servers reject heat to the air. The CRAC unit used for this experiment is a Liebert downflow CHW cooled unit, piped with three-way valve configuration, and includes a Toshiba VF-S11Variable Frequency Drive (VFD). Appendix A.4 shows photographs of the CRAC. The chilled water loop is composed of a main loop of 76 mm (3 in) piping and 50 mm (2 in) runouts to CRAC units. A constant speed pump circulates chilled water. The lab chiller is a Trane RTAA-130 rated at 450 kW (130 tonnes). The chiller has two equally sized compressors, but normally only one needs to operate, since the data center is not yet fully populated. For more details on cooling equipment, refer to Table A 1.

Since the main component of the CRAC unit is an air-to-water heat exchanger (HX), the important measurements for characterizing its performance were inlet and outlet temperatures of each fluid and fluid flow rates. Therefore, thermocouples were placed at the following locations on the CRAC: at air inlet filters, after the HX, at the fan outlet, on the inlet chilled water (CHW) header, and at the outlet CHW header. In addition, an ultrasonic flow meter was placed on the return CHW pipe leaving the HX and a thermal anemometer (TA) was situated at the fan outlet. These locations are shown in Figure 1.1-7. Temperature values were recorded in Labview on a laptop computer via National Instruments Field Point data collection equipment as listed in Table 2. Probable ranges for uncertainties in sensing equipment and sources these uncertainties are also listed in Table A 2 and associated notes.

Figure 5.1-3 shows the sequence of events involved in the experiment. The experiment was performed as follows: The CHW pump was turned off. The room air temperature was allowed to reach 29°C. Then the pump was restarted. Temperature and water flow rate data were collected every 5 seconds and airflow data were collected manually as close as possible to every 5 seconds. The time period of greatest interest spans the change from approximately steady state before pump restart to approximately

steady state after pump restart. Data from this time period will be analyzed in the next section.

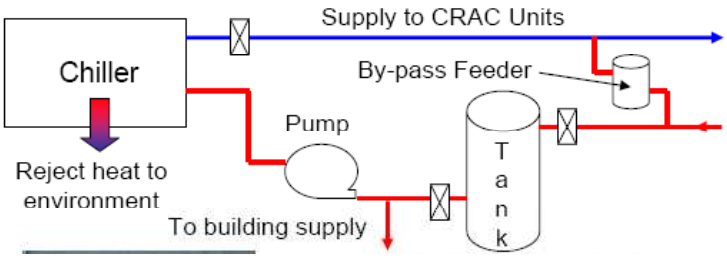

Figure 5.1-1: CHW piping to data center [6]

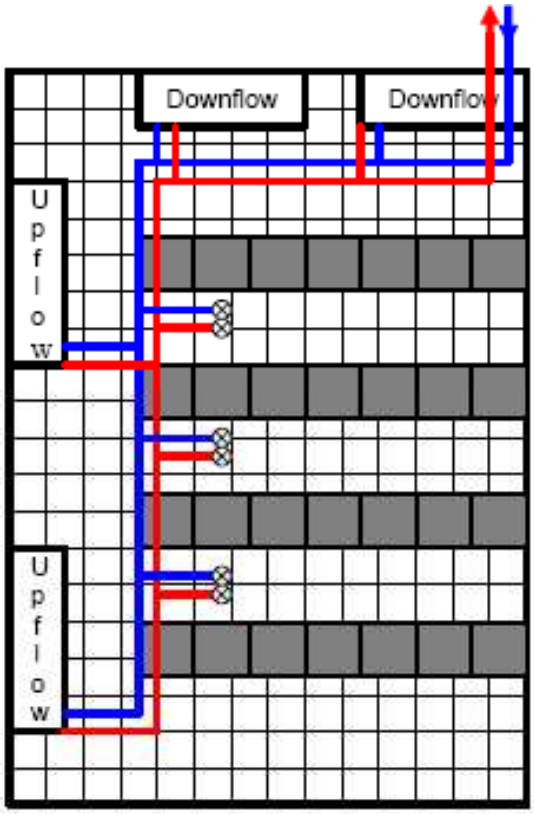

Figure 5.1-2: CHW piping within data center [6]

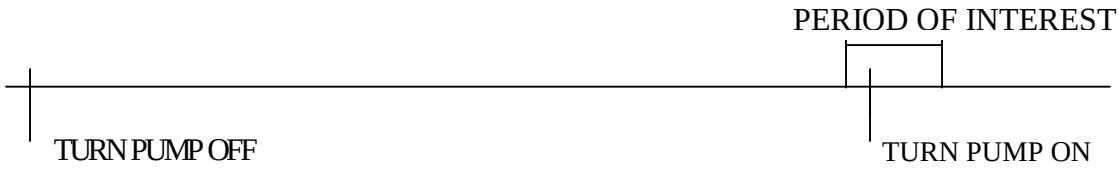

Figure 5.1-3: Timeline of events during CRAC response experiment.

5.2 AH HX Response Results

Figure 5.2-1 shows water flowrate data from a typical pump restart experiment. The graph shows that the flowrate jumps sharply from 0.2 L/s at time 5 s to 3.7 L/s at time 10 s. Because the change is much faster than the dynamic thermal response of the outlet air, the water flow approximates the shape of a step function.

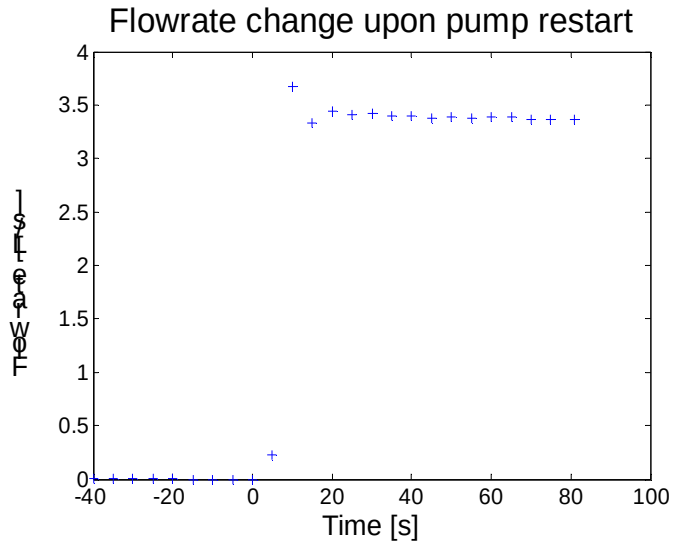

Figure 5.2-1: Flowrate change upon pump restart.

The equivalent heat transfer coefficient (UA) of the HX depends on the flowrate of both air and CHW. Thus, it was necessary to determine how the VFD regulated the airflow during the experiment. Measurements showed that the face velocity at the outlet of the CRAC fan remained between 24.0 m/s and 25.2 m/s during the period of interest.

Since the magnitude of velocity fluctuations was small, airflow will be assumed to remain constant during the time period of interest.

Figure 5.2-2 shows temperature changes during pump restart. During the time that the pump is off, the CHW inlet header is exposed to airflow and its temperature rises. However, CHW at about 8°C remains in the CHW loop. When the pump is restarted, this CHW enters the HX and the incoming CHW temperature remains between 8-10°C during the time period of interest. Due to the large volume of warm air present in the room before the experiment, the CRAC inlet air temperature also remains relatively constant over the period of interest. Considering these factors we can conclude that the experiment approximates a step change in CHW flowrate with all other variables held constant. As might be expected under the circumstances the temperature responses of the HX, fan, and CHW outlet header all follow the familiar 1st order curve. In particular we can assign a time constant of 10 seconds to each response, where the time constant is defined as in Equation 6.1.1 [11].

$$T = T_{\max} - (T_{\max} - T_{\min}) \bullet (1 - e^{-t/\tau}) \qquad (6.2.1)$$

However, there is a delay of approximately 20 seconds between the air outlet responses and the CHW outlet responses, presumably because it takes a finite amount of time for the CHW to progress through the HX. It can also be noted that although measurements of outlet air taken just after the HX have a lower temperature than those taken at the fan outlet, both data sets lead to the same time constant for the CRAC.

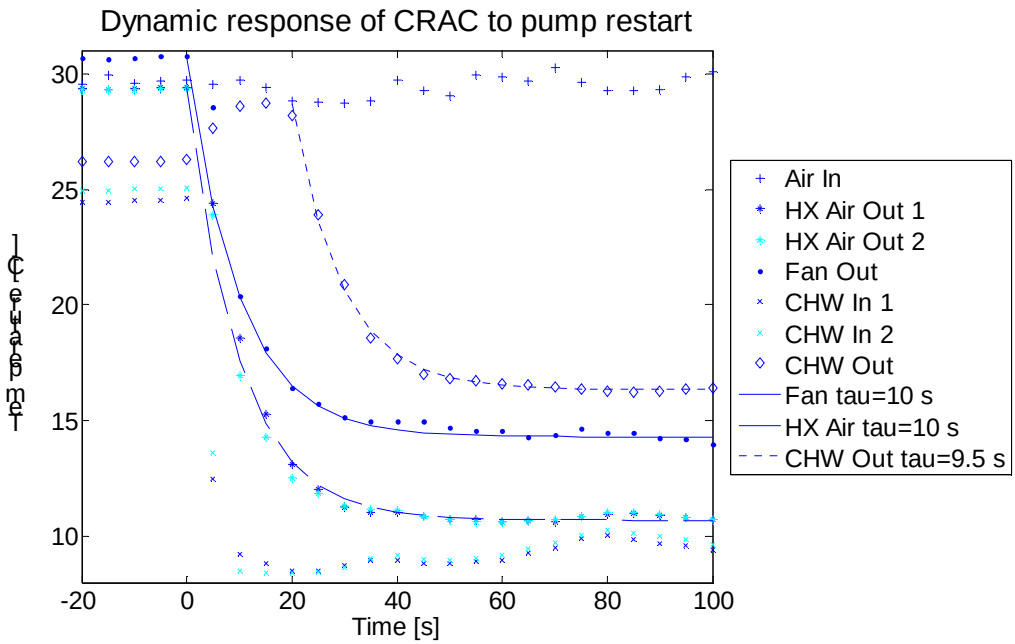

Figure 5.2-2: Dynamic response of CRAC to step change in water flowrate.

It is interesting, but not unexpected, that a dynamic thermal system should follow a first order response curve when only one boundary condition is changed, which can be approximated as a step function.

5.3 Implications

The CRAC HX used for this experiment has a time constant of 10 seconds using a first order analysis. This time constant was determined by visually fitting a first order curve to the data set. The first order curve fits the data well; this is expected, since the only boundary condition which changed significantly was the CHW flowrate. Modeling of events with a short time scale (close to the time constant for the CRAC) would likely be improved by using a first order response to describe the response of the CRAC unit. However, models that describe events taking place over hours or days might see little improvement over a steady state model.

6 Conclusion

Data from the CRAC transient response experiment suggest that the steady state assumption across the AH HX may be acceptable for events with time scales much larger than the ~10 second time constant. However, the data from the server time constant experiment suggest that the thermal capacitance of the CE should be taken into account for modeling events taking place over intervals comparable with the 130-380 second time constant; the exact time constant depends on the particular CE chosen. When an event is characterized by continually increasing temperatures, as is the case during data center power failure events, the time rate of increase in the temperatures may have more influence on whether a transient model of various equipment should be used than the total duration of the event. The first two time constant experiments never reached a steady state temperature rise across the server, despite the fact that measurements took place over a time scale an order of magnitude larger than the time constant. However, the third control volume experiment showed quasi-steady state behavior, due to the slow changes in temperature within the system. Previous first order models have predicted the rise in the average air temperature based on only the thermal capacitance of the air; the first order model proposed here shows that much better estimates can be obtained by including the thermal mass of servers, AH HX if fans run, and CHW in the piping loop if the CHW pump runs. Various power failure scenarios are likely, based on what cooling equipment receives emergency generator or uninterrupted power supply. It is not possible to make a blanket recommendation as to what level of backup power provisioning holds the most value for a particular facility. Various factors influence the decision, including: level of reliability desired by owner, power density of equipment within facility, past data showing probable outage duration, and cost of data center downtime. Each of these factors calls for additional lists of factors and further investigation; this thesis has experimentally explored the effects of providing UPS to various cooling equipment for a given mock facility.

7 Recommendations for Future Work

This thesis has explored the temperature response of a relatively small control volume with relatively few servers to various transient scenarios representing various levels of provisioning of cooling systems for power failure in a data center. The temperature distributions within the control volume are largely a result of the airflow patterns; these have not been explored experimentally. Through the use of particle image Velocimetry (PIV) it would be possible to experimentally investigate the airflow patterns developed when the AH is running, or when the AH is not running and the server fans drive the flow. In addition, it is suggested that a CFD/HT CM of a server that accounts for thermal capacitance be developed. Compact models could also be developed to account for the transient behavior of the AH, and even the temperature rise of the CHW in the piping loop. Future work could be devoted to characterization of time constants for various compute equipment subjected to a step change in inlet temperature, or to the effect of increasing the power density within a room level control volume under typical power failure scenarios. Current research has only touched the tip of the iceberg that symbolizes the unknowns involved in predicting data center response to various power failure scenarios, but further research could be of great help while designing the data center to respond appropriately with cost effective level of power infrastructure. This thesis hopes to serve as an introduction to and call for further work in a field that has seen little research, should others seek to extend the body of knowledge.

Appendix A.1: Cooling Infrastructure outside Data Center Space

Table A 1: Cooling infrasutructure equipment

Device	Details
CRAC	Liebert, Downflow, Chilled Water, 3-Way Valve: FH529VCAG00
VFD	Toshiba VF-S11
Chiller	Trane 130 Tonne: RTAA-130
Pump	Armstrong 5 HP: AVN184TTDR7356DV E

Figure A 1: RTAA-130 chiller at CEETHERM laboratory.

Figure A 2: Pump, tanks, makeup water, and associated piping in CEETHERM lab.

Appendix A.2: Measurement Equipment

Table A 2: Data collection equipment

Device	Model	Estimated Uncertainty Bounds
Thermocouples	OMEGA Type-T 40 AWG	Note 1
Thermocouple Reader	NI FP-TC-120 and FP-1601	Note 1
Laptop	Dell Inspiron 600m	Note 1
Programs	Labview 7.1 with MAX	Note 1
Ultrasonic Flowmeter	GE Panametrics PT878	Note 3
Thermal Anemometer	TSI Velocicalc 8350	Note 4
Stopwatch	Casio 756 AQ-47	Note 5

Notes:

1. **Thermocouple data collection equipment calibrated to combined uncertainty range of ±0.3°C [15]**
2. **Not used**
3. **±2% to 5% of reading [15]**
4. **±5 of reading or ±0.025 m/s (5 FPM), whichever is greater [36]**
5. **Human error estimated less than 1 second**

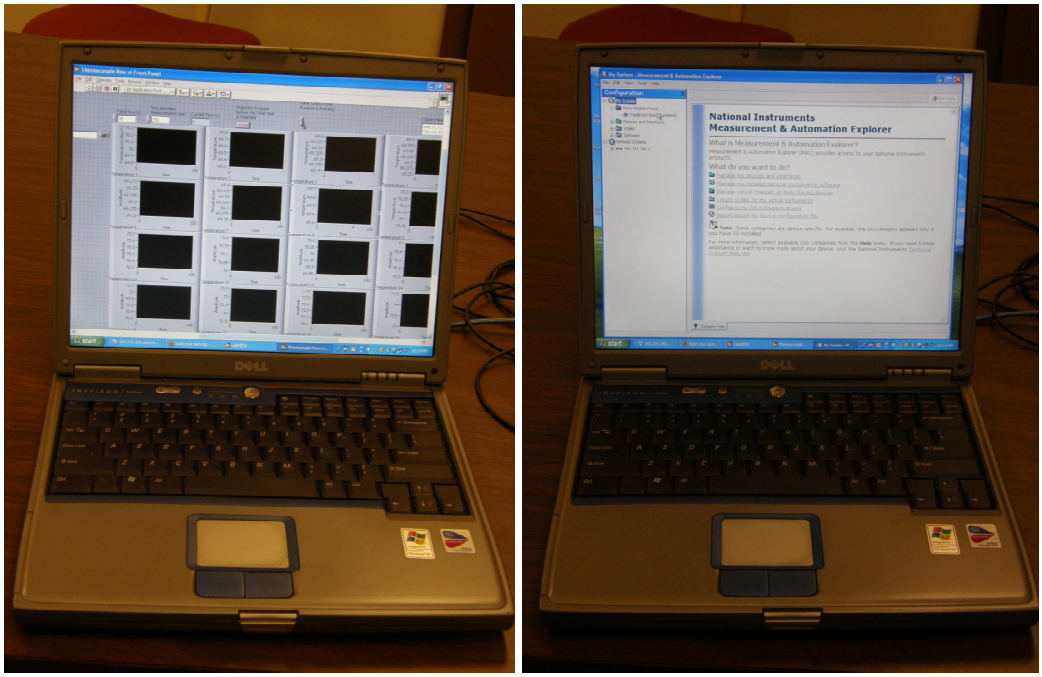

Figure A 3: Laptop running Labview and MAX programs to record data.

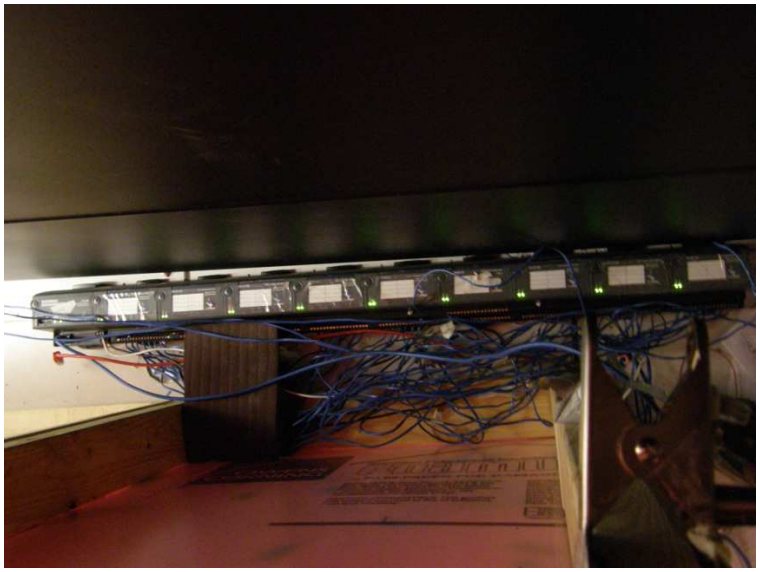

Figure A 4: NI Field Point data collection equipment.

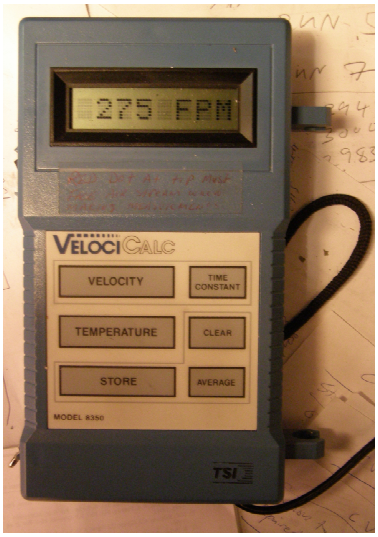

Figure A 5: TSI Velocicalc handheld thermal anemometer.

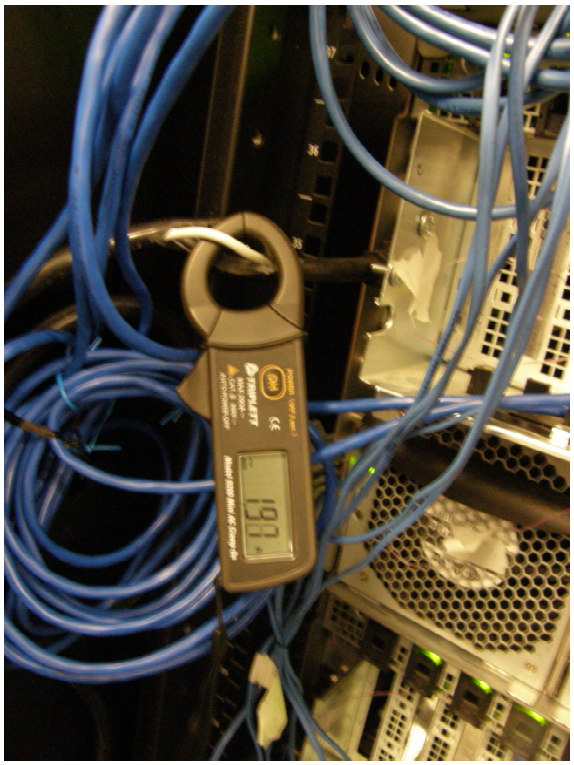

Figure A 6: Measurement of current with clamp-on ammeter.

Appendix A.3: Pictures of Server Simulator Heater Unit

Figure A 7: Inlet to server simulator heater unit.

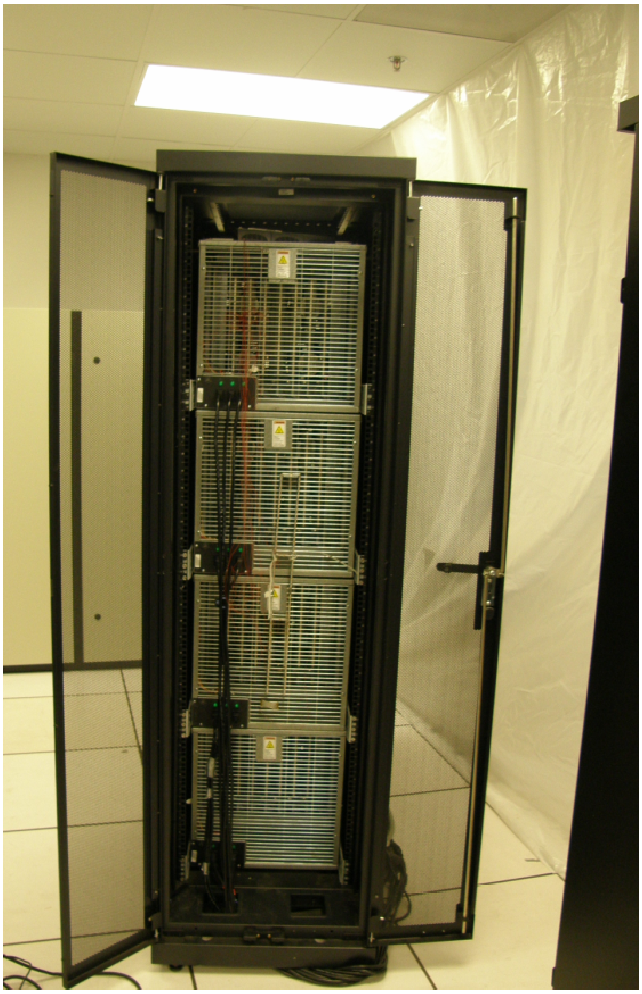

Figure A 8: Outlets of server simulator heater unit.

Figure A 9: Sample controls to server simulator heater unit.

Appendix A.4: Pictures of Legacy Servers

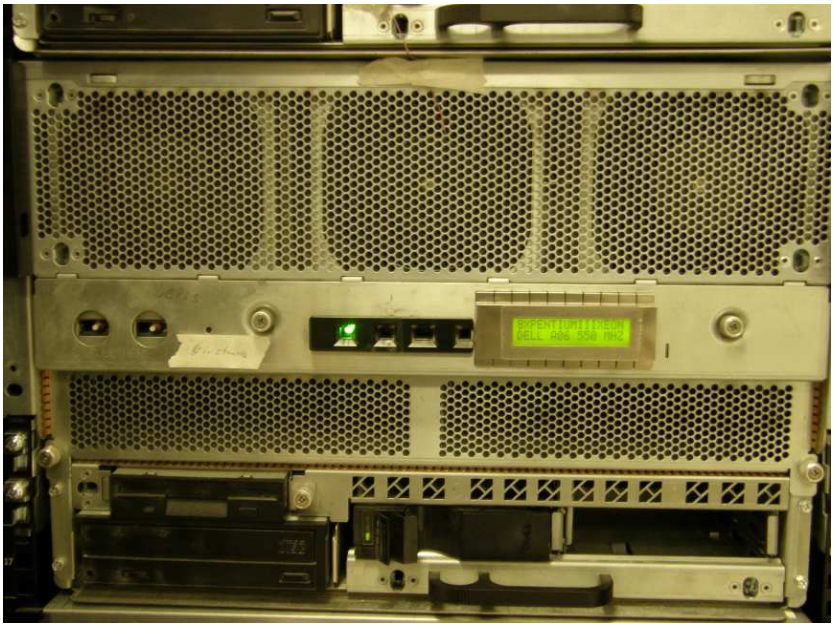

Figure A 10: Inlet to sample legacy server.

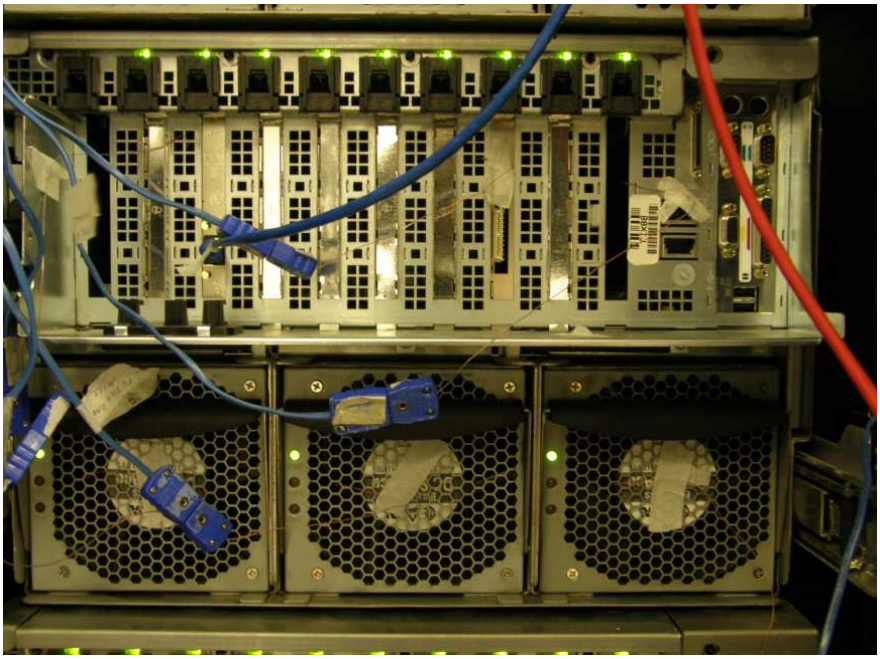

Figure A 11: Outlet from sample legacy server.

Figure A 12: Power supply cooling fan in legacy server.

Figure A 13: Processor cooling fan in legacy server.

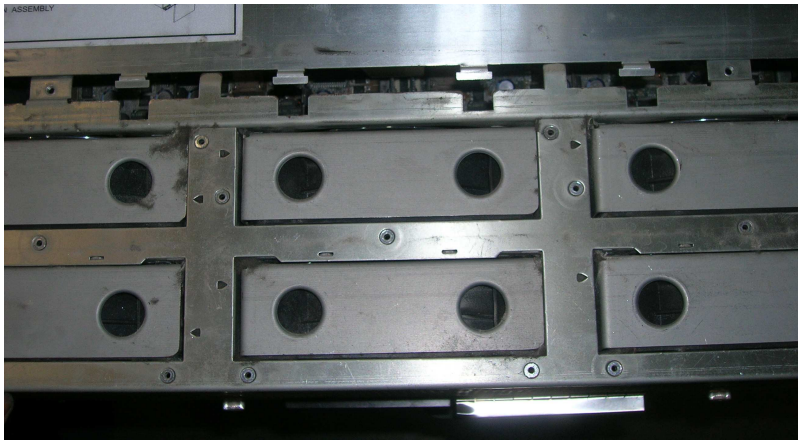

Figure A 14: Processor fans installed in series behind server inlet.

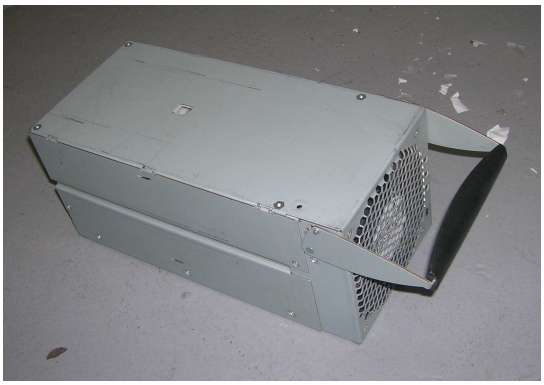

Figure A 15: Sample power supply for legacy server.

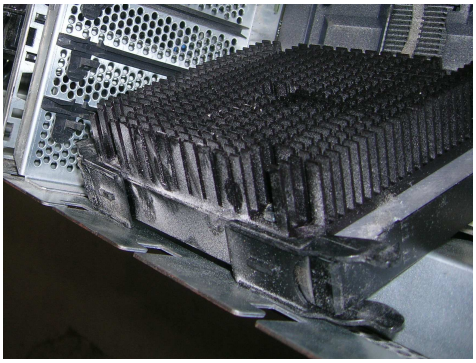

Figure A 16: Sample Processor package with heat sink for legacy server.

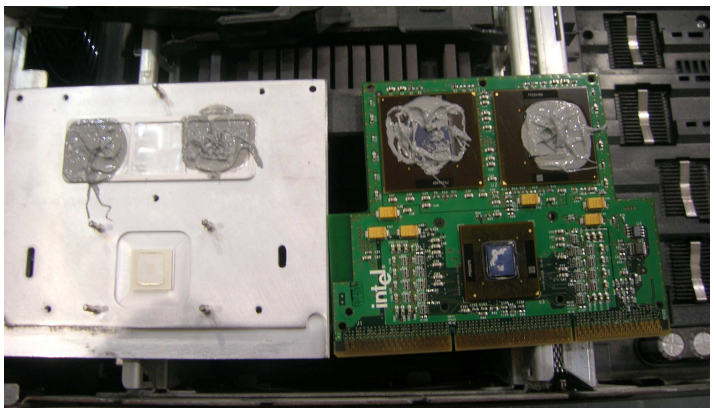

Figure A 17: Processor package with heat sink removed to show mateials.

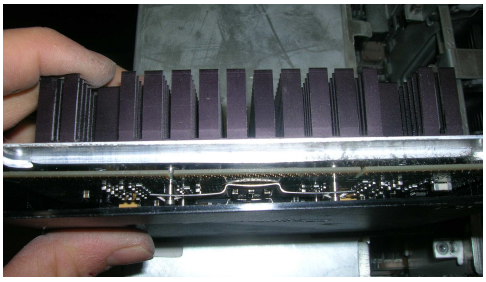

Figure A 18: Side view of processor package.

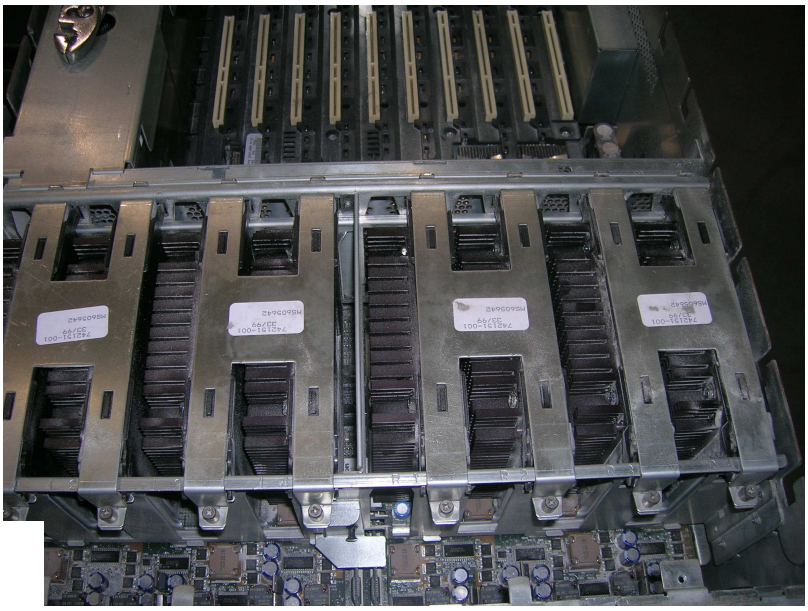

Figure A 19: Processors installed with circuit boards aligned vertically behind cooling fans.

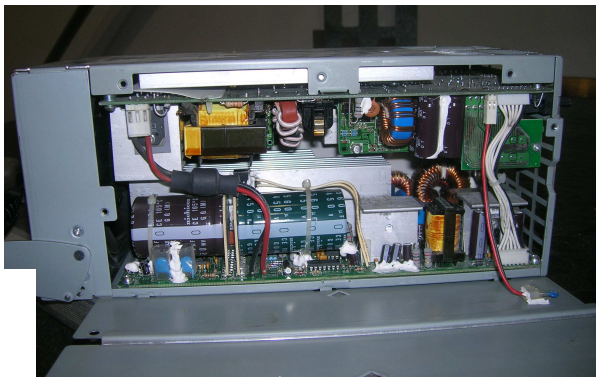

Figure A 20: Heat sinks, circuit boards, and electronic components within power supply.

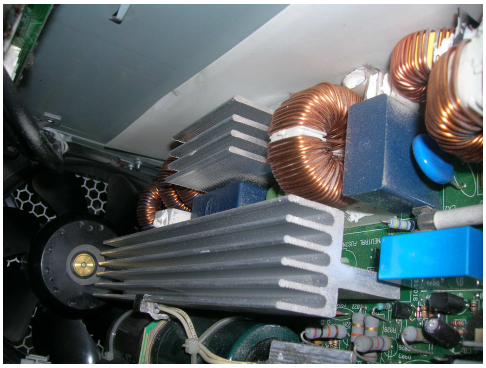

Figure A 21: Heat sinks with thicker, more widely spaced fins within power supply.

Appendix A.4: Pictures of Air Handler (CRAC)

Figure A 22: Inlet of CRAC unit with filters.

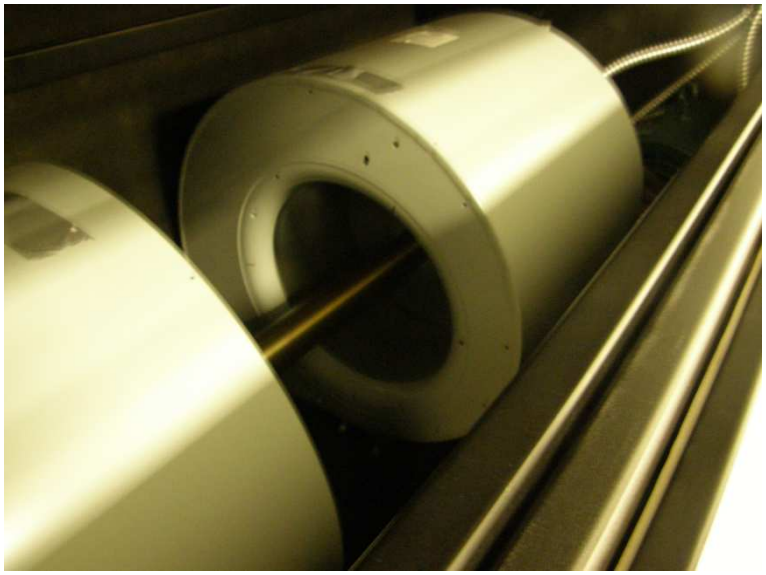

Figure A 23: Centrifugal fans within CRAC.

Figure A 24: Outlet of centrifugal fan beneath CRAC unit.

Figure A 25: CRAC with front panels removed to show heat exchanger.

Figure A 26: Side view of heat exchanger showing six tube passes.

Figure A 27: CRAC control window showing current conditions.

Figure A 28: Three-way valve within CRAC unit.

Appendix A.5: Server time constant room division and thermocouple placement.

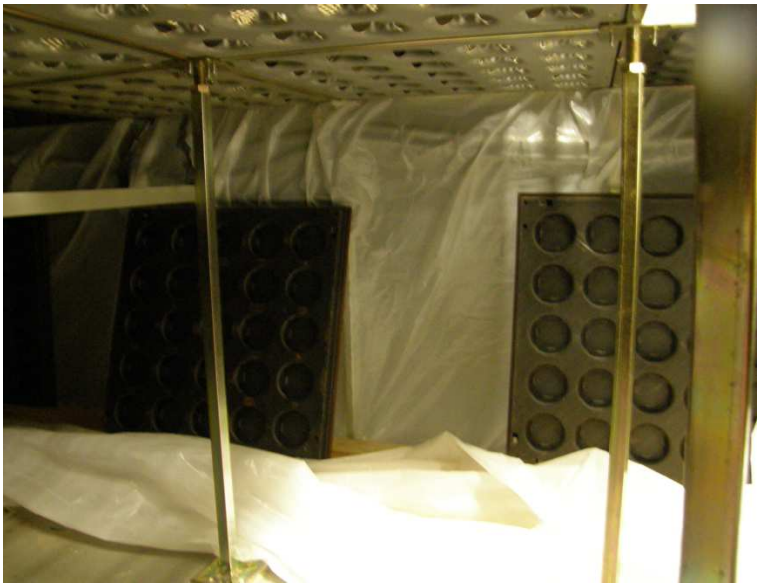

Figure A 29: Raised floor plenum divided by polyethylene sheeting to segregate air into hot and cold side.

Figure A 30: Opening in polyethylene sheet to allow for passage of human operator.

Figure A 31: Rack of servers adjacent to opening used to transfer said rack during transition between hot and cold sides.

Figure A 32: Chilled water piping and cable trays were among the obstacles encountered while separating the hot and cold sides within the raised floor plenum.

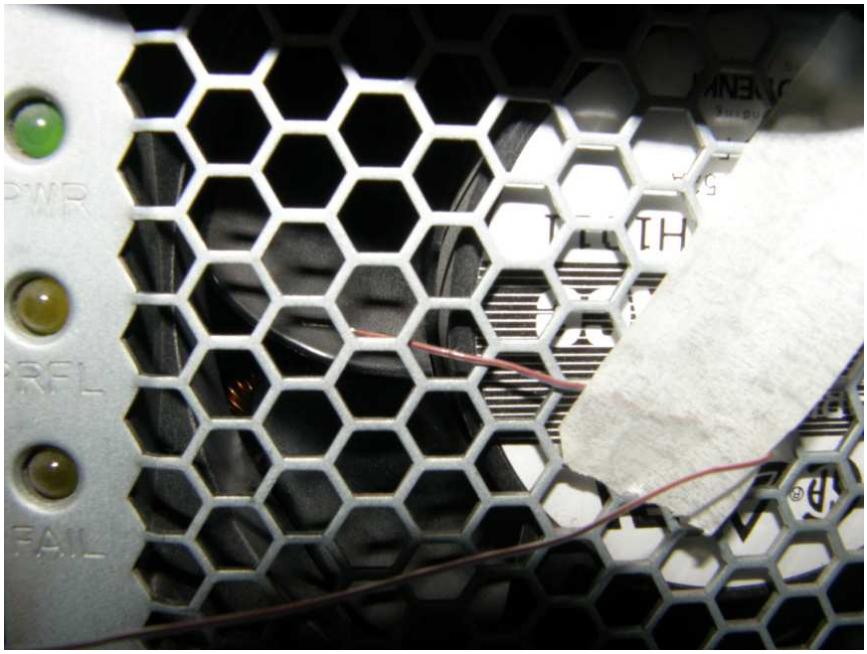

Figure A 33: Thermocouple placed at power supply outlet of legacy server.

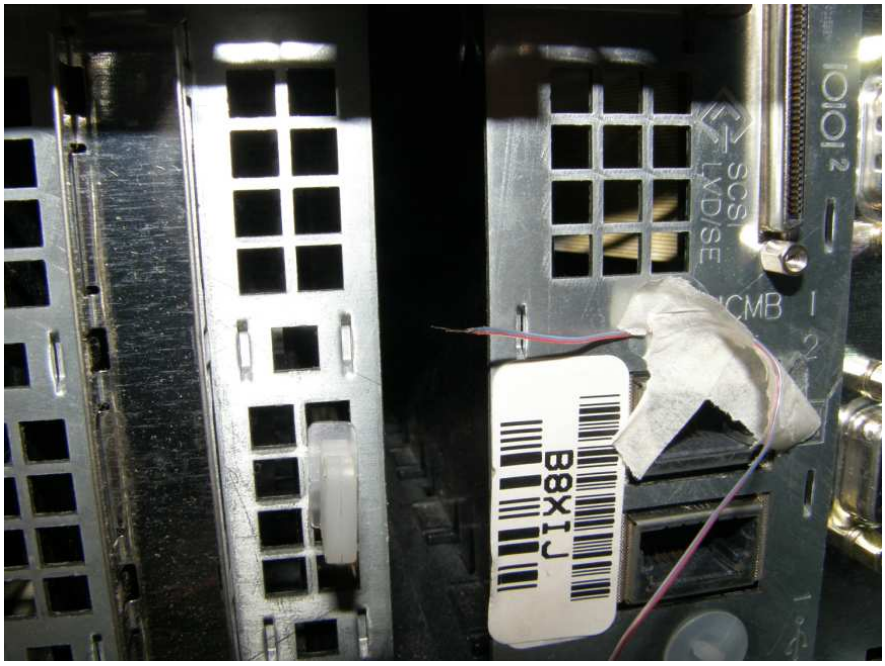

Figure A 34: Thermocouple placed at processor outlet of legacy server.

Appendix A.5: Pictures of Control Volume Used in Experiments

Figure A 35: Outside of control volume enclosure viewed from above floor.

Figure A 36: Interior of control volume showing server outlets and return air path to air handler as viewed through opening in enclosure.

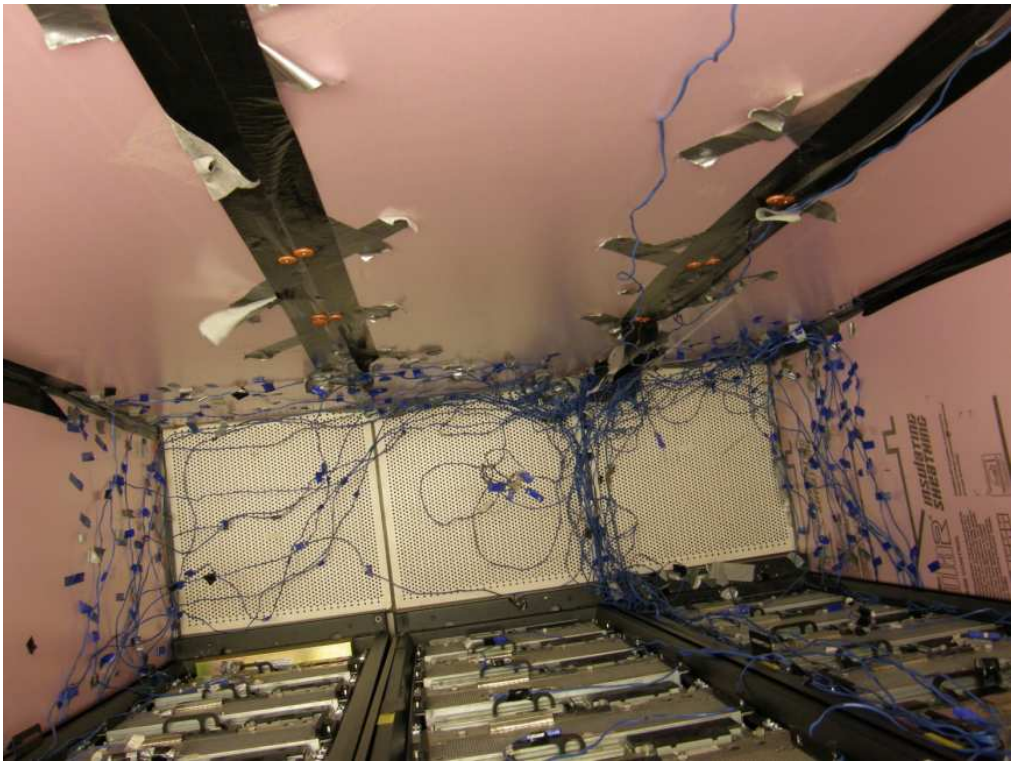

Figure A 37: Overhead view of cold aisle within control volume.

Figure A 38: Outside of control volume enclosure viewed from within raised floor plenum.

Figure A 39: Interior of control volume viewed from within raised floor plenum.

Appendix A.6: Order of Magnitude Estimate of Conduction and Infiltration Heat Losses

```
T_outer=20;
T_inner=50
R_val_US=3;
convert=(5/9*(0.0254*12)^2/0.29307);
R_val=R_val_US*convert;
delta_T=T_outer-T_inner;

%dimensions
cvLength_ft=8*2;
cvWidth_ft=6*2;
cvHeight_ft=9;
conversion=0.0254*12;
cvLength_m=8*2*conversion;
cvWidth_m=6*2*conversion;
cvHeight_m=9*conversion;
area=cvWidth_m*cvHeight_m+cvLength_m*cvHeight_m;
U=1/R_val;
q=U*area*(T_outer-T_inner);
plenumHeight_ft=3;
plenumHeight_m=plenumHeight_ft*conversion;
areaPlenum=2*(cvLength_m+cvWidth_m)*plenumHeight_m;
area_w_plenum=area+areaPlenum;
q_w_plenum=area_w_plenum*U*(T_outer-T_inner)

cfm_p_ft2=0.5;
area_ft=2*(cvLength_ft+cvWidth_ft)*12;
cfm=cfm_p_ft2*area_ft
m3ps=cfm/60*(12*0.0254)^3; %s/min in/ft m/in vol
```

```
rho=1.2929; %air kg/m3 at sea level
cp=1000; %J/kgK air
q_inf=rho*m3ps*cp*delta_T
```

Appendix A.7: First Order Prediction of Rate of Temperature Rise Considering Various Thermal Capacitances

```
%volume of air from gambit model

abv_fl=10.533897; %m3
bel_fl=6.626142; %m3
racks=2.5;
crac=3; %estimate
 cv=abv_fl+bel_fl+racks

%eng_toolbox kg/m3 kJ/kg K air props Al from Incrop &
Dewitt
%http://www.engineeringtoolbox.com/air-properties-
d_156.html

rho=1.205;
c_p=1.005;
rho_al=2702;
c_p_al=0.903;
m=rho*cv;  %kg

percent=1
servers=16;
m_al=51.4*percent*servers

m_racks=3*275.20*2.2;
%http://www.apc.com/resource/include/techspec_index.cfm?bas
e_sku=AR3100
c_p_steel=0.450;
```

```
m_coil=500; %kg
%http://www.liebert.com/product_pages/ProductTechnical.aspx
?id=13&hz=60
q_dot=10; %kW approx

pass=6;
tubes=28;
l=6*12*0.0254;
area=(0.5*0.0254)^2*pi;
v_water=area*l*pass*tubes;
rho_water=998.3;
m_water=v_water;
c_p_water=4.183;

%air only
airDelT=q_dot/(m*c_p)
air_servDelT=q_dot/(m*c_p+m_al*c_p_al)
CVexp1DelT=(23.87-17.39)/(1090-510)
air_serv_racksDelT=q_dot/(m*c_p+m_al*c_p_al+m_racks*c_p_ste
el)
air_serv_h2oDelT=q_dot/(m*c_p+m_al*c_p_al+m_water*c_p_water
)

air_serv_h2o_coilDelT=q_dot/(m*c_p+m_al*c_p_al+m_water*c_p_
water+m_coil*c_p_al)
CVexp2DelT=(22.38-19.12)/(2750-2190)
air_serv_h2o_coil_racksDelT=q_dot/(m*c_p+m_al*c_p_al+m_wate
r*c_p_water+m_coil*c_p_al+m_racks*c_p_steel)


tank=200*0.00378541178; %[m^3]
d_2=2*0.0254; %[m]
```

```
d_3=3*0.0254; %[m]
d_5=5*0.0254; %[m]
A_2=pi*d_2^2/4;
A_3=pi*d_3^2/4;
A_5=pi*d_5^2/4;
L_2=12*12*0.0254; %[m]
L_3=120*12*0.0254; %[m]
L_5=16*12*0.0254; %[m]
pipe_2=A_2*L_2;
pipe_3=A_3*L_3;
pipe_5=A_5*L_5;
Vol_sys=pipe_5+pipe_3+pipe_2+tank;
m_chw=Vol_sys*rho_water;

air_serv_h2o_coil_chwDelT=q_dot/(m*c_p+m_al*c_p_al+m_water*
c_p_water+m_coil*c_p_al+m_chw*c_p_water)
CVexp3DelT=(25.64-21.1)/(4710-2750)
air_serv_h2o_coil_racks_chwDelT=q_dot/(m*c_p+m_al*c_p_al+m_
water*c_p_water+m_coil*c_p_al+m_racks*c_p_steel+m_chw*c_p_w
ater)

%%RESULTS:
% airDelT =    0.4200
%
% air_servDelT =    0.0130
% CVexp1DelT =    0.0112
% air_serv_racksDelT =    0.0063
%
% air_serv_h2oDelT =    0.0130
%
% air_serv_h2o_coilDelT =    0.0082
% CVexp2DelT =    0.0058
```

```
% air_serv_h2o_coil_racksDelT =    0.0049
%
% air_serv_h2o_coil_chwDelT =    0.0019
% CVexp3DelT =    0.0023
% air_serv_h2o_coil_racks_chwDelT =    0.0016
```

Appendix A.8: Method of Loading the Processors Using Prime95

[root@Jedi1 1]# ./mprime2
Welcome to GIMPS, the hunt for huge prime numbers. You will be asked a few simple questions and then the program will contact the primenet server to get some work for your computer. Good luck!

Attention OVERCLOCKERS!! Mprime has gained a reputation as a useful stress testing tool for people that enjoy pushing their hardware to the limit. You are more than welcome to use this software for that purpose. Please select the stress testing choice below to avoid interfering with the PrimeNet server. Use the Options/CPU menu choice to make sure your cpu type was detected properly, then use the Options/Torture Test menu choice for your stress tests. Also, read the stress.txt file.

If you want to both join GIMPS and run stress tests, then Join GIMPS and answer the questions. After the server gets some work for you, stop mprime, then run mprime -m and choose Options/Torture Test.

Join Gimps? (Y=Yes, N=Just stress testing (Y): n
Main Menu

1. Test/Primenet
2. Test/User Information
3. Test/Vacation or Holiday
4. Test/Status
5. Test/Continue
6. Test/Exit
7. Advanced/Test
8. Advanced/Time

9. Advanced/P-1
10. Advanced/ECM
11. Advanced/Priority
12. Advanced/Manual Communication
13. Advanced/Unreserve Exponent
14. Advanced/Quit Gimps
15. Options/CPU
16. Options/Preferences
17. Options/Torture Test
18. Options/Benchmark
19. Help/About
20. Help/About PrimeNet Server

Your choice: 17

Beginning a continuous self-test to check your computer.

Please read stress.txt. Press Ctrl-C to end this test.

Test 1, 400 Lucas-Lehmer iterations of M19922945 using 1024K FFT length.

Test 2, 400 Lucas-Lehmer iterations of M19922943 using 1024K FFT length.

Appendix A.9: Method of Collecting Data from Motherboard Sensors

```
#bin/bash
# test file for scripting the output of OpenManage
Reporting program
# by Eric Burgett
rm testf1
rm testt1
clock
while [ 1 ]
do
omreport chassis fans -fmt ssv -outa testf1
clock>>testf1
omreport chassis temps -fmt ssv -outa testt1
clock>>testt1
echo "done with round 1">>testf1
echo "done with round X"
sleep 10
done
clock
omreport chassis fans -fmt ssv -outc testf2
omreport chassis temps -fmt ssv -outc testt2
echo "done with round 2"
```

Appendix A.10: Suggested Guidelines for Modeling of Servers with Thermal Capacitance.

This thesis has endeavored to show that compute servers have significant thermal capacitance and should not be modeled using the steady state assumption for prediction of temperatures during power failure events. This appendix is written to suggest a method that a researcher might pursue in constructing a model of a server that includes thermal capacitance.

A steady state model of a server that is often used in CFD/HT simulations is a plane perpendicular to the air stream that uses the porous jump boundary condition. The boundary condition is then set to yield either a constant heat input or a constant temperature rise. This model has the advantage of being simple enough to converge quickly with a relatively coarse mesh. More complexity is required for the suggested model.

It is proposed that solid material be introduced into the model of the server in order to provide a thermal storage medium. This might be done in a number of ways. However, as much simplicity as possible is desired for the sake of meshing and convergence. Therefore, the proposed method is to construct a rectangular model of the same outer dimensions as the physical server and fill some portion of the interior with porous media. Figure A 29 illustrates and example geometry for such a model. The properties of the porous media must be tuned to obtain the correct thermal capacitance.

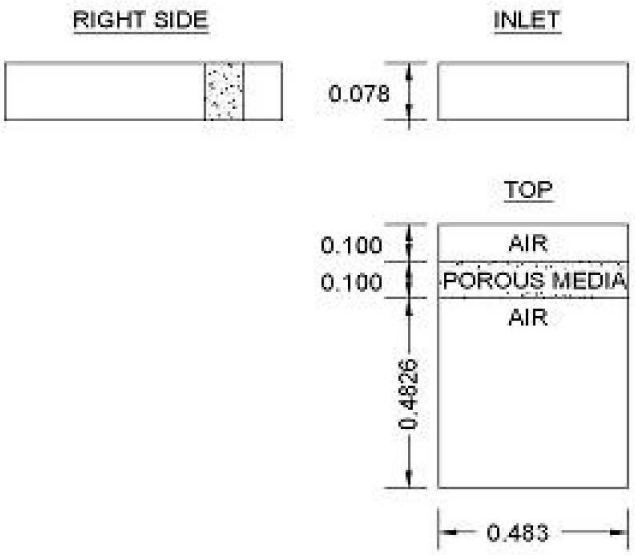

Figure A 40: Example layout of a porous media model of a server.

The volume, heat capacity, density, and porosity all affect the time constant of the server model. It is necessary to fix all but one of these parameters and then tune the model to achieve the desired response. If possible, the parameters that are fixed should be chosen to represent the physical properties of the server. The tuning process proposed mocks the server time constant experiment. Server inlet velocity should be set to match the experimentally determined (or otherwise estimated) value. The outlet boundary condition could be set to pressure outlet. The mesh size need not be fine; the details of airflow patterns are not as important as the mass averaged outlet temperature. A mesh study should be performed to determine the smallest mesh that yields consistent results. The model should be run in steady state until the desired convergence is obtained. The model should then be run in a dynamic simulation. The inlet temperature boundary condition can be changed to produce a step change in inlet temperature. The response of the outlet temperature will be first order. The size of the time step is most critical right after the step change, since the slope of the curve will be sharpest at that point. After the slope decreases, the time step size can be increased. After multiple runs, the time constant can be plotted versus the variable property and a best fit line can be used to determine what value should be used.

Once the desired time constant is achieved, the model should be validated in a room level simulation. The geometry of the room level model should mimic that of the

control volume experimental setup. However, if meshing proves problematic, it is more important model the correct volume of air be included and less important to model every detail of the airflow. A starting point for the geometry of the room level model is shown in Figure 4.1-1. Simplifications, such as the elimination of sharp, triangular edges, may be necessary to aid in convergence. Boundary conditions should be chosen to simulate those of the CRAC unit. It should be determined whether it is necessary to simulate conduction through the walls of the control volume. The model should first be run at steady state. Then a boundary condition may be changed within a dynamic simulation in order to simulate one of the control volume experiments. Control volume experiments 2 and 3 may require user-defined functions for simulation. After simulation results are obtained, they can be compared with the experimental results from the control volume experiment in question. Error bounds for the validation might best be obtained from Appendix A.6 and the associated discussion within Section 4.2.

Appendix A.11: Fitting and Error Analysis of Server Time Constant Curves

Fully Loaded Head Node:

```
NUMERIC = xlsread('server_tau_jedi_2U_3rd_run','Sheet1');
t=NUMERIC(:,1);
inlet=NUMERIC(:,19);
T_proc=NUMERIC(:,4);
T_pwr=NUMERIC(:,8);


%%PROCESSOR ANALYSIS
t_0 = 430;
T_min_proc = 18;
T_max_proc = 56;

delta_T_max_proc=T_max_proc-T_min_proc;
delta_T_proc=T_proc(44:94)-T_min_proc;
theta_proc=delta_T_proc./delta_T_max_proc;
x_proc=log(1-theta_proc);
y=t(44:94)-t_0;
tau_proc=-inv(x_proc'*x_proc)*x_proc'*y
T_proc_an=T_min_proc+delta_T_max_proc*(1-exp(-(t-t_0)/tau_proc));

%error estimate
%sum of squared in x
%actual-predicted
x_bar_proc=log(1-(T_proc_an(44:94)-T_min_proc)/(T_max_proc-
T_min_proc));
res_x_proc=x_proc-x_bar_proc;
res_x_sq_proc=res_x_proc.^2;
ss_x_proc=sum(res_x_sq_proc);
den_proc=size(res_x_sq_proc);

ss_T_proc=sum((T_proc(44:94)-T_proc_an(44:94)).^2);
sigma_T_proc=sqrt(ss_T_proc/den_proc(1))
conf_proc_95=2*sigma_T_proc

percent_T_proc=conf_proc_95/delta_T_max_proc*100

%%POWER SUPPLY ANALYSIS
T_min_pwr = 17;
T_max_pwr = 52;

delta_T_max_pwr=T_max_pwr-T_min_pwr;
delta_T_pwr=T_pwr(44:94)-T_min_pwr;
theta_pwr=delta_T_pwr./delta_T_max_pwr;
x_pwr=log(1-theta_pwr);
y=t(44:94)-t_0;
tau_pwr=-inv(x_pwr'*x_pwr)*x_pwr'*y
```

```
T_pwr_an=T_min_pwr+delta_T_max_pwr*(1-exp(-(t-t_0)/tau_pwr));

%error estimate
%sum of squared in x
%actual-predicted
x_bar_pwr=log(1-(T_pwr_an(44:94)-T_min_pwr)/(T_max_pwr-T_min_pwr));
res_x_pwr=x_pwr-x_bar_pwr;
res_x_sq_pwr=res_x_pwr.^2;
ss_x_pwr=sum(res_x_sq_pwr);
den_pwr=size(res_x_sq_pwr);

ss_T_pwr=sum((T_pwr(44:94)-T_pwr_an(44:94)).^2);
sigma_T_pwr=sqrt(ss_T_pwr/den_pwr(1))
conf_pwr_95=2*sigma_T_pwr

percent_T_pwr=conf_pwr_95/delta_T_max_pwr*100

plot(t(35:300),inlet(35:300),t(35:300),T_proc(35:300),t(35:300),T_pwr(3
5:300),t(44:300),T_proc_an(44:300),t(44:300),T_pwr_an(44:300))
axis tight
legend('Inlet','Processor Outlet','Power Supply Outlet','Processor
Fit','Power Supply Fit','location','southeast')
title('Jedi Back Node Time Constants: Processor and Power Supply')
xlabel('Time [s]')
ylabel('Temperarure [C]')


%%RESULTS

% tau_proc =  342.0743
% sigma_T_proc =    0.5858
% conf_proc_95 =    1.1715
% percent_T_proc =    3.0830
%
% tau_pwr =  382.4745
% sigma_T_pwr =    2.4911
% conf_pwr_95 =    4.9822
% percent_T_pwr =   14.2348
```

Back Node with Idle Processors:

```
NUMERIC = xlsread('server_tau_jedi_2U_3rd_run','Sheet1');
t=NUMERIC(:,1);
inlet_proc=NUMERIC(:,22);
inlet_pwr=NUMERIC(:,23);
T_proc=NUMERIC(:,10);
T_pwr=NUMERIC(:,11);


%%PROCESSOR ANALYSIS
t_0 = 430;
T_min_proc = 17;
T_max_proc = 56.75;

delta_T_max_proc=T_max_proc-T_min_proc;
delta_T_proc=T_proc(44:94)-T_min_proc;
theta_proc=delta_T_proc./delta_T_max_proc;
x_proc=log(1-theta_proc);
y=t(44:94)-t_0;
tau_proc=-inv(x_proc'*x_proc)*x_proc'*y
T_proc_an=T_min_proc+delta_T_max_proc*(1-exp(-(t-t_0)/tau_proc));

%error estimate
%sum of squared in x
%actual-predicted
x_bar_proc=log(1-(T_proc_an(44:94)-T_min_proc)/(T_max_proc-
T_min_proc));
res_x_proc=x_proc-x_bar_proc;
res_x_sq_proc=res_x_proc.^2;
ss_x_proc=sum(res_x_sq_proc);
den_proc=size(res_x_sq_proc);

ss_T_proc=sum((T_proc(44:94)-T_proc_an(44:94)).^2);
sigma_T_proc=sqrt(ss_T_proc/den_proc(1))
conf_proc_95=2*sigma_T_proc

percent_T_proc=conf_proc_95/delta_T_max_proc*100


%%POWER SUPPLY ANALYSIS
T_min_pwr = 17;
T_max_pwr = 49.5;

delta_T_max_pwr=T_max_pwr-T_min_pwr;
delta_T_pwr=T_pwr(44:94)-T_min_pwr;
theta_pwr=delta_T_pwr./delta_T_max_pwr;
x_pwr=log(1-theta_pwr);
y=t(44:94)-t_0;
tau_pwr=-inv(x_pwr'*x_pwr)*x_pwr'*y
T_pwr_an=T_min_pwr+delta_T_max_pwr*(1-exp(-(t-t_0)/tau_pwr));

%error estimate
```

```
%sum of squared in x
%actual-predicted
x_bar_pwr=log(1-(T_pwr_an(44:94)-T_min_pwr)/(T_max_pwr-T_min_pwr));
res_x_pwr=x_pwr-x_bar_pwr;
res_x_sq_pwr=res_x_pwr.^2;
ss_x_pwr=sum(res_x_sq_pwr);
den_pwr=size(res_x_sq_pwr);

ss_T_pwr=sum((T_pwr(44:94)-T_pwr_an(44:94)).^2);
sigma_T_pwr=sqrt(ss_T_pwr/den_pwr(1))
conf_pwr_95=2*sigma_T_pwr

percent_T_pwr=conf_pwr_95/delta_T_max_pwr*100


plot(t(35:300),inlet_proc(35:300),t(35:300),inlet_pwr(35:300),t(35:300)
,T_proc(35:300),t(35:300),T_pwr(35:300),t(44:300),T_proc_an(44:300),t(4
4:300),T_pwr_an(44:300))
axis tight
legend('Processor Inlet','Power Supply Inlet','Processor Outlet','Power
Supply Outlet','Processor Fit','Power Supply
Fit','location','southeast')
title('Jedi Back Node Time Constants: Processor and Power Supply')
xlabel('Time [s]')
ylabel('Temperarure [C]')

%%RESULTS

% tau_proc =  365.5931
% sigma_T_proc =    0.6523
% conf_proc_95 =    1.3046
% percent_T_proc =    3.2820
%
% tau_pwr =  296.4472
% sigma_T_pwr =    0.6867
% conf_pwr_95 =    1.3733
% percent_T_pwr =    4.2257
```

2U Modern Server with Processors Fully Loaded:

```
NUMERIC = xlsread('server_tau_jedi_2U_3rd_run','Sheet1');
t=NUMERIC(:,1);
inlet_proc=NUMERIC(:,15);
inlet_pwr=NUMERIC(:,16);
T_proc=NUMERIC(:,2);
T_pwr=NUMERIC(:,3);


%%PROCESSOR ANALYSIS
t_0 = 430;
T_min_proc = 16.4444;
T_max_proc = 55;

delta_T_max_proc=T_max_proc-T_min_proc;
delta_T_proc=T_proc(44:64)-T_min_proc;
theta_proc=delta_T_proc./delta_T_max_proc;
x_proc=log(1-theta_proc);
y=t(44:64)-t_0;
tau_proc=-inv(x_proc'*x_proc)*x_proc'*y
T_proc_an=T_min_proc+delta_T_max_proc*(1-exp(-(t-t_0)/tau_proc));

%error estimate
%sum of squared in x
%actual-predicted
x_bar_proc=log(1-(T_proc_an(44:64)-T_min_proc)/(T_max_proc-
T_min_proc));
res_x_proc=x_proc-x_bar_proc;
res_x_sq_proc=res_x_proc.^2;
ss_x_proc=sum(res_x_sq_proc);
den_proc=size(res_x_sq_proc);

ss_T_proc=sum((T_proc(44:64)-T_proc_an(44:64)).^2);
sigma_T_proc=sqrt(ss_T_proc/den_proc(1))
conf_proc_95=2*sigma_T_proc

percent_T_proc=conf_proc_95/delta_T_max_proc*100


%%POWER SUPPLY ANALYSIS
T_min_pwr = 27.556;
T_max_pwr = 65;

delta_T_max_pwr=T_max_pwr-T_min_pwr;
delta_T_pwr=T_pwr(44:94)-T_min_pwr;
theta_pwr=delta_T_pwr./delta_T_max_pwr;
x_pwr=log(1-theta_pwr);
y=t(44:94)-t_0;
tau_pwr=-inv(x_pwr'*x_pwr)*x_pwr'*y
T_pwr_an=T_min_pwr+delta_T_max_pwr*(1-exp(-(t-t_0)/tau_pwr));

%error estimate
```

```
%sum of squared in x
%actual-predicted
x_bar_pwr=log(1-(T_pwr_an(44:94)-T_min_pwr)/(T_max_pwr-T_min_pwr));
res_x_pwr=x_pwr-x_bar_pwr;
res_x_sq_pwr=res_x_pwr.^2;
ss_x_pwr=sum(res_x_sq_pwr);
den_pwr=size(res_x_sq_pwr);

ss_T_pwr=sum((T_pwr(44:94)-T_pwr_an(44:94)).^2);
sigma_T_pwr=sqrt(ss_T_pwr/den_pwr(1))
conf_pwr_95=2*sigma_T_pwr

percent_T_pwr=conf_pwr_95/delta_T_max_pwr*100


plot(t(35:300),inlet_proc(35:300),t(35:300),inlet_pwr(35:300),t(35:300)
,T_proc(35:300),t(35:300),T_pwr(35:300),t(44:300),T_proc_an(44:300),t(4
4:300),T_pwr_an(44:300))
axis tight
legend('Processor Inlet','Power Supply Inlet','Processor Outlet','Power
Supply Outlet','Processor Fit','Power Supply
Fit','location','southeast')
title('Jedi Head Node Time Constants: Processor and Power Supply')
xlabel('Time [s]')
ylabel('Temperarure [C]')

%%RESULTS

% tau_proc =  131.8339
% sigma_T_proc =    3.5526
% conf_proc_95 =    7.1052
% percent_T_proc =   18.4284
%
% tau_pwr =  994.0129
% sigma_T_pwr =    0.3139
% conf_pwr_95 =    0.6278
% percent_T_pwr =    1.6765
```

References

1. ASHRAE, *Thermal Guidelines for Data Processing Environments*. 2004, ASHRAE.

2. AFCOM. 2006. Five Bold Predictions For The Data Center Industry That Will Change Your Future [Keynote Slides]. AFCOM Data Center Institute. March. http://www.afcom.com/files/PDF/DCI_Keynote_Final.pdf . (accessed April 6, 2007).

3. US EPA Energy Star Program. Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431. August 2, 2007.

4. Beitelmal, A. and Patel, C., Thermo-Fluids Provisioning of a High Density Data Center, Distributed and Parallel Databases, v 21, n 2-3, June, 2007, High Density Data Centers, p 227-238, Kluwer Academic Publishers, 2007.

5. Turner, W., Seader, J., and Brill, K., Industry Standard Tier Classifications Define Site Infrastructure Performance. ©2001-2005 The Uptime Institute Inc. http://upsite.com/TUIpages/tuiwhite.html. (Accessed September 3, 2007).

6. Consortium for Energy Efficient Thermal Management. http://www.me.gatech.edu/CEETHERM/Downloads/Files/CEETHERM_lab_info .pdf. (Accessed October 20 2007).

7. Nelson, G. Development of and Experimentally Validated Compact Model of a Server Rack. Master's Thesis: Georgia Institute of Technology. Dec, 2007.

8. GE Sensing. TransPort® PT878GC Panametrics Ultrasonic Portable Gas Flowmeter. http://www.gesensing.com/products/resources/datasheets/pt878gc.pdf (Accessed October 25, 2007).

9. Pearson, J. Leonard, R. McCutchan, R. Gain and Time Constant for Finned Serpentine Crossflow Heat Exchangers. ASHRAE Transactions 80 (2) (1974) 255-267.

10. Kim, Y. An Experimental Study on the Quantitative Interpretation of Local Convective Heat Transfer for a Plate Fin and Tube Heat Exchanger using the Lumped Capacitance Method. International Journal of Heat and Mass Transfer. Vol 49. Pg 230-239. 2006.

11. Incropera, F. et al. Fundamentals of Heat and Mass Transfer. 6th ed. Hoboken, NJ: John Wiley & Sons, 2007. pp. 256-270.

12. Telecordia Technologies, Thermal Management in Telecommunications Central Offices: GR-3028-CORE, p 5-10 thru 5-13, Compact Disc, 2001.

13. U. S. Department of Energy. "Energy Plus Simulation Software." http://apps1.eere.energy.gov/buildings/energyplus/.

14. Liebert: Precision Cooling. Libiert Deluxe System 3 System Design Manual – 50 & 60 Hz. http://shared.liebert.com/SharedDocuments/Manuals/sl_18110826.pdf

15. Somani, A. Advanced Thermal Management Strategies for Energy Efficient Data Centers. Master's Thesis: Georgia Institute of Technology. Dec, 2008.

16. Hawick, K. "Commodity Cluster Computing for Computational Chemistry." University of Adelaide. http://www.dhpc.adelaide.edu.au/reports/073/html/.

17. Cray, Inc. " Cray-2." http://www.craysupercomputers.com/cray2.htm

18. “Press Room.” Capital Outsourcing.” http://www.c-o.com/NewsDetails.aspx?ID=457&PID=340. (Accessed November 7, 2007).

19. Madara, S. “FAQ: How do I Cool High Density Racks.” Searchdatacenter.com. http://searchdatacenter.techtarget.com/generic/0,295582,sid80_gci1241135,00.html. (Accessed November 15, 2007).

20. Schroff. Effective heat dissipation from server cabinets**.** http://www.varistar.co.uk/coldaislecontainment.html. (Accessed November 18, 2007).

21. KD Electric, Inc. “Generators and UPS Projects.” http://www.4teamkd.com/generators.php. (Accessed November 23, 2007).

22. Tokyo Electric Power Company. “Network.” Power Supply and Network. http://www.tepco.co.jp/en/challenge/energy/images/03-03b.jpg. (Accessed December 15, 2007).

23. Golson, J. “How Rackspace Really Went Down.” Vallywag. http://valleywag.com/tech/confirmed/how-rackspace-really-went-down-322828.php. (Accessed January 4, 2008).

24. Patterson, M.The Effect of Data Center Temperature on Energy Efficiency. Intel Corporation. Itherm 2008.

25. Samadini, E. Numerical Modeling of Perforated Tile Flow Distribution in a Raised-Floor Data Center. Proceedings of IPACK2007. ASME InterPACK '07, Vancouver, British Columbia, Canada. 2007

26. Kankari, K. “CoolSim.” AFCOM Chapter Meeting. Lecture: November 14, 2007.

27. APC. “Netshelter SX Enclosures: Technical Specifications.” Apc.com. http://www.apc.com/resource/include/techspec_index.cfm?base_sku=AR3100. (Accessed January 12, 2008).

28. Blackshear, E. The evolution of build-up package technology and its design challenges. IBM Journal of Reserch and Development. 2005. http://www.research.ibm.com/journal/rd/494/black11.gif.. (Accessed January 27, 2008).

29. The Uptime Institute; see the whitepaper, “Heat Density Trends in Data Processing, Computer Systems and Telecommunication Equipment,” *www.uptimeinstitute.org*, 2000. (Accessed March 29, 2008).

30. Schmidt, R. “Challenges of Data Center Thermal Management.” IBM Journal of Research and Development. 2005. http://www.research.ibm.com/journal/rd/494/schmidt.html. (Accessed June 8, 2008).

31. Hiles, A. “Five Nines: Chasing the Dream.” Continutity Central. http://www.continuitycentral.com/feature0267.htm. (Accessed August 13, 2008).

32. Dell. “System Overview: Dell™ PowerEdge™ 8450 Systems Service Manual.” http://support.dell.com/support/. (Accessed November 17, 2008).

33. Ziegler, J. “M02: Mechanical Plan.” Georgia Institute of Technology: CEETHERM Test Lab. Construction Documents. 2004.

34. ASHRAE TC 9.9. “2008 ASHRAE Environmental Guidelines for Datacom Equipment.” http://tc99.ashraetcs.org. 2008. (Accessed October 1, 2008).

35. Nie. Q. Experimentally Validated Multiscale Thermal Modeling of Electronic Cabinets. PhD Thesis: Georgia Institute of Technology. Aug, 2008. (Accessed December 10, 2008).

36. TSI. “Velocicalc Air Velocity Meters.” http://www.tsi.com. (Accessed December 28, 2008).