**iTad: An I/O-Aware Thermal Model for
Data Centers**

by

Tausif Muzaffar

A thesis submitted to the Graduate Faculty of
Auburn University
in partial fulfillment of the
requirements for the Degree of
Master of Science

Auburn, Alabama
May 4, 2014

Keywords: Data Centers, Thermal Energy, Heat Model

Approved by

Xiao Qin, Chair, Associate Professor of Computer Science Software Engineering
Cheryl Seals, Associate Professor of Computer Science Software Engineering
Alvin Lim, P Associate Professor of Computer Science Software Engineering

Abstract

With the ever-growing cooling costs of large-scale data centers, thermal management must be adequately addressed. Thermal models can play a critical role in thermal management that helps in reducing cooling costs in data centers. However, existing thermal models for data centers can overload I/O activities. To address this issue, we developed an I/O-aware thermal model called iTad for data centers. The iTad model captures the thermal characteristics of servers in a data center, offering a much finer granularity than the existing models. In addition to CPU workloads, iTad incorporates the I/O load in order to accurately estimate the thermal footprint of the servers with I/O-intensive activities. We validate the accuracy of the iTad model using real-world temperature measurements acquired by an infrared thermometer. Our empirical results show that I/O utilizations have a significant impact on internal temperatures of data servers. We show that thermal management mechanisms can quickly retrieve the thermal information of servers from iTad before making important workload placement decisions in a real-time manner.

Acknowledgments

As my studies have gone the more humbled I am but the generosity and support to those in my life. First I would like thank Dr. Qin for his belief in me that encouraged me to pursue higher education. Without his support and guidance none of this would be possible.

Of course I would also like to thank all the great professor here Auburn CSSE department, for challenging me and teaching me all the things I know about computer science today. A special thanks of course to Dr. Seals and Dr. Lim for their work at teaching me but also serving as committee members for my defense.

I would be doing an injustice not to thank all those in my lab group for their on going support and guidance, especially those I have known since the beginning of my masters program: Xunfei, Sanjay, Ajit.

Lastly I would like to thank God and my family for all their support and protection. My hope is that my actions and this thesis is done in the way that would please them.

Table of Contents

List of Figures

List of Tables

Auburn Auburn University

LoA   List of Abbreviations

iTad: I/O Thermal Aware Datacenter

I/O: Input and Output

CPU: Central processing unit

CRAC: Computer room air conditioning

Chapter 1

Introduction

Recent studies show that thermal management is an important issue to data centers due to ever-increasing cooling cost [24]. Cooling costs contribute to a significant portion of the operational cost of large-scale data centers; therefore, increasing the size of a data center leads to huge amount of energy consumed by the center's cooling system. An efficient way to combat the high cost of cooling systems is to develop thermal-aware management techniques that place jobs and data on servers to minimize temperatures of data centers.

Thermal management aims at reducing cooling costs of data centers; thermal management mechanisms largely rely on thermal information to make intelligent job and data placement decisions. Thermal information can be acquired in the following three means:

1. Temperature sensors measure inlet and outlet temperatures of servers

2. Computational fluid dynamics simulators (see, for example, Flovent) simulate temperatures of servers in data centers

3. Thermal models estimate a server's temperature based on the server's workloads.

After looking through these options we decided to create an CPU and I/O aware thermal model called iTad. iTad standing for I/O Thermal Aware Data center. The reason we decided to make such a model than use the other two options will be explained in the following sections.

1.1 Reducing Monitoring Cost

The first approach is to monitor server inlet temperatures by deploying sensors in a number of locations in a data center [22]. This approach faces a dilemma; while high levels

of accuracy can be achieved by increasing the number of sensors, this leads to an expensive monitoring solution. Reversely reducing the number of sensors may cause inaccuracies, and an algorithm would need to be developed to extrapolate the heat from individual nodes thus taking away the simplicity that makes this route so appealing. For large-scale data centers, this approach is not very practical for two reasons. First, it is prohibitively expensive to deploy hundreds of thousands of sensors to offer accurate temperature measurements. Each server needs at least two sensors; each sensor may cost up to $100 [7]. Second, wiring and maintenance cost of the large number of sensors can further increase the operational cost of data centers.

1.2 Reducing Monitoring Time

To reduce the high cost of deploying an excessive number of sensors, data center managers can make use of the computational fluid dynamics simulators to simulate and collect inlet temperatures of servers [21]. Although this simulation approach offers accurate thermal information at low cost without employing any sensor, it is time consuming (e.g., several hours) to run each simulation study. Thus, the simulation studies must be conducted offline, indicating that thermal management mechanisms are unable to retrieve thermal information from the simulators at run-time.

1.3 Benefits of Thermal Model

Thermal models are arguably a more promising approach to providing thermal management mechanisms; they can provide temperature information of servers at run-time without incurring any cost to purchase and maintain sensors. Thermal models offer the following four major benefits for data centers. First, thermal models significantly reduce thermal monitoring costs. Second, unlike thermal simulators, thermal models offer temperature information to thermal management schemes in a real-time manner. For example, our iTad thermal model is able to profile the thermal characteristics of a data center in a matter of

Figure 1.1: Temperature of Processor when CPU utilization 100% vs Time [26].
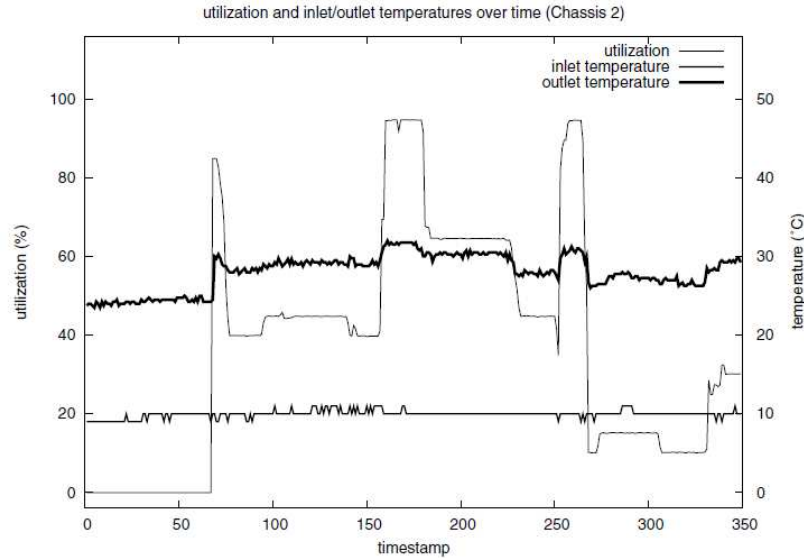
seconds. Third, thermal management powered by thermal models helps cut cooling costs and boosts system reliability. Last, thermal models allows data center designers to quickly make intelligent decisions on thermal management in an early design phase.

Most existing thermal models in the market treat servers as a uniform black box because it is unclear what all factors are involved in the heat distribution of a data center [20]. There are a few thermal models (see, for example, [26]) that can derive power consumption and necessary cool power from inlet and outlet temperatures of servers. However, implementing these models requires (1) many thermometers and (2) the management of thermal information in real time (3)only based on CPU work. CPU Workload has been known to affect thermal load, for example Figure 1.1 shows that when CPU utilization is increased, there is a large increase in temperature [26], but it neglects I/O workload. I/O-intensive activities in servers are commonly overlooked in these models. One of the goals of this research is to demonstrate that I/O utilization plays an important role in a server's thermal dissipation. We believe I/O-intensive applications running in data centers impose heavy load on servers, making disks of the servers hot-spots.

3

1.4 Contributions

Contributions. The major contribution of this study are summarized as follows:

1. We develop the iTad thermal model that provides outlet temperatures of servers in a data center. We show that both CPU and I/O thermal outputs can be extrapolated from radiation heat and convection heat applied to a server. With iTad in place, thermal management schemes can quickly make workload management decisions at run-time based on I/O and CPU utilizations.

2. We validate the accuracy of the iTad model using a server's real-world temperature measurements obtained by an infrared thermometer (Figure 1.2).

Figure 1.2: Thermometer used in testing

3. Our experimental results suggest that iTad is an accurate model to derive server outlet temperatures according to I/O and CPU activities.

4. We show that this model is easily can be plugged into any data centers.

5. We analytically study the relationship between I/O load and server outlet temperatures. Our analysis confirms that I/O-intensive workloads have significant impact on temperatures of servers.

Chapter 2

Related Work

2.1 Energy-Efficient Data Centers

Large data has become on the hottest topics in computer. With refocus there has been more interest in energy efficient data centers [10] [11], because a recent study shows that 1.2% of all energy consumption in U.S. is attributed to data centers [17]. To minimize the effect of the data centers on the national consumption there has been many energy-saving approaches, one that relates to this research is the work of Bieswanger et al. where they deploy sensors to analyze the power consumption instantaneously, our research deals with real-time thermal management has some overlap [8].

2.2 Thermal Aware Data Centers

Energy aware data centers has been the classical way of thinking about reducing the effect of data centers on the environment. Another school of though is if we manage the thermal outputs of the data centers, thus reducing the cooling cost we can have the same impact that as energy efficient data centers. [25]

2.3 Thermal Simulations

Most of the research related to thermal management in data centers use a commercial simulation software *FloVent*, which provides detailed 3D visualization of airflow and temperature throughout the server room [2]. It can get very accurate heat recirculation results. The downside is that, it is very complicated to setup or configure and it takes huge amount of time to run each simulation. Such software is very useful for machine learning because

of the time needed to implement machine learning techniques but not very effective on split second decision making. We use iTad to implement a low cost and less time consuming management technique.

2.4 Thermal Models

Eibeck *et al* [13] developed a model to predict the transient temperature profile of an IBM 5-1/4-in. fixed disk drives by experimentally determining the thermal characteristics of the disk drive. Tan *et al* presented a 3D finite element modeling technique to predict the transient temperature under frequent seeking [23]. Gurumurthi *et al* investigated the thermal behavior of the hard disk and presented an integrated disk model. Their model calculates the heat generated from the physical components of the disk drive like spindle motor, voice-coil motor and disk arms [14]. Kim *et al* studied thermal behavior of disks by varying the platter types and number of platters and established a relationship between seek time and the disk temperature [16]. However, the impact of the disks utilization on the disk temperature and contribution of disks to the outlet temperature of nodes have not been investigated. Even though clearly thermal footprints of computing has a breathe of research.

Microsoft research and Carnegie Mellon University [20] presented a model which predicts the future temperature of servers through machine learning. As this model relies on the sensor data, it will be costlier for large data centers to buy large amount of sensors. In our research, instead of predicting future temperature we want a model to calculate the current temperature based on the workload without using sensors.

Tang et al [27] [26] developed an interesting model demonstrating the effect of heat recirculation on the inlet temperature of servers in a data center, and in turn, on the efficiency of cooling system. They calculate inlet temperature of servers based on the temperature of the air supplied by the cooling room air condition (CRAC) and CPU utilization. Li et al [19] showed that CPU intensive applications cause dramatic heat change for processor. We believe that data intensive applications running in data centers will have the similar effect

6

on the disks of storage nodes, which has to be taken into account while calculating the total heat generated by the node. Kozyrakis [12] studied the effect of different application and observed the power consumption of the nodes. It showed that disk and memory consumes significant amount of power, even as compared to CPU (as shown in Example 2). As power consumption has direct impact on heat generated, there is a need to investigate the thermal load of I/O intensive applications on the nodes in data centers.

Chapter 3

Methodology

3.1   Determine Recirculation Factors

We achieve the aforementioned goal by focusing on heat recirculation of active data centers. Figure 3.1 depicts a general model for a data center, where each blade server's outlet temperature affects room temperatures. The outlet temperature of the server depends on its inlet air that enters the front of the server's rack. The inlet air temperature is the computer room temperature cooled by an air conditioning system.

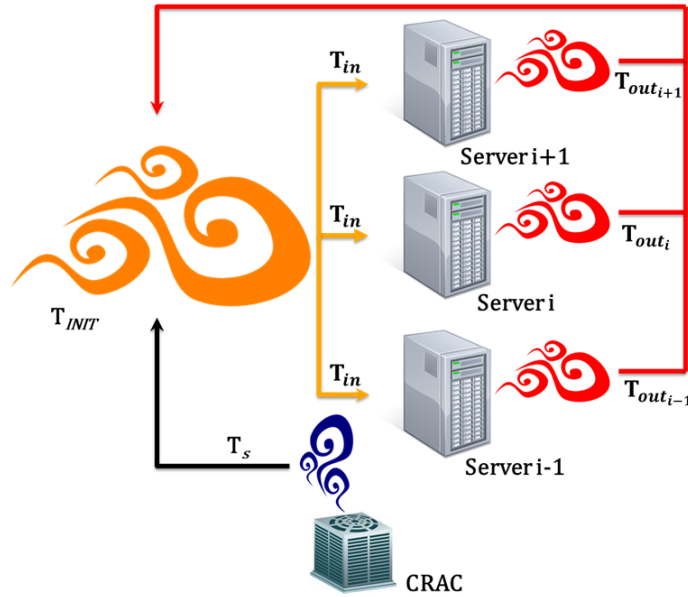Figure 3.1: An Overview of a data center.

Figure 3.1 shows that heat recirculation in a data center can be derived as the sum of each server's outlet temperature. To build a model representing the heat recirculation of a data center (see Figure 3.1), we start this study by paying attention to constructing a thermal model for each individual server. Here we are using the assumption that since recirculation

is the sum of single servers, then if we can model a single server, we just need to apply it to all the servers in the data center and add it all together. Essentially we will modeling two different things, the first will be the server heat transfer and the the inlet temperatures value before the server heat transfer. Figure 3.2 shows how the initial temperature will feed into our first model and they feeds into our inlet temperature model.

Figure 3.2: Model Overview

In our iTad model, there are three components (see Figure 3.3, where T_{out} denote outlet temperature) affecting the outlet temperatures of a blade server. These three affecting factors are inlet temperature, CPU utilization, and I/O workload.

Figure 3.3: Three factors affect the outlet temperature of a single blade server.

Our iTad model makes use of these factors to estimate the outlet temperature for server i, thereby enabling thermal management schemes to place workloads to control outlet temperatures. The iTad model is orthogonal to existing thermal management schemes; iTad can be seamlessly integrated with any thermal management scheme to either minimize outlet temperatures or minimize heat recirculation in a data center. In this study, we focus on the

9

accuracy of iTad by validating it against real-world temperature measured by an infrared thermometer.

Figure 3.4: Three factors affect the inlet temperature of a single blade server.

A challenge in the development of iTad is the measurements of inlet temperatures of servers. More specifically, Figure 3.1 indicates that the air entering the servers is not equivalent to initial temperature. Rather, the inlet temperature equals to the initial temperature subtracted by some factor of air supplied by the air conditioning system. The inlet temperature of a server is affected by three factors, namely, computer room temperature, cooling supply air temperature, and the outlet temperatures of other servers (see Figure 3.4). For this model we decided to model only the current server outlet temperature an instantaneous moment so its the only one that affects input temperatures. In one of our current studies, we are extending the iTad model to investigate the heat recirculation effect by considering the impact of all nodes outlet temperatures on inlet temperatures.

3.2 Determine Hardware Factors

After dealing with actual inlet temperatures, we incorporate I/O and CPU workloads into iTad. In this part of the study, we show how the outlet temperature of a server changes based on I/O-intensive activities. The iTad model has to deal with heat transfer, especially convection heat transfer. Convection heat transfer [28] is based on temperature and specific heat, all of which have a linear relationship. A study conducted by Barra and Ellzey demonstrates how a wide range of shapes affect heat transfer [9]. iTad is the first model

that attempts to incorporate I/O-intensive workload therefore, we consider cases where all the components in a data center have the same transfer rate. Nevertheless, we do not imply by any means that all the components have an identical transfer rate. In our future work, we will extend iTad to consider multiple heat transfer rates to further improve the accuracy of iTad.

The iTad model helps in improving the energy efficiency of data centers because thermal information offered by iTad assists dynamic thermal management to reduce the energy consumption in cooling systems in data centers. We show that thermal management mechanisms can quickly make workload placement decisions based on thermal information facilitated by iTad.

Chapter 4

Modeling

4.1 Assumptions and Notation -

We described the plan of our model as well as the basic components necessary for the model. In this section, we will present the assumptions and the notations we used in the model. Following are the assumptions :

1. Initial temperature is always consistent throughout the data center.

2. The air flow is static in all parts of the data center.

3. Supplied temperature strength is linearly proportional to the distance from the vent.

4. Our model is models temperature at an instantaneous moment so nothing is being circulated in our model.

5. The adjacent nodes will not heat up enough to cause an effect to the node in question.

6. PC components are all similar in shape so the heat transfer is consistent.

7. The entire experiment is based on the premise that taking a single node from a cluster and running our experiments we can grasp the important factors in thermal change in computers. With this information we will able to model large scale environment.

After laying out the assumptions, the notations used in the model are described in Table 4.1.

This equation can be reorganized to solve for outlet temperature.

Table 4.1: Model Notation

Variables	Description
i	Number of Server Node
Q	Heat generated (J)
p	Density of air (kg/m^3)
f	Flow rate (m^3/s)
c_p	Specific Heat (J/kg/c)
T_{out}	Outlet Temperature (c)
T_{in}	Inlet Temperature (c)
T	Change in temperature (c)
h_r	Heat Transfer Coffiecnt (J/s*m^2*c)
A	Surface area of PC components(m^2)
Z	Percent of added temperature after workload
R	Ratio of distance
k	The amount outlet affects inlet temperature
d_i	Distance of the server from AC vent (m)
d	Height of room (m)
T_{INIT}	Room Initial Temperature (c)
T_s	Supplied temperature from CRAC (c)
$T_{workload}, T_w$	Surface Temperature at a workload (c)
T_{idle}	Surface Temperature at a idle (c)
W	Workload supplied (%)
T_{Max}	Max Temperature the components(c)

4.2 Modeling impacts of heat on temperature

The heat transfer in a data center node can be expressed by Equation 4.1 [20] [26] [24]. There are two kinds of heat transfer in this system: convective heat transfer adn radiant heat transfer. We organized Equation 4.1 to solve for outlet temperature. In Equation 4.1, Q_i is the convective heat transfer of server i, which means as the inlet air passes through the amount of heat is builds up is Q_i.

$$Q_i = pfc_p(T_{out_i} - T_{in_i})$$

$$T_{out_i} = \frac{Q_i}{pfc_p} + T_{in_i} \tag{4.1}$$

The heat generated in the chassis is actually the heat being radiated from the components of a server, which also know as radiant heat. So, in this case, the convective heat transfer of inlet temperature and outlet temperature is equal to the radiant heat transfer of the PC components. Figure 4.1 shows you the model how all the heat that radiates off the components mixes into the air to form the outlet temperature and its convective heat gain.

Figure 4.1: Radiant heat equals convective heat

The equation 4.2 shows the formula for radiation heat transfer [1]. The radiant heat is dependent on the surface of the object and the heat it generates on its surface.

$$Q_i = h_r A \triangle T_i \tag{4.2}$$

In Equation 4.3, the $\triangle T_i$ is the change of temperature caused by the PC components, which we modeled as the change in temperature of the server at the specified workload ($\triangle T_{workload}$) plus difference between inlet and outlet temperature of server at idle state ($T_{out_{idle}} - T_{in_{idle}}$).

$$\triangle T_i = \triangle T_{workload_i} + (T_{out_{idle}} - T_{in_{idle}}) \tag{4.3}$$

14

To help simplify what we need to find we set Equation 4.1 and Equation 4.2 equal to each other to give us the variable Z, thus letting us relate T_{out} to T_{in} and $\triangle T_i$, as shown Equation 4.4

$$h_r A \triangle T_i = p f c_p (T_{out_i} - T_{in_i})$$
$$Z = \frac{h_r A}{p f c_p} = \frac{T_{out_i} - T_{in_i}}{\triangle T_i}$$
$$T_{out_i} = Z \triangle T_i + T_{in_i} \tag{4.4}$$

4.3 Modeling impacts of workload on temperature

In the article [24], they define T_{in} as dependent on T_s and a vector which models the exact strength of T_s at each height. We simplified the model further by declaring the T_{in} as the room temperature subtracted by the a percentage of the temperature supplied by the CRAC as shown in Equation 4.5.

The amount that T_{out} effects the inlet temperature is proportional to k which is something that is outside the scope of our paper. That being said the way it is implemented now, only the a current server outlet temperature will effect the outlet temperature.

$$R = \frac{d_i}{d}$$
$$T_{in} = T_{INIT} - R T_s + k T_{out} \tag{4.5}$$

Also in Equation 4.3, the other variable that defines T_{out} is $\triangle T_i$, and we modeled $\triangle T_i$ after the Figure 4.2.

The theory behind our proposed model is that some components of the server get more heated by I/O intensive applications while others get more heated by CPU intensive applications; and based on the percent of CPU or I/O utilization the components will get to

Figure 4.2: Visual representation of workload effects outlet temperature

some percentage of its maximum temperature. After the calculations of Equation 4.5 we are given $\triangle T_{workload_i}$ which is the increase in the temperature as compared to idle server.

$$\triangle T_{CPU}^{MAX} = T_{workload_{MAXCPU}} - T_{idle}$$

$$\triangle T_{I/O}^{MAX} = T_{workload_{MAXI/O}} - T_{idle}$$

$$\triangle T_{workload_i} = \triangle W_{CPU} \triangle T_{CPU}^{MAX} + \triangle W_{I/O} \triangle T_{I/O}^{MAX} \tag{4.6}$$

In the end all of these equation are the components needed in modeling a single server node. This is important because, as we discussed before, getting each single server outlet temperature can help to model a data center thermal profile. Before we can do that we need to verify that these equations are accurate.

Chapter 5

Experimental Parameters

In this section we will be determining the parameters for the model we created in the modeling section for server. Do this we need to prove that all the factors describe in the model will indeed have an effect, and then solve for the constants described in the previously in the modeling section.

5.1   Set up

Since our models, described in the previous section, model a single server node, we decided to verify the equations by setting up an experimental machine. The machine we used is an OptiPlex 360 whose specifications are listed in Table 5.1. So in this section we will define the characteristics of our machine and later use those constants to verify how accurate our models are.

Table 5.1: Server Specifications

Dimensions	15.65 x 4.59 x 14.25
RAM	1GB
Chipset	Intel G31/ICH7
DC Power Supply	255 W
Processor Type	Intel Core 2 Duo
Memory	800 MHz DDR2 SDRAM

To test our server we used a command called "stress" in Ubuntu, which can spawn multiple CPU workers or I/O workers. This process would allow us to get a estimate how a computer would act under such a load. To get an estimate of CPU utilization impact we just used "stress" to call only CPU workers [6]. To get an estimate of I/O utilization impact we just used "stress" to call only I/O workers. Finally to find a mixture of I/O and CPU

utilization impacts we called a ratio of CPU workers to I/O workers. (i.e 80 CPU workers and 20 I/O workers will be 80% CPU utilization and 20% I/O utilization).

After running our stress tests, we used 3 different tools to help design experiment. First, we used the Linux command "iostat", which gave us details about server usages. The most important pieces of information in "iostat" were the "CPU user%" which displays the percentage of CPU utilization, "system%" which displays the percentage of I/O utilization [4] [5].

Another tool we used "HDDTemp"; a software that can measure the temperature of the hard drives [3]; more of as a reassurance, to make sure our thermometer was working.

When we say thermometer, we are referring to the HDE Temperature Gun Infrared Thermometer w/ Laser Sight. This thermometer measures the surface temperature of whatever surface it is pointed on. The thermometer has a reading ratio that is 12:1, which means for every 12 cm away we have a 1 cm radius of temperature. We used this tool to measure all kinds of temperatures used for verification.

5.2   Period

Since it takes time for different components to heat up to its max temperature, we needed to test how long it would take for our each application to reach it hottest point. For CPU intensive application we can assume that the processor would be the most highly active component. So we ran our stress test at 100% CPU utilization and periodically checked the processors temperature. We plotted the temperatures over time as shown in Figure 5.2(a). If you look at Figure 5.2(a), you can see the temperature plateaus around 30 minutes, but we wait till the blue line to turn off the stress test.

In Figure 5.2(a), we did similar procedure, but we used I/O intensive application. During the I/O intensive application, instead of monitoring the processor we monitored the I/O controller. From Figure 5.2(b), it is clear that I/O application takes longer to heat up and once we turned it off, at the blue line, it takes longer to cool down. So for all our other

Figure 5.1: Utilization of Components at fixed utilization
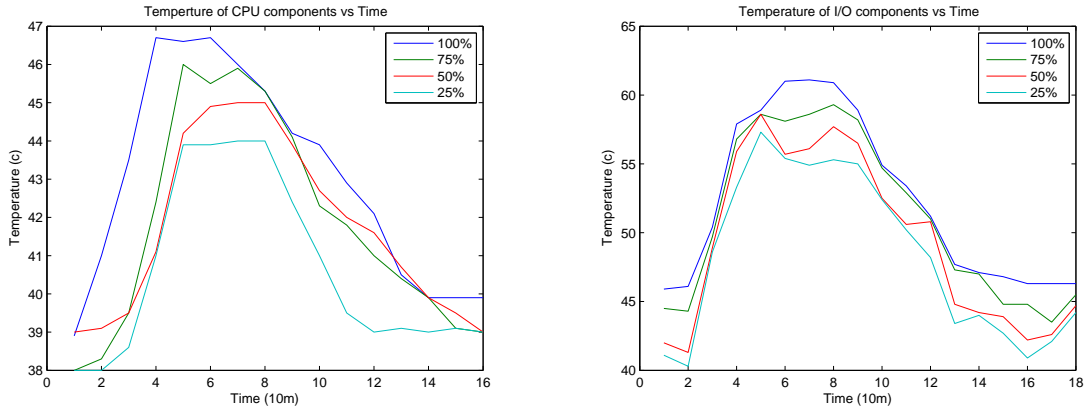
(a) Temperature of Processor when CPU utilization is changed vs Time

(b) Temperature of Processor when CPU utilization is changed vs Time

tests we run them for 1 hour before taking temperature readings, to give the CPU and I/O components ample time to heat up. We waited 1.5 hours between tests to allow server to cool down.

5.3 Determining a baseline temperature

After finding out how long it takes to run an experiment we were able to run our tests. The first experiment we needed to run was one to figure out the thermal impact of the idle machine. In order to do this, we decided to take an array of temperatures and extrapolate the information we need.

Figure 5.2 shows the insides of the our server; the numbers 1-32 are areas where we measured the temperature, 33 is the place we measured the hard drive, 34 is the place where we measured power supply, 35 is the place where measured inlet temperature and finally 36 is the outlet temperature. So once we determined what to measure, we measured each grid area with our thermometer for the server at the idle state. The measurements are given in Figure 5.4(a).

In Figure 5.3, we arranged the gathered data to graphically match Figure 5.2. In Figure5.3, the inlet temperatures are temperatures in the middle on the far left. All the

Figure 5.2: Measured Area of Temperatures

numbers in the middle row are the temperatures of grid spots of 1-32. The two temperatures on the bottom left are the temperature of the disk drive measured by two different methods, one using *HDDtemp* and other using thermometer. After gathering the temperatures we calculated the average, which we labeled off to the bottom on the far right. This gives us a baseline value, which is called T_{idle} in our model. We compared this baseline value with the other values. Another number that we needed to keep for later calculations is the difference between the idle inlet and idle outlet temperatures $(T_{out_{idle}} - T_{in_{idle}})$, which is 1.8^oC.

5.4   Impact of I/O utilization on temperature

Once we have the idle data as a reference we started to test an active machine. We started with an I/O intensive machine, and to make the machine I/O intensive we use the following command:

$stress io 3 vm 7

The ratio of –io to –vm and –io is 30% which means that the I/O is set to 30% utilization. Using this pattern we created the a figure similar to 5.3. We then allowed the computers to

Figure 5.3: Surface Temperatures

(a) Idle Temperatures

CPU: 0-5%
I/O: 0-5%
Avg: 33.7

						28.2
29.9	29.7	32.1	32	31.4	33.6	
29.8	29.9	30.3	34.2	34.9	33.4	
24.3 ... 36.9 40.2 46.4 39.4 ... 25.5
38 37 37.3 34.7
30.5	30,2	29.8	30.2	30.5	30.1
31.9 ... 29.8 29.9 29.7 29.8 29.9 30.1

(b) Max CPU Temperatures

CPU: 100%
I/O: 0-8%
Avg: 35.2

						29.9
30.6	32.6	33.1	33	33.2	33.2	
30.1	32.5	37.1	37.2	32	34.5	
27.9 ... 45.2 44.4 59.6 39.4 ... 30.4
46.7 43.2 47.3 35.1
30.1	30.1	31.6	29.4	30.9	31.2
33 ... 29.2 30.2 29.7 29.9 29.9 30.1

(c) Max I/O Temperatures

CPU: 0-10%
I/O: 100%
Avg: 34.6

						29.9
30.6	30.6	32.9	32.3	32	32.6	
30.1	31.8	35.1	34.1	31	34.2	
26.9 ... 40.2 40.4 49.7 39.4 ... 29.9
36.7 36.1 57.8 33.8
30.1	30	29.2	29.2	30.1	29.4
35 ... 29.8 29.8 30.4 31 31.2 31.4

cool down and then rechecked the temperatures at 60% utilization, then once more for 100% utilization. Although we kept and organized the data for these runs similar to the format of the idle calculations above for simplicity to show how the utilization effected the heat with Figure 5.4. Each bar represents a group of points from Figure 5.2. Table 5.2 lists each group to help interpret the results.

Table 5.2: Temperature Zones

Group Number	Number of points from Figure 5.2	Reasoning
1	34	Power Supply
2	1-12	Rarely changing
3	13,14,17,18	CPU Controls
4	15,16,19,20	I/O Controls
5	33	Hard drive

Figure 5.4: Utilization Temperature

(a) I/O Utilization

(b) CPU Utilization

(c) Mixture of CPU and I/O Utilization
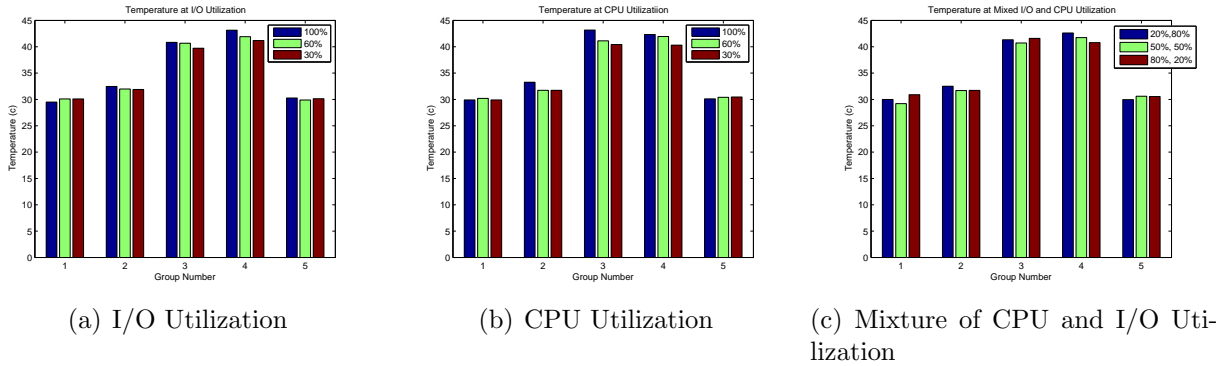
5.5 Impact of CPU utilization on temperature

After conducting the I/O tests it was time to test CPU utilization impacts. To do this
we used the stress command in the following way:

$stress cpu 3 vm 7

The ratio of –cpu to –vm and –cpu is 30% which means that the CPU is set to 30%
utilization. Using that pattern we created a figure similar to 5.3. Just as we did for I/O we
checked the temperatures at 3 different utilization levels; making sure to give the machine
time to cool down before each experiment. We found the average temperature, like we did
for the idle case and I/O, and also created a group heat graph as seen in Figure 5.5(b).
After plotting the average temperatures we were able use curve fitting techniques which we
discuss in section 5.7.

Figure 5.5: Relationship between Utilization and Outlet Temperature
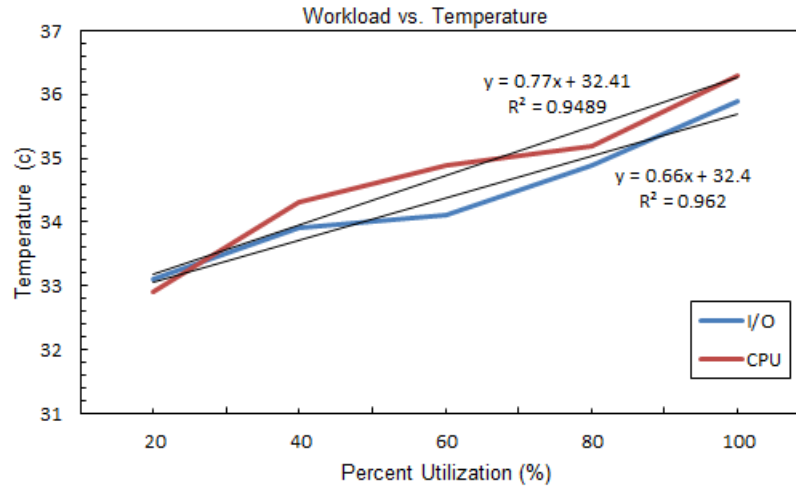
5.6 Shared I/O and CPU Utilization

Finally we decided to test a mixture of both I/O and CPU utilizations. We did this by calling the command:

$ stress cpu 50 io 50

This would set utilization for each component to 50%. So we did the same procedure as we did on the I/O intensive and CPU intensive. After running the experiments for the mixed utilization we created the Figure 5.5(c), which we will discuss in section 5.7.

5.7 Determining Constants

After our experiments, we are left with many impressions about iTad. First of all, there is a clear relation between utilization and temperature which is shown in Figures 5.5(a), 5.5(b). This relationship seemed to be linear relationship shown by the curve fitting techniques we used on the Figure (5.5) where we graph the change in outlet temperature change over the percent utilization.

After using the data from the I/O runs we were able to calculate, the slope of the line is 2.7 which represents the speed with which the temperature was increasing with R^2 value of 0.981, where R^2 represents the accuracy of the slope. In the case of CPU data, the slope of the line is 3.5 which represents the speed with which the temperature was increasing with a R^2 value of 0.961. Which indicates that a CPU intensive application will get server hotter than an I/O intensive one. This is nearly confirmed by the mixed data because it clearly shows that when CPU is higher than I/O the server is warmer.

Once we accept that the relationship is linear we can start to figure out some of the values on constants in the equation that we proposed. The proposed Equation 4.4 has consolidated all the constants of the experiment into one variable and with all our data readings we can solve for Z.

Table 5.3: Compilation of all the values gathered

I/O Intensive									
W_{io}	T_W	T_{Idle}	$\Delta T_{Workload}$	$T_{idle_{In-Out}}$	Q	T_{in}	T_{out}	T_{In-Out}	Z
30%	34.021	33.692	0.315	1.800	2.115	26.700	28.700	2.000	0.946
60%	34.237	33.692	0.630	1.800	2.430	24.700	26.900	2.200	0.905
100%	34.742	33.692	1.050	1.800	2.850	26.900	29.900	3.000	1.052

CPU Intensive									
W_{cpu}	T_W	T_{Idle}	$\Delta T_{Workload}$	$T_{idle_{In-Out}}$	Q	T_{in}	T_{out}	T_{In-Out}	Z
30%	34.039	33.692	0.442	1.800	2.242	26.100	28.300	2.200	0.981
60%	34.400	33.692	0.884	1.800	2.684	26.900	29.700	2.800	1.043
100%	35.166	33.692	1.474	1.800	3.274	27.900	31.400	3.500	1.069

I/O and CPU Intensive									
W_{cpu},W_{io}	T_W	T_{Idle}	$\Delta T_{Workload}$	$T_{idle_{In-Out}}$	Q	T_{in}	T_{out}	T_{In-Out}	Z
20%,80%	34.347	33.692	1.135	1.800	2.935	27.700	29.900	2.200	0.750
50%,50%	34.326	33.692	1.262	1.800	3.062	26.600	28.800	2.200	0.718
80%,20%	34.639	33.692	1.389	1.800	3.189	26.200	28.800	2.600	0.815

In Table 5.3, we consolidated all the information we gather while trying to determine the experimental parameters process. In Table 5.3, column 2 is the average surface temperature of the machine at that utilization, while column 3 is the average surface temperature with no utilization; the difference in column 2 and 3 is the observed difference in the values. This shows how much extra heat is generated once the server is pushed to that specific utilization. Column 4 shows what the extra temperature generated should would be, from iTad. As you can see the values are closely related. Any difference could be accounted by the change in air flow of the room or own server fans.

Column 5 of Table 5.3 is $T_{workload}$ we calculated plus the difference between inlet and outlet temperature for the idle server. This is essentially $\triangle T$ from Equation 4.4. And since we were able to actually measure the final inlet and outlet temperature for server, we were able to calculate the value of Z. The value of Z is ratio of $\triangle T$ from column 9 to the surface heat listed in column 6. As you can see, in Figure 5.6, the value of Z has an average just under 1.

Figure 5.6: Values of Z in all experiments

This simply means that the current arrangement of hardware has a one to one relationship between heat exuding from the server and outlet temperature. The change in Z for different utilization maybe an indication that the air flow is changing, but our assumption

25

is that the air flow stays constant and since the values don't vary that much it doesn't contradict iTad.

Through our results the accuracy seems close enough to say that the equations used to model the outlet temperature can work as a basis to for thermal management based on workload, or even just a starting point for future expansion.

Chapter 6

Usage

In the verification section we took our server and determined all the constants of the machine. In the following section we will discuss how accurate the numbers from the verification process are. First we would simply like to explain a use case how this model can be used a in data center.

As we you can tell by Equation 4.4 the most important value you need to solve for the the Z variable. So when building a data center you should find the Z for every machine then you can plug it into equations.

After gathering the Z values like we did in the verification section for all the machines we will need to run another task that ill keep track of CPU and I/O utilization of each machine

After setting up a monitor you take the values and plug them into Equation 4.6 and the outcomes will give the individual outlet temperature. Now this is where the model is lacking there is no heat circulation model, that will need to extended research to help control thermal output.

6.1 Verification

Now that we developed the model and got all the constants of the experiment solved the only thing left to do was to actually run iTad model on a server with random amount of CPU and I/O utilization. In Figure 6.1 you can where we ran a server for 5 hours, with utilization variation. The model had a tendency to over estimate the the temperature especially at the beginning of the process.

Figure 6.1: Verification of Model

We would account for this problem to a poor recirculation constant and the fact that the model isn't time variant. That being said the longer the machine ran the better the results were because this would return back to the state how we tested our machines to find Z value. The variation of the model and actual is a not perfect because at points our model is over 2 degrees off, but the good thing is the trend stays very close to the original so we have a model that is going to err in the safe side of calculations.

6.2   MPI

One of the goals of this experiment was to make that this model could work for any kind of data center. So if the data center was to using a C version of MPI, message passing interface, then iTad should work for them. So we set up a method called "iTad" inside an toy MPI project. iTad would return an outlet temperature, and based on that value it would make a decision about data movement. The method iTad would pull the utilization from the OS and get the other values from system to determine the its output, and all of this happened seamlessly without any noticeable change in performance.

Figure 6.2: Sample MPI Usage

```
if(iTad() < 29)
{
    if(myid != 0)
    {
        MPI_Recv(&number, 1, MPI_INT, myid-1, 0, MPI_COMM_WORLD,MPI_STATUS_IGNORE);
    } else {
        MPI_Recv(&number, 1, MPI_INT, world_size, 0, MPI_COMM_WORLD,MPI_STATUS_IGNORE);
    }
} else {
    while(iTad() > 29){
        sleep(10);
    }
    if(myid != 0)
    {
        MPI_Recv(&number, 1, MPI_INT, myid-1, 0, MPI_COMM_WORLD,MPI_STATUS_IGNORE);
    } else {
        MPI_Recv(&number, 1, MPI_INT, world_size, 0, MPI_COMM_WORLD,MPI_STATUS_IGNORE);
    }
}
```

6.3 Hadoop

We also wrote the model for JAVA, this would allow the model to used for other data centers that would use JAVA. This version had to use a runtime shell to get CPU and I/O utilization because the JAVA virtual machine doesn't have direct access to the OS information like C does. This version of the model would work for a JAVA MPI implementation, but for Hadoop each node is not responsible for its own data so we needed to see if our model could be used for an Hadoop data center. In the work of our group member paper [18] shows that the scheduler and heartbeat of Hadoop can be updated to take account of CPU and I/O utilization to make thermal decisions. So implementing our model is as straight forward as replacing there thermal method with out iTad method.

Chapter 7

Experience

7.1 Improvements

Our experiment is not beyond criticism, but no criticism that is so heinous that will crush the results of this experiment. The first critic that can be made about this experiment is that the stress command is not pushing the computer enough, especially the I/O commands seems to be inaccurate so the temperatures we are reading are off. This criticism is one that has some true concerns but the fact that when I/O intensive the I/O controller is the hottest spot proves that the I/O is being pushed. Some people may see the fact that our machine is an isolated machine not actually on a rack as concern. This in fact a true issue, while other reports [15] an increase 10 degree different in inlet and outlet with CPU utilization our machine shows a max of 1.7. This concern is invalid because our computer is prototype to monitor trends to prove our model. The next concern is the method of gathering temperature readings, in which we use the surface temperature using IR thermometer. This concern is slightly justified because of the cables and clutter on a machine the reading can be off. But each section in Figure 5.3 is actually a sample of points in that block thus giving us a fairly good representation of the machine, and since most of the materials stay the same and we average out the hotspots out with the cold spots before using any of the data it takes care of small temperature reading issues. The last concern may hold the truest of all the concerns, the way we created a recirculation variable k which without doing much research on how that variable must be used. In our experiment we pick a very small number for k because our setup had alot of room so the outlet temperature would dissipate very easily so a improvement would be to aggregate all the servers and layouts and update the k variable in real time.

7.2 Extension

The next step would be to run these similar results in a full blown data center, or at least one with AC unit and multiple servers. Also for the future research we would like to use more sophisticated pieces of hardware for testing temperatures. The follow up research should have a flowing air thermometer for the inlet and outlet temperatures, some kind of heat measurement to keep from extrapolating heat from surface temperatures, and lastly mounts for our sensors and programmatic way to gather these multiple pieces of data without manually scanning them to get more accurate results. Future research should look at neighboring nodes and heat recirculation more closely to see how they will affect the output in a more exact fashion. Equation 5 is a possible enhancement to the model that takes in account of the neighboring nodes.

Another place for future research could be to look at the model as function of time. This could be useful because some application may run short while other run long. Since this model was made to help facilitate quick modeling of thermal pattern of data centers, a model that will run in real time could help in develop even more robust algorithms. As a quick preview equation 6 explains what a time based model could look like.

Chapter 8

Conclusion

Growing evidence show that cooling costs contribute a significant portion of the operational cost of large-scale data centers. Therefore, thermal management techniques aim to reduce energy consumption and cooling costs of data centers. A thermal management mechanism relies on thermal information to make intelligent decisions. Thermal information can be acquired in three ways: 1) using "Computational Fluid Dynamics" simulators (e.g., Flovent is a commercial product), 2) deploying temperature sensors to measure inlet and outlet temperatures of servers, and 3) applying a thermal model to estimate temperatures based on specific workloads.

We advocate that thermal models are a cost-effective and practical approach to providing information on server temperatures to thermal management mechanisms. In this study, we develop a thermal model - iTad - that enables thermal management techniques to quickly make management decisions based on intensive I/O activities. We show that in light of iTad, both CPU and I/O thermal outputs can be extrapolated from radiation heat and convection heat applied to a server. The iTad model helps in improving the energy efficiency of data centers, because thermal information offered by iTad assists dynamic thermal management to reduce the energy consumption in cooling systems in data centers. We validate the accuracy of the iTad model using a server's real-world temperature measurements obtained by an infrared thermometer. Our experimental results suggest that I/O-intensive workloads have significant impacts on temperatures of servers. We demonstrate that thermal management mechanisms can quickly make workload placement decisions based on thermal information facilitated by iTad.

Bibliography

[1] Convective and radiant heat transfer equations. `http://blowers.chee.arizona.edu/cooking/heat/convection.html`.

[2] Flovent. `http://www.mentor.com/products/mechanical/products/upload/flovent-data-center.pdf`.

[3] hddtemp. `http://manpages.ubuntu.com/manpages/natty/man8/hddtemp.8.html`.

[4] iostat. `http://sourceforge.net/apps/mediawiki/filebench/index.php?title=Main_Page/`.

[5] iostat. `http://linux.die.net/man/1/iostat`.

[6] stress. `http://www.unixref.com/manPages/stress.html`.

[7] Amazon. http://www.amazon.com/pyle-pma90-anemometer-thermometer-temperature/dp/b009tq6ilq. *Amazon*, 10 March 2012.

[8] Hendrik F. Hamann Andreas Bieswanger and Hans-Dieter Wehle. Energy efficient data center. Technical Report 1, 2012.

[9] Amanda J Barra and Janet L Ellzey. Heat recirculation and heat transfer in porous burners. *Combustion and Flame*, 137(1):230 – 241, 2004.

[10] Luiz André Barroso and Urs Hölzle. The case for energy-proportional computing. *Computer*, 40(12):33–37, December 2007.

[11] David J. Brown and Charles Reams. Toward energy-efficient computing. *Commun. ACM*, 53(3):50–58, March 2010.

[12] Dimitris Economou, Suzanne Rivoire, and Christos Kozyrakis. Full-system power analysis and modeling for server environments. In *In Workshop on Modeling Benchmarking and Simulation (MOBS*, 2006.

[13] P.A. Eibeck and D.J. Cohen. Modeling thermal characteristics of a fixed disk drive. *Components, Hybrids, and Manufacturing Technology, IEEE Transactions on*, 11(4):566 –570, dec 1988.

[14] Sudhanva Gurumurthi, Anand Sivasubramaniam, and Vivek K. Natarajan. Disk drive roadmap from the thermal perspective: A case for dynamic thermal management. *SIGARCH Comput. Archit. News*, 33(2):38–49, May 2005.

[15] Urs Hoelzle and Luiz Andre Barroso. *The Datacenter As a Computer: An Introduction to the Design of Warehouse-Scale Machines.* Morgan and Claypool Publishers, 1st edition, 2009.

[16] Youngjae Kim, S. Gurumurthi, and A. Sivasubramaniam. Understanding the performance-temperature interactions in disk i/o of server workloads. In *High-Performance Computer Architecture, 2006. The Twelfth International Symposium on*, pages 176 –186, feb. 2006.

[17] Jonathan G. Koomey. Estimating total power consumption by servers in the U.S. and the world. Technical report, Lawrence Derkley National Laboratory, February 2007.

[18] Sanjay Kulkarni. *Cooling Hadoop: Temperature Aware Schedulers in Data Centers.* PhD thesis, Auburn University, 2013.

[19] J Li, G.P Peterson, and P Cheng. Three-dimensional analysis of heat transfer in a micro-heat sink with single phase flow. *International Journal of Heat and Mass Transfer*, 47(190):4215 – 4231, 2004.

[20] Lei Li, Chieh-Jan Mike Liang, Jie Liu, Suman Nath, Andreas Terzis, and Christos Faloutsos. Thermocast: a cyber-physical forecasting model for datacenters. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '11, pages 1370–1378, New York, NY, USA, 2011. ACM.

[21] J. Moore, J. Chase, and P. Ranganathan. Consil: Low-cost thermal mapping of data centers. In *the First Workshop on Tackling Computer Systems Problems with Machine Learning (SysML)*, 2006.

[22] R. Sharma, C. Bash, C. Patel, R. Friedrich, and J. Chase. Balance of power: Dynamic thermal management for internet data centers. *IEEE Internet Computing*, 9(1), 2005.

[23] C.P.H. Tan, J.P. Yang, J.Q. Mou, and E.H. Ong. Three dimensional finite element model for transient temperature prediction in hard disk drive. In *Magnetic Recording Conference, 2009. APMRC '09. Asia-Pacific*, pages 1 –2, jan. 2009.

[24] Q. Tang, S. Gupta, and G. Varsamopoulos. Thermal-aware task scheduling for data centers through minimizing heat recirculation. In *Cluster Computing, 2007 IEEE International Conference on*, pages 129 –138, sept. 2007.

[25] Q. Tang, S. Gupta, and G. Varsamopoulos. Thermal-aware task scheduling for data centers through minimizing heat recirculation. In *Cluster Computing, 2007 IEEE International Conference on*, pages 129 –138, sept. 2007.

[26] Qinghui Tang, S.K.S. Gupta, D. Stanzione, and P. Cayton. Thermal-aware task scheduling to minimize energy usage of blade server based datacenters. In *Dependable, Autonomic and Secure Computing, 2nd IEEE International Symposium on*, pages 195 –202, 29 2006-oct. 1 2006.

[27] Qinghui Tang, S.K.S. Gupta, and G. Varsamopoulos. Energy-efficient thermal-aware task scheduling for homogeneous high-performance computing data centers: A cyber-physical approach. *Parallel and Distributed Systems, IEEE Transactions on*, 19(11):1458 –1472, nov. 2008.

[28] W. Wechsatol, S. Lorente, and A. Bejan. Dendritic heat convection on a disc. *International Journal of Heat and Mass Transfer*, 46(23):4381 – 4391, 2003.