PROACTIVE THERMAL-AWARE MANAGEMENT IN CLOUD DATACENTERS

BY EUN KYUNG LEE

A dissertation submitted to the

Graduate School—New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Electrical and Computer Engineering

Written under the direction of

Professor Dario Pompili

and approved by

_____

_____

_____

_____

New Brunswick, New Jersey

January, 2015

ABSTRACT OF THE DISSERTATION

Proactive Thermal-aware Management in Cloud Datacenters

by Eun Kyung Lee

Dissertation Director: Professor Dario Pompili

The complexity of modern datacenters is growing at an alarming rate due to the rising popularity of the cloud-computing paradigm as an effective means to cater to the ever increasing demand for computing and storage. The management of modern datacenters is rapidly exceeding human ability, making autonomic approaches essential. In the meanwhile, the increasing demand for faster computing and high storage capacity has resulted in an increase in energy consumption and heat generation in datacenters. Due to the increased heat generation, cooling requirements have become a critical concern, both in terms of growing operating costs as well as their environmental and societal impacts. (e.g., increase in CO_2 emissions, overloading the electric supply grid resulting in power cuts, heavy water usage for cooling systems causing water scarcity)

In this thesis, proactive thermal-aware datacenter management solutions, which include thermal- and energy-aware resource provisioning, cooling system optimization, and anomaly detection, are proposed to help minimize both the impact on the environment and the Total Cost of Ownership (TCO) of datacenters, making them energy efficient and green. For the proactive thermal-aware solutions, a novel architecture endowed with different abstract components is introduced, which is composed of four

layers: the environment layer (which detects, localizes, characterizes, and tracks thermal hotspots), the physical-resource layer (which manages the hardware and software components of servers), the virtualization layer (which instantiates, configures, and manages VMs), and the application layer (which is aware of the workload's and applications' characteristics and behavior).

Our solutions autonomically manage datacenters using cross-layer information collected from the four-layered architecture and make decisions based on various application-specific optimization goals (e.g., performance, energy efficiency, anomaly detection rate). A sensing infrastructure to measure the datacenter's environmental change and methods to acquire thermal awareness (using real-time measurements and heat- and air-circulation models) are discussed. Then, specific proactive thermal-, energy-, and anomaly-aware solutions are proposed, which i) optimize cooling systems (i.e., air conditioner compressor duty cycle and fan speed) to prevent heat imbalance and minimize the cost of cooling, ii) maximize computing resource utilization to minimize datacenter energy consumption, and iii) differentiate servers' thermal map (temperature) frequently to maximize the thermal anomaly detection rate.

Acknowledgements

I have been dreaming about the moment of writing dissertation and finally getting Ph.D. degree, but I realize that I could have never gotten here alone without the support from a number of important people. My work would not be complete without the supervision of my advisor Professor Pompili. His infectious enthusiasm for research and perception into problems have taught me how to conduct research and be an influential researcher. He has motivated and pushed me with good discipline during numerous discussions and meetings, always providing research-oriented guidance. He has helped me to understand the challenges of research in this fast-paced world, and to come up with novel ideas to tackle such challenges. Moreover, he has been a good friend whenever I struggled with personal issues.

I would also like to thank my committee members Dr. Manish Parashar, Dr. Ivan Rodero, and Dr. Renato Figueiredo for their support over the years. They have first inspired me, provided innovative ideas on different approaches for datacenter management, and led me to the issues on the virtualized datacenters and solutions for green computing. Their professionalism have encouraged me to complete my Ph.D. studies. Also, I would like to take this opportunity to thank to all the professors at Rutgers and Cloud and Autonomic Computing Center (CAC), who have offered stimulating ideas and courses on communications, computer engineering, networking, programming, and mathematics. I have been lucky to have many supportive friends, and I especially thank to Hariharasudhan Viswanathan for rescuing me whenever possible.

Finally, I sincerely thank my parents for urging me to excel in every sphere of life with warm heart and my brother for being a good friend. And the last goes to my wife and daughter, who continue to support me with their big smiles.

Dedication

To my wife Rebecca Choi

my daughter Elise Yena Lee

my mother Mrs. Kum Ryun Kim

my father Mr. Kun Chang Yi

and my brother Mr. Yangkyoung Yi

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

Big data analytics is a key enabler of the paradigm shift towards online data-driven decision making in the fields of health, banking, commerce, defense, research and education, and entertainment. The increase in demand for fast computing and large storage to handle big data – large *volumes* of raw data aggregated at high *velocities* from a *variety* (i.e., heterogeneous) of sources – has led to a tremendous increase in the number and size of cloud datacenters (in terms of number of servers). Datacenters have become key components of society's IT infrastructure; therefore, their *energy consumption*, *heat generation*, and *cooling requirements* have become critical concerns both in terms of the growing operating costs as well as their environmental and societal impacts [4].

1.1 Motivations and Objectives

Datacenter energy consumption surpassed 237 billion kWh/year worldwide and 76 billion kWh/year in the US in 2010 [1], which correspond to 1.3% and 2% of the total electricity usage worldwide and in the US, respectively. The impact of the proliferation of datacenters on the environment and society includes increase in CO_2 emissions, overload of the electricity supply grid, and rise in water usage for cooling leading to water scarcity [5], all of which have initiated studies about green-computing practices, especially, energy and power management [6, 7]. The scale and complexity of datacenters are growing at an alarming rate and their management is rapidly exceeding human ability, making *autonomic* (self-configuration, self-optimization, self-healing, and self-protection) management approaches essential.

Figure 1.1 shows the average share of the different components of datacenter-energy consumption based on independent studies from different agencies [8] in 2006 and shows

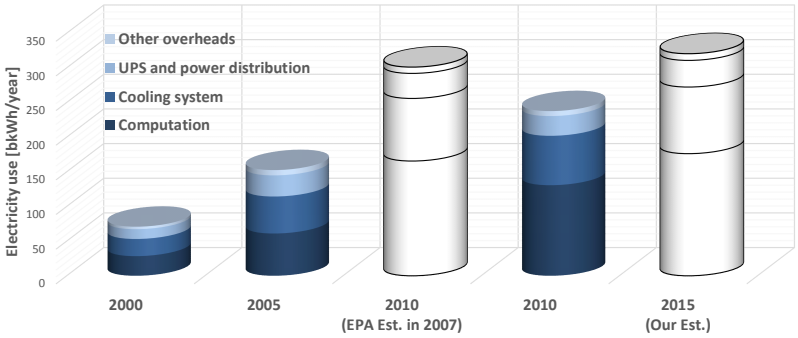

Figure 1.1: Average share of the different components of datacenter energy consumption based on independent studies [1, 2].

that large amount (40%) of the energy is consumed by the cooling infrastructure. Even though the best-case Power Usage Efficiency (PUE, i.e., the ratio of total power consumption to the power usage for computing) improved owing to green-computing efforts with advances in cooling technologies and operational changes (e.g., innovations at Facebook and Google datacenters), majority of old datacenters are not equipped yet with the latest technologies to decrease their total energy consumption [2].

Recent advances in optimization of computing resource and cooling system utilization have strived to increase energy and cooling efficiency while decreasing equipment failure rates so to minimize both the impact on the environment and the Total Cost of Ownership (TCO) of datacenters. However, the solutions for optimizing computing resource utilization, e.g., workload consolidation, Dynamic Voltage and Frequency Scaling (DVFS), pinning (CPU affinity) [9–11], that are oblivious of their *thermal* implications (i.e., heat generation) may adversely affect the cooling system efficiency, the lifetime of computing equipment due to overheating, and hence the TCO of datacenters. Research on improving cooling system efficiency (i.e., heat extraction) has focused on determining the optimal compressor duty cycle and fan speeds of CRAC units [12, 13] and on designing datacenter layouts that minimize recirculation of heat generated by computing equipment inside the datacenter [14, 15]. However, cooling system optimization solutions provide benefits only when they are aware of the spatial distribution of heat generation, which is usually not the case.

From our feasibility study and proof-of-concept experiments conducted at our machine room in the NSF Cloud and Autonomic Computing Center (CAC), we have inferred that one of the fundamental problems in cloud datacenters is the *local unevenness in heat-generation* and *heat-extraction rates*: the former can be attributed to the non-uniform distribution of workloads (of different types and intensities) among servers and to the heterogeneity of computing hardware; the latter can be attributed to the non-ideal air circulation, which depends on the layout of server racks inside the datacenter and on the placement of Computer Room Air Conditioning (CRAC) unit fans and air vents.

The heat-generation and -extraction rates may differ, which over time causes what we call *heat imbalance*. This heat imbalance will be large if the rates are significantly different from each other or if their difference prolongs over extended time periods. A large negative heat imbalance at a particular region inside a datacenter will result in energy-inefficient overcooling and, hence, a significant decrease in temperature. Conversely, a large positive heat imbalance at a particular region will create a *thermal hotspot*. Thermal hotspots may also result in a thermal fugue characterized by a continuous increase in the rate of temperature rise, which in turn leads to server operation in the unsafe temperature range [16]. Thermal awareness, which is the knowledge of heat imbalance at different regions inside a datacenter, is essential to maximize energy and cooling efficiency as well as to minimize server system failure rates. In this dissertation, we study the research efforts on thermal-aware management of datacenters.

Moreover, due to their growing importance, datacenters are strategic targets [17] for *denial-of-service attacks* (running illegitimate workloads) and *cooling-system attacks* aimed at causing thermal runaways and, hence, costly outages, which can potentially cripple society's critical IT infrastructure. Furthermore, due to their large scale and high server density, the probability of computing- and cooling-system misconfigurations as well as of cooling-equipment and server-fan failures is high [18]. Such unpredictable events may result in unexpected high temperature areas/regions (hotspots) or excessively cooled low temperature areas/regions (coldspots), also referred to as *thermal anomalies*.

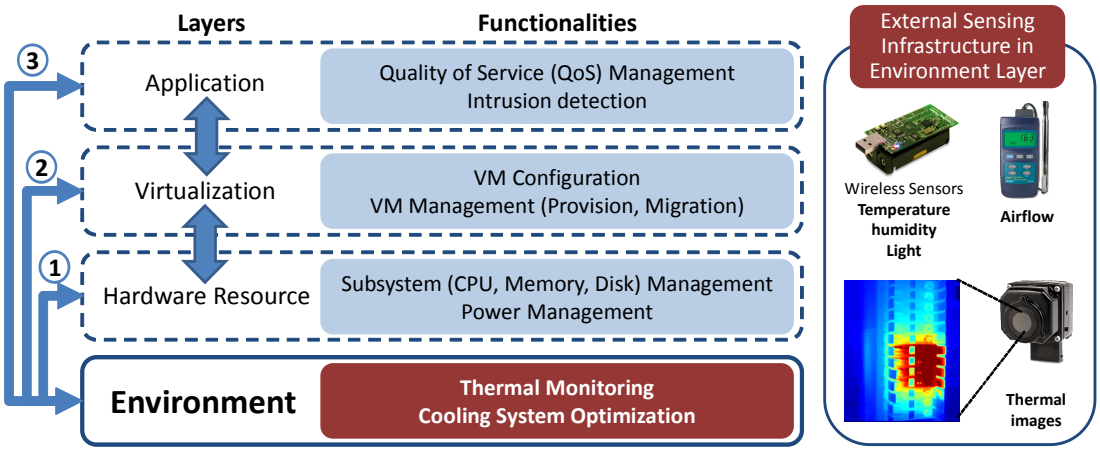

Figure 1.2: Our proposed environment layer and layered architecture mode: ①, ②, and ③ show the interactions between hardware-resource and environment layer, virtualization and environment layer, and application and environment layer, respectively.

1.2 Contributions

The contribution of the dissertation is to introduce an environment layer in multi-layered architecture, and proactive thermal-aware resource management techniques.

1.2.1 Environment Layer in Multi-Layered Architecture

The long-term goal of our approach is to autonomically manage datacenters using the information from sensors and taking decisions at different levels (through controllers) based on the optimization goals (e.g., performance, energy efficiency, cost). To do this, we consider an architecture composed of layers belonging to different abstract components with different responsibilities but with the same common objectives. The layered approach to datacenter management was first proposed in [19] to absorb the wide range of emergent datacenter implementations. Even though there is no consensus in the community regarding the 'names' of the different layers, there is one as far as the set of functionalities they should encapsulate. For example, the *physical-resource layer* referred to in this article is similar to the physical-infrastructure layer introduced in [19]. In this article, we choose names that best convey the functionalities encapsulated by a layer. Typically, the layered architecture is composed of three layers: physical (hardware) resource layer, virtualization layer, and application layer, as shown in Fig. 1.2.

The interactions and dependencies between the layers are systematically specified.

In addition to the existing layered architecture, we introduce a new layer in the datacenter management architecture, the so-called *environment layer*, whose functionalities aim at detecting, localizing, characterizing, and tracking thermal hotspots using a external sensing infrastructure composed of in-built as well as external scalar and multimedia sensors. Thus, our proposed architecture (see Figure 1.2) is composed of four layers: environment layer (which detects, localizes, characterizes, and tracks thermal hotspots using a hybrid sensing infrastructure composed of in-built as well as external scalar temperature and humidity sensors), physical resource layer (which manages the hardware and software components of servers), virtualization layer (which instantiates, configures, and manages VMs), and application layer (which is aware of the workload's and applications' characteristics and behavior).

Heat imbalance at different regions inside a datacenter, inferred from the data gathered and processed at the environment layer, can be employed to improve the overall operational efficiency of datacenters. This is achieved through interactions with the other layers rendering the different functionalities thermal aware. Such thermal awareness, which is acquired using the environment-layer-specific functionalities, can be leveraged for 1) reduction of the TCO of datacenters as well as their environmental and societal impacts, and 2) early detection of thermal anomalies inside datacenters so to improve the availability, security, and lifetime of computing systems. However, maximizing energy efficiency and utilization of Cloud-datacenter resources, avoiding undesired thermal hotspots, and ensuring QoS are *all conflicting objectives* that require careful consideration of multiple pairwise trade-offs.

For autonomic thermal management we focus on the possible interactions among all the four layers (*cross-layer*) based on temperature, power consumption, and application QoS requirements. Note that the environment layer comprising of the hybrid sensing infrastructure monitors both the micro (chip and server level) as well as the macro (rack and datacenter level) phenomena, i.e., temperature distribution, in the datacenter. The proposed cross-layer approach represents a significant and transformative shift towards cross-layer autonomics for datacenter management problems, which have so far

been considered mostly in terms of individual layers. The environment layer provides information of thermal behavior of datacenters, and then, the controllers in each layer jointly take decision to minimize the energy consumption while alleviating thermal anomalies. In this dissertation, the cross-layer solutions are characterized by interaction between two layers - physical and virtualization layers, physical and environment layers, application and virtualization layers, or application and environment layers.

1.2.2 Proactive Thermal-aware Resource Management Techniques

The proactive approach is also proposed as the reactive approach causes delay in propagating remedial action in thermal domain. Since the reactive approach directly observes the temperature at the place where we want to control the temperature, depending on the size of the datacenters and placement of the temperature sensors this approach could incur significant delay for cooling. The larger the datacenter size the more the cooling delay incurs and the farther the temperature sensor from the controller the more the cooling delay incurs. If the delay increases, the temperature of any blade rises up to the critical point and damage the hardware components before the air conditioning system reacts and cools down the computing system. Also, the cooling cost could be more expensive if only a few overheated machines near the sensor can trigger an entire cooling system.

In proactive approach, the actions (i.e., adjusting fan speed, adjusting air compressor cycle, placing workloads) are jointly optimized to achieve different objectives (i.e., minimizing energy, minimizing thermal anomalies) before taking any action. As the different actions in different layers are optimized ahead, delay does not incur, and hence, minimize the risk of damaging hardware components in datacenters. Our proactive solution employs *heat-imbalance model* to estimates the heat that will be generated in the future and take actions. This way, we can achieve different optimization objectives while preventing 'risk of overheating' and 'excessive cooling'.

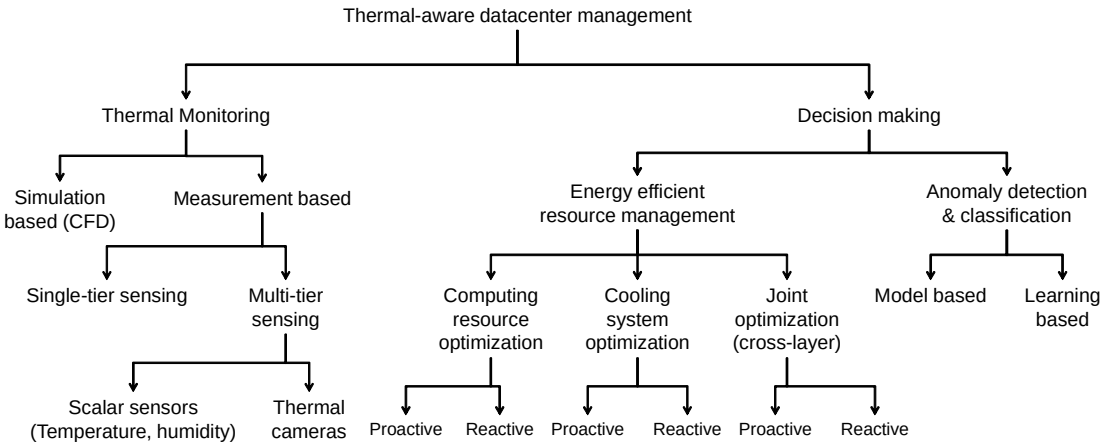

Figure 1.3: Components of thermal-aware datacenter management.

1.3 Scope of Dissertation

The scope of the dissertation is shown in Fig. 1.3. Two categories of thermal-aware datacenter management are mainly discussed in this dissertation: thermal monitoring and thermal-aware decision making.

Thermal Monitoring: Monitoring is essential to characterize the thermal behavior of a datacenter. The *thermal signature*, which is the expected distribution (locations and intensities) of thermal hotspots for a specific workload distribution among servers and a specific cooling system setting. The thermal signature is exploited to make any control (i.e., energy-aware, anomaly-aware) decisions. Most thermal-aware datacenter management solutions assume inputs from a single-tier sensing infrastructure (e.g., temperature and humidity scalar sensors, and air flow meters). However, this extracted information from the single-tier sensing infrastructure cannot capture the complex thermodynamic phenomena of heat and air circulation inside a datacenter.

Thus, we propose a multi-tier sensing infrastruecure – composed of thermal cameras, scalar temperature and humidity sensors, and airflow meters – into a multi-tier sensing infrastructure that monitors the complex thermodynamic phenomena inside a datacenter. The sensing infrastructure exploits the spatio-temporal correlation in the observed phenomena and the temperature distribution map provided by thermal images, and decides on-the-fly the granularity at which it should sample over space and

time (*adaptive sampling* [20]), thus enabling efficient real-time monitoring of datacenters. Incorporation of real-time measurements (such as temperature, air-flow rate, etc.) into heat generation and extraction rate models helps estimate the heat imbalance, which allows prediction of the future temperature map of the datacenter. Also, it helps the accurate understanding of thermal signature, which is crucial for designing efficient thermal-management solutions

Monitoring could be simulation based, e.g., using Computational Fluid Dynamics (CFD), or measurement based, e.g., temperature and humidity sensors, airflow meters, thermal cameras, etc. Furthermore, datacenter managers may choose to employ either single-tier sensing or multi-tier sensing strategy depending on their needs. The dissertation advocate multi-tier sensing infrastructure to measure the heat transfer in datacenters.

Decision making (Energy efficient resource management): Thermal Management decision can be made base on the information provided from the monitoring phase. Majority of prior research efforts on thermal management of datacenters has focused exclusively on only one of the two fundamental approaches: management of heat generation (computing resource optimization) or management of heat extraction (cooling system optimization) inside a datacenter. The first approach focuses on optimization of computing resource utilization and how to balance or migrate workloads in such a way as to avoid overheating of computing equipment while minimizing energy consumption. The second approach aims at improving cooling system efficiency by effectively distributing cold air inside the datacenter.

However, we propose a joint approach so to maximize resource utilization and cooling efficiency while simultaneously minimizing the risk of overheating of servers. Thermal-aware decision making can either be reactive or proactive in nature. Reactive management solutions involve remedial actions in response to undesired thermal behavior deduced from measurements. In contrast, our proposed proactive management solutions use measurements and predictive models so to avoid undesired thermal behavior.

Decision making (Anomaly detection & classification): Ensuring datecenter

equipment stays up and operating is essential for the seamless service in datacenters. As the service failure can incur huge losses in productivity, business, and opportunity for both service providers and consumers, datacenter manager's main goal is to prevent failures and provide reliable service. Unexpected changes in the local heat generation and extraction rates due to cooling equipment failures, misconfigurations, and attacks (such as illegitimate workloads) may over time cause large heat imbalances resulting in *unexpected thermal hotspots.* Thermal anomalies such as unexpected hotspots and fugues lead to system operation in unsafe temperature regions, which will increase the server failure rate and the TCO of datacenters. Furthermore, the probability of occurrence of thermal anomalies in modern datacenters is high because of their large scale and high server density. Monitoring and thermal-aware decision making enable easier detection and localization of anomalous events manifesting as unexpected thermal hotspots.

1.4 Dissertation Organization

In Chapter 2, we discuss the preliminaries to motivate proactive thermal-aware datacenter management approaches. We introduce changes of datacenter design to explain the importance of cooling system and need for thermal management, and the state of the art datacenter monitoring (sensing) and decision making techniques as well as cross-layer techniques in thermal-aware datacenter management. Then, we take a closer look at two example cross-layer management solutions. In Chapter 3, we propose a proactive control approach that jointly optimizes cooling systems (i.e., air conditioner compressor duty cycle, and fan speed) to prevent heat imbalance - the difference between the heat generated and extracted from a machine - thus minimizing the cost of cooling. In Chapter 4, we propose VMAP, an innovative proactive thermal-aware virtual machine consolidation technique to maximize computing resource utilization to minimize datacenter energy consumption for computing, and to improve the efficiency of heat extraction. Finally, in Chapter 5, we conclude the dissertation and summarize the open research problems in datacenter monitoring, in autonomic thermal-aware datacenter management, and in thermal anomaly detection and classification in datacenters.

Chapter 2

Preliminaries

In this chapter, preliminaries to motivate proposed proactive and thermal-aware datacenter management approaches are discussed. In sect. 2.1, the evolution of datacenter design over the last two decades is retraced to explain the need for thermal management. In sect. 2.2, the state of the art datacenter monitoring (sensing) techniques and a novel multi-tier sensing infrastructure (which enable thermal-aware management) are introduced. In sect, 2.3, the state of the art cross-layer techniques in thermal-aware datacenter management are discussed.

2.1 Evolution of Datacenter Design

Datacenter design and management has evolved over the years to enable efficient thermal management (extracting heat at a low cost from computing resources) for minimizing server failure rates due to overheating, for minimizing energy expenditure, and for minimizing CO_2 footprint. Innovations in datacenter design and management can be roughly categorized as belonging to one of the two eras in which they were made. The first is the "dot-com era" [21] (late 1990s up to mid 2000s) when a number of small datacenters appeared to support the first online services of dot-com (e-commerce-driven) companies. The main focus of datacenter designers and managers during this period was reliability (the goal usually was to ensure 99.99% availability). The second is the "post dot-com era" (mid-2000s onwards) characterized by the growing popularity of online search, social networking, online gaming, and a host of cloud service offerings (from storage to high-performance computing support). Due to the advent of warehouse-scale datacenters [22] during this era, minimizing the TCO and the environmental impact of datacenters became critical concerns in datacenter design and management apart from

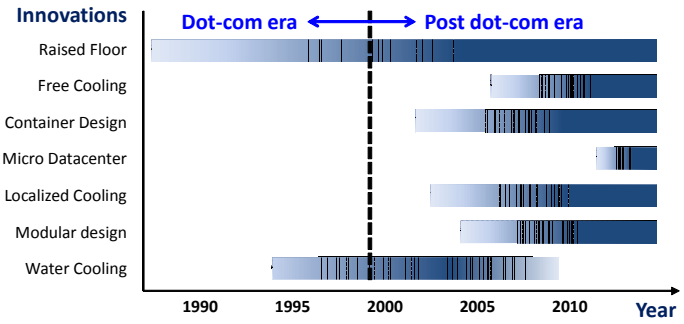

Figure 2.1: Different design choices of of datacenters in different scopes.

reliability.

Figure 2.1 shows the different innovations and the year (hence the era) in which they were made. To show the major shift in the design focus from "reliability" to "energy savings and environmental awareness" we separate the period of last two decades into "dot-com era" and "post dot-com era". During the dot-com era, the innovations were geared towards lowering the room temperature effectively to prevent overheating of computing resources. As the majority of datacenters is air cooled, cooling system management techniques focused on optimizing air flow inside the datacenter by preventing the mixing of cold air from the CRAC and the hot air from the servers [23]. Innovations in the post dot-com era include modular design (to ease management and scale out) and free cooling (for energy efficiency). Innovations made during the dot-com era are still relevant because modern warehouse-scale datacenters still incorporate some of the innovations (such as raised floor design) and also because as many as 80% of small- to medium-sized datacenters in the US were built in the late 1990s and early 2000s [22].

Figure 2.2 shows evolving researches trends on the thermal management. We used keywords of 'thermal management, green computing, cooling optimization, and computing optimization' using Google Scholar (http://scholar.google.com) and count indexes to show the increasing academic interests on thermal management in datacenters.

2.2 Thermal Monitoring

There have been two approaches for monitoring datacenters; one is simulation-based and the other one is measuring-based approach. The former focuses on how to model

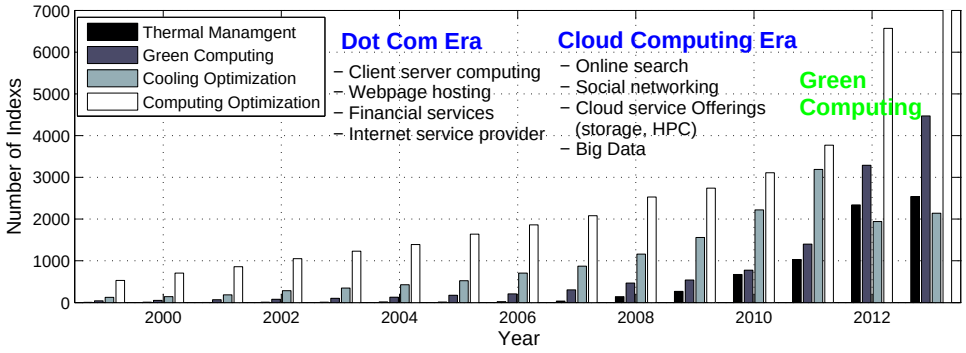

Figure 2.2: Number of indexes of keywords through Google Scholar (http://scholar.google.com); Thermal Management (keywords: datacenter thermal management), Green Computing (keywords: datacenter green computing), Cooling Optimization (keywords: datacenter cooling optimization), and Computing Optimization (keywords: datacenter computing optimization).

and simulate the thermal behavior of datacenter while the later focuses on how to effectively collect data using sensors. The goal of both approaches is to profile thermal behavior (i.e.,temperature, humidity, etc.), and hence, to minimize the cost for cooling or to detect anomalies in datacenters.

2.2.1 Simulation-based approach

Most of the simulation-based approaches use CFD that thermally models 3-D temperature map in datacenters. CFD predicts the condition of the air (i.e., temperature, airflow, and pressure) in data center to assess performance and energy consumption, using numerical modeling [24] [25].It has been primarily used for analyzing the effectiveness of cooling system (i.e., CRAC, fan, and ventilation system) in datacenters. CFD provides 3-D map of cold air moving through a data center and it identifies potential "hot spots" where the cold air does not reach enough. Also, it detects "over-cooling spots" where the area receives more cold air than needed, resulting in wasting energy. When datacenter manager needs to find potential problem (hotspots) ahead of time, the CFD tool could be very effective.

However, the simulation-based approach is computationally expensive. This approach incurs time delay to solve complex differential equations to generate the thermal map. It is not suitable for the online cooling solutions, which manage cooling system on

the fly depending on condition of datacenters. Moreover, input parameters of the simulation (i.e., workload, temperature, power level) are highly time-varying and location-varying, the simulation gives the results frequently to catch up the variance. Otherwise, the temperature of the datacenter can readily go up to high temperature if the cooling system is not properly activated [26]. Also, simulation parameters should be tuned to give more accurate simulation, which gives more overheads. Hence, simulation-based approach is not appropriate for the real-time cooling solution, which requires immediate remedial action.

Simulation-based approach is especially useful in deciding an energy-efficient datacenter design among different design choices. As different building has different limitations (e.g., different cooling system, limited space, insulation), deciding one solution that works well for other datacenters may not be the best solution. Also, complex thermodynamic makes hard to decide the design (e.g., location of servers and cooling system, location of the vent, cooling power). In these cases, simulation-based approach works very well. Commercial CFD solutions are available at 6Sigma from Future Facilities [27], FloVENT from Mentor Graphics [28], and TileFlow from Innovative Research Inc. [29].

In [30], As server power and density continue to increase, efficient cooling has become key to controlling costs. Predictive modeling based on computational fluid dynamics enables enterprises to configure their data centers for optimal cooling, helping maximize efficiency, reduce costs, and meet both current and future IT requirements. (CFD-ENV Layer) Find best design. and

In [31], This paper provides an overview of some of the ongoing work to operate within the thermal environment of a data center. Some of the factors that affect the environmental conditions of data-communication (datacom) equipment within a data center are described. Since high-density racks clustered within a data center are of most concern, measurements are presented along with the conditions necessary to meet the datacom equipment environmental requirements. A number of numerical modeling experiments have been performed in order to describe the governing thermo-fluid mechanisms, and an attempt is made to quantify these processes through performance

metrics.

In [32], This paper focuses on the effect on rack inlet air temperatures as a result of maldistribution of airflows exiting the perforated tiles located adjacent to the fronts of the racks. The flow distribution exiting the perforated tiles was generated from a computational fluid dynamics (CFD) tool called Tileflow (trademark of Innovative Research, Inc.). Both raised floor heights and perforated tile-free areas were varied in order to explore the effect on rack inlet temperatures. The flow distribution exiting the perforated tiles was used as boundary conditions to the above-floor CFD model. A CFD model was generated for the room with electronic equipment installed on a raised floor. (CFD-ENV Layer)

[33] thermally profiled a non-raised-floor data center in a small office and then compared the results to a Computation Fluid Dynamics (CFD) model of the space. The results compared favorably. Data from other electronic equipment rooms are available, but it is very difficult to glean important information or correlate or compare data. [24, 25] developed a three-dimensional model of a laboratory data center and experimentally verified the numerical results to ensure the specified inlet air temperatures to the computer systems met the temperature limits.

In [34], the authors emulate the thermal behavior of server systems to manage datacenter cooling. Using the emulator, a system named Freon monitors temperature changes and, if the temperature of a machine crosses a threshold (defined thermal emergency), Freon redistributes the jobs; however, cooling cost is not considered in this study.

Those simulation approaches produce thermal distribution of the datacenter in space, and does not require any physical sensors or systems to process measurements. Several commercial simulators are available currently (e.g., ANSYS HPC [35]). However, the cost for these simulators are expensive and it takes many ours to days to simulate large quantitative systems, so that cooling system cannot get a chance to react depending the profiled thermal distribution. Also, those thermal distributions from simulators are not workload dependent. Since the heat generated from the servers are heavily dependent on location and the types of workloads (i.e., CPU intensive, memory

intensive, disk intensive, etc.), the features of different workloads should be considered when running simulations. However, it is hard for the simulators to factorize all the complexities of different types of workloads. They typically simulate only the temperature based on steady state power consumption, not considering workload information.

In [36], we propose an abstract heat flow model which uses temperature information from onboard and ambient sensors, characterizes hot air recirculation based on these information, and accelerates the thermal evaluation process for high performance datacenters. This is critical to minimize energy costs, optimize computing resources, and maximize computation capability of the datacenters. Given a workload and thermal profile, obtained from various distributed sensors, we predict the resulting temperature distribution in a fast and accurate manner taking into account the recirculation characterization of a datacenter topology. Simulation results confirm our hypothesis that heat recirculation can be characterized as cross interference in our abstract heat flow model. Moreover, fast thermal evaluation based on cross interference can be used in online thermal management to predict temperature distribution in real-time. In [32], the air flow was modeled using the k-e turbulence model to provide some guidance on the design and layout of a data center.

2.2.2 Measuring based approach

The source of the cool air in traditional data centers is only part of the problem. Companies have also begun demonstrating the importance of managing circulation within the space. When Synapsense audits a facility, its technicians install wireless sensors throughout the building to measure temperature, pressure, humidity, and more. Pandey says Synapsense often identifies intense hot spots warmer areas that force the fans and mechanical chillers to work harder to manipulate the temperature, thus increasing energy usage. "You might have enough cool air, but it's not going to the right places," he explains. "There might be mixing of the air or there might be areas where it's leaking."

In one case, Synapsense installed 3,674 sensors throughout a 100,000-square-foot data center. The sensors fed Synapsense's control system a stream of data on temperature, pressure, and humidity, and the company's software built a live-updated map of

these metrics throughout the facility. With this data, Synapsense was able to figure out how to optimize energy use by turning up certain fans or shutting down specific air conditioning units. It ended up saving the company 8,244 megawatt hours per year $766,000 in annual electrical bills.

As a multi-point measurement method, Microsoft research group suggests Wireless Sensor Networks (WSNs) platform to monitor data center environment. In [37], authors present a wireless platform named Microsoft Research Genomotes designed for data centers. Genomotes are composed of wireless and wired hardware devices which embeds temperature and humidity sensors inside. In [38], a company (MAXIM) propose 1-Wire system which can route all sensor readings over a single wire. However, the information from those scalar sensors is not sufficient to capture "heat" because it is too sparsely deployed to capture the heat propagation. The authors stress that 1-Wire system saves more energy so reduces cost than adding complex wiring overhead. Building wireless sensor networks costs more/less than 10$ per node, but 1-Wire system keeps costs to couple of dollars per nodes.

There are two types of Genomotes, masters and slaves. A master is a wireless node with sensing capability, having a RS232 port connected to a slave node. A slave only has a sensing capability without wireless device, wired to the master forming daisy chain network. They present that the daisy chain formation is ideal for installing on server racks because this model can measure the temperature of the various heights of the rack without tangling wires as well as reduce wireless nodes which create bandwidth saturation.

The same group suggests RACNet [13] defined as a 'large-scale sensor network for high-fidelity data center environmental monitoring' which is the first attempt to monitor cooling behavior of data center. In this paper, they present Reliable Data Collection Protocol (rDCP) that dynamically constructs spanning trees to reliably route sensing measurement from the racks in a data center. Their empirical results show that rDCP achieves over 99% data yield for over 95% of the sensors out of 174-Genomotes deployments.

Based on the measurement from scalar sensors, there are efforts to thermally profile

data centers. Basic mathematical modeling and parameters for profiling datacenter are proposed in [39]. Since many data centers employ raised floors with perforated tiles to distribute the chilled air to racks, providing proper distribution of chilled air is one of the challenges. In [15, 31, 40, 41], the authors provide some insights into the airflow distribution from perforated tiles and raised floor design in datacenters. In [14, 41], the authors thermally profiled a datacenter in space and time, and analyzed trends and correlations among collected measurements. In [31], the authors comment on the challenges associated with thermal management in datacenters. In [42], the authors review the existing literature on datacenter thermal management work. However, due to the complexity of the thermodynamics [24, 25], the research done in profiling using real measurements from scalar sensors are limited. It is difficult to find a suitable model to describe the complex phenomena of heat transfer and air distribution in datacenters.

In [43], authors benchmark and profiled 22 data centers about their energy usage. They stress that energy benchmarking using a metric, which compares energy used for IT equipment to energy used for Heating, Ventilating, and Air Conditioning (HVAC) systems, reveals some data centers perform better than the others and the key part is air management for effective and efficient cooling. This paper mainly compares the energy cost for different datacenters, which have different layouts, but relying on only heuristic analysis.

Thermally profiling datacenters based on measurements may be expensive because there are extra costs for sensing, monitoring and processing measurements(i.e., temperature and humidity sensors, thermal cameras etc.). However, this measurement-based approach provides high accuracy and less delay compared to simulation based approach. In using measurement-based approach, cooling system directly receives measurements from the sensors in real-time while simulation-based approach requires time to solve complex differential equations. Hence, measurement-base approach is able to capture the dynamic workload distribution in space easier than simulation-based approach.

One of the challenges to this arrangement is to provide the proper distribution of chilled air to the racks provided by computer room air-conditioning (CRAC) units situated on the raised floor. Authors in [41] show the perforated tile distribution on

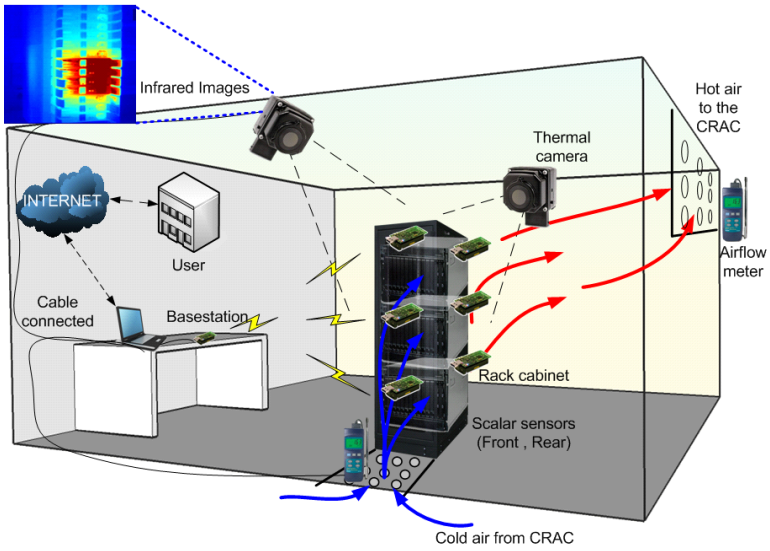

Figure 2.3: Network of heterogeneous sensors (scalar temperature and humidity sensors, airflow meters, and thermal cameras) monitoring a rack of servers in a machine room.

the inlet temperatures of racks located on the raised floor. [24] developed a three-dimensional model of a laboratory data center and experimentally verified the numerical results to ensure the specified inlet air temperatures to the computer systems met the temperature limits. [33] thermally profiled a non-raised-floor data center in a small office and then compared the results to a Computation Fluid Dynamics (CFD) model of the space. The results compared favorably. Data from other electronic equipment rooms are available, but it is very difficult to glean important information or correlate or compare data. [14] thermally profiled a data center in space in detail. They provide air flow rate of the perforated floor tiles to analyze the relation between temperature and air flows. [42] reviewed the existing literatures and presents future directions to increase data center energy efficiency.

We advocate hybrid approach — Logical Bridge - Sensing Infrastructure, Thermal cameras are needed for accurate measurement and reaction, modeling, proactive actions.

2.2.3 Sensing Infrastructure

We present innovative *communication and coordination solutions* for enabling self-organization of a network of external heterogeneous sensors – composed of thermal cameras, scalar temperature and humidity sensors, and airflow meters (Fig. 2.3) – into

a multi-tier sensing infrastructure that monitors the complex thermodynamic phenomena inside a datacenter [44]. The sensing infrastructure exploits the spatio-temporal correlation in the observed phenomena and the temperature distribution map provided by thermal images, and decides on-the-fly the granularity at which it should sample over space and time (*adaptive sampling*), thus enabling efficient real-time monitoring of datacenters. Incorporation of real-time measurements (such as temperature, air-flow rate, etc.) into heat generation and extraction rate models helps estimate the heat imbalance, which allows prediction of the future temperature map of the datacenter.

Energy-efficient proactive autonomic resource provisioning decisions as well as cooling system optimization in datacenters can leverage this thermal awareness. We have shown in [12], how heat-imbalance-based proactive datacenter management (cooling system optimization) is superior in terms of energy-efficiency and minimization of risk of equipment failures compared to its conventional temperature measurement-based reactive counterpart. Our novel multi-tier approach eliminates the need for collecting and processing large volumes of data from the entire datacenter. Instead, it enables to intuitively *zoom-in* only into the problem areas like thermal hotspots (higher temperature regions), collect, and analyze thoroughly the most relevant data from multiple sources (i.e., server and network element logs, internal and external sensors). This drastically improves the responsiveness of algorithms for time-critical management autonomic thermal-aware anomaly detection and classification.

As mentioned earlier, information from a network of external heterogeneous sensors is essential for efficient autonomic thermal-aware datacenter management. However, such a network would not scale in terms of overhead (communication, computation, and energy) and cost when the size of the datacenter and its server density increase significantly. For example, consider instrumenting a large High-Performance Computing (HPC) datacenter consisting of 1000 racks and 50 blade servers in each rack, with external temperature and humidity sensors on each server (50000 in total). The amount of sensed information collected and processed every second at a monitor node would be of the order of gigabits, thus increasing the strain on the communication and computation

resources. We have designed two innovative solutions, namely, autonomic adaptive sampling and coordinated hotspot detection and localization, for enabling self-organization of these heterogeneous sensors into an intelligent multi-tier sensing infrastructure that adaptively samples data to be fed into our heat imbalance model. As a result, the sensors function not only as passive measurement devices monitoring thermal phenomena in a datacenter, but as intelligent data processing instruments capable of data quality assurance, statistical synthesis, and hypotheses testing as they stream data from the physical environment to the computational world.

2.3 Thermal Decision Making

In this section, firstly, we review the state of the art in thermal management [31, 42] at the environment as well as application layers. Secondly, we review power management techniques (such as DVFS and pinning at the physical resource layer) that we propose to exploit for energy-efficient thermal management in place of costly VM migrations. Thirdly, we discuss prior work in the field of energy-efficient VM allocation, which is key for VM migration decisions in reaction to thermal anomalies. Finally, we review existing cross-layer solutions, which are characterized by pair-wise interactions between layers.

2.3.1 Thermal management at the environment layer

Prior research efforts in this area focus on management of heat extraction in a datacenter [15, 31, 40, 41]. In [43], Greenberg et al. profiled and benchmarked the energy usage of 22 datacenters and concluded that the key to energy efficiency is air circulation management (for effective and efficient cooling). Lee et al. [12] propose a proactive control approach that jointly optimizes the air conditioner compressor duty cycle and fan speed to prevent heat imbalance and to minimize the cost of cooling in data centers. As many datacenters employ raised floors with perforated tiles to distribute the chilled air to racks, researchers have tried to gain valuable insights into efficient airflow distribution strategies in such datacenter layouts [15, 31, 40, 41]. Other research

efforts were aimed at improving the efficiency of cooling systems through thermal profiling [37–39, 45], which is the extraction of knowledge about air and heat circulation using measurements from scalar sensors and mathematical models. However, capturing this complex thermodynamic phenomena using compute-intensive models [24, 25] (e.g., computational fluid dynamics) is prohibitive in terms of computational overhead.

2.3.2 Thermal management at the application layer

Another popular approach to thermal management has been controlling the heat generation inside a datacenter [46–50] through thermal-aware workload placement. Moore et al. [46] propose the use of "offline experiments" to characterize the thermodynamic phenomena (heat recirculation) inside the datacenter and schedule workloads by taking the temperature distribution into account. In [47], Moore et al. eliminate the need for the aforementioned offline experiments and propose an online machine-learning-based method to model the thermal behavior of the datacenter. Bash et al. [49] propose a policy to place the workload in areas of a datacenter that are easier to cool, which results in cooling power savings. They use scalar temperature sensor measurements alone to derive two metrics that help decide whether to place workload on a server or not. Tang et al. [50] develop a linear, low-complexity process model to predict the equipment inlet temperatures in a datacenter given a server utilization vector and formalize (mathematically) the problem of minimizing the datacenter cooling cost as the problem of minimizing the maximal (peak) inlet temperature through task assignment. In [51], Mukherjee et al. explore a spatio-temporal thermal-aware job scheduling as an extension to spatial thermal-aware solutions like [46–48, 50].

2.3.3 SLA-management at the virtualization layer

VM migrations are performed either reactively [52] or proactively [53] in such a way as to avoid equipment overheating and/or SLA violations. Kochut et al. [54] provide an estimate of the expected improvement in response time due to a migration decision and determines which VMs are best candidates to be placed together. In [55], Hermenier et al. determine the order in which the VM migrations should occur in addition to

deciding which VMs to migrate so to minimize the impact on application performance in terms of execution time. Stoess et al. [56] developed a multi-tiered infrastructure that enables intra-node virtual CPU (vCPU) migration and inter-node live VM migration for workload consolidation and thermal balancing. Voorsluys et al. [57] present an evaluation on the effects of live migration of virtual machines on the performance of applications running inside Xen VMs.

On-demand server resource provisioning techniques monitor the workloads on a set of VMs and adjust the instantaneous resources availed by VMs. Song et al. [58] propose an adaptive and dynamic scheme for adjusting resources (specifically, CPU and memory) among virtual machines on a single server to share the physical resources efficiently. Menasce et al. [59] proposed an autonomic controller and showed how it can be used to dynamically allocate CPUs in virtualized environments with varying workload levels by optimizing a global utility function. Meng et al. [60] exploited statistical multiplexing of VMs to enable joint VM provisioning and consolidation based on aggregated capacity needs. However, all the aforementioned techniques are aimed at satisfying the resource utilization level guarantees and do not consider the application-level performance (execution time).

2.3.4 Cross-layer solutions for power-, thermal-, and QoS-management

Nathuji et al. [61] investigate the integration of power management at the physical layer and virtualization technologies. In particular they propose VirtualPower to support the isolated and independent operation of virtual machines and control the coordination among virtual machines to reduce the power consumption. In [62], Nathuji et al. consider the heterogeneity of the underlying platforms (in terms of processor and memory subsystem architecture) to efficiently map the workloads to the best fitting platforms. Laszewski et al. [63] present a scheduling algorithm for VMs in a cluster to reduce power consumption using DVFS. Kumar et al. [64] present vManage, a practical coordination approach that loosely couples platform and virtualization management aimed at improving energy savings and QoS and at reducing VM migrations.

[EXPLAIN]advocate a joint approach so to minimize the risk of overheating of

servers while simultaneously maximizing resource utilization and cooling efficiency [65] [66]

Heath et al. [48] propose emulation tools (Mercury and Freon) for investigating the thermal implications of power management. In [67], Ramos et al. present C-Oracle, a software infrastructure that dynamically predicts the temperature and performance impact of different thermal management reactions (such as load redistribution and dynamic voltage and frequency scaling) into the future, allowing the thermal management policy to select the best reaction.

Recently, researchers have started to focus on application- or workload-aware VM consolidation that not only achieves all the objectives as its traditional resource-utilization- and energy-aware counterparts but also ensures minimum degradation to application performance due to resource multiplexing and virtualization overhead [68–70]. Application-aware VM allocation is not only aimed at energy-efficient VM consolidation but also at co-locating VMs that are "compatible" so that further gains can be achieved in terms of energy savings and overhead for virtualization. In [68, 69], the authors propose to consolidate VMs with similar memory content on the same hosts for higher memory sharing and Govindan et al. [70] propose to consolidate based on inter-process communication patterns. Zhu et al. [71] propose pSciMapper, a power-aware method for consolidation virtual machines that host scientific workflow tasks. They use a dimensionality reduction technique, Kernel Canonical Correlation Analysis (KCCA), to associate temporal features extracted from time series data about resource requirements of a workflow task with its power consumption and execution time (performance). Information about the tasks' power consumption and performance are exploited in an online consolidation algorithm.

All the aforementioned cross-layer solutions are characterized only by pair-wise interaction between two layers - physical and virtualization layers, physical and environment layers, or application and virtualization layers. In our proposed thermal-aware management solution interactions between the application and the virtualization layers are leveraged for proactive VM allocation while interactions among all the four layers

are exploited for reactive thermal management. In other words, in our thermal management solution, the environment layer alerts the virtualization layer and physical layer about undesired thermal behavior. These two layers jointly decide whether to use power management techniques at the physical layer or VM migration with help from application layer to alleviate the thermal anomalies.

Chapter 3

Proactive Thermal Management in Green Datacenters

3.1 Overview

The increasing demand for faster computing and high storage capacity has resulted in an increase in energy consumption and heat generation in datacenters. Because of the increase in heat generation, cooling requirements have become a critical concern, both in terms of growing operating costs as well as their environmental and societal impacts (e.g., increase in CO_2 emissions, overloading the electric supply grid resulting in power cuts, heavy water usage for cooling systems causing water scarcity, etc.) Many current datacenters are not following a sustainable model in terms of energy consumption growth as the rate at which computing resources are added exceeds the available and planned power capacities. For these reasons, there is a need for realizing environment friendly computing systems that maximize energy and cooling efficiency. Technical advances are leading to a pervasive computational ecosystem that integrates computing infrastructures with embedded ssensors and actuators, thus giving rise to a new information/sensor-driven and autonomic paradigm for managing datacenter cooling systems.

Due to the increasing costs and high energy consumption of current cooling systems in datacenters, energy efficient and intelligent cooling solutions are required to minimize these costs and consumption. Empirical data from Little Blue Penguin cluster shows that every 10^oC rise in temperature results in a doubling of the system failure rate, as per Arrhenius' equation applied to microelectronics, which increases the Total Cost of Ownership (TCO) significantly [72]. Overheating of components causes *thermal cycling*, which eventually leads to device failure, thus affecting the TCO [16]. Cooling

systems aim at effectively maintaining the temperature of the datacenter. For robustness and safety, over-provisioning is often implemented to avoid any loss to property due to unforseen perturbations. According to Lawrence Livermore National Laboratory (LLNL), for every watt of power its IBM BlueGene/L consumes, 0.7 W is required to cool it [73, 74]. As the number of datacenters with high processing power increases, the expense to run cooling equipment such as chillers, compressors, and air handlers also increases. According to [75], it is predicted that datacenter energy consumption in the U.S. will reach 100 billion kWh/year by 2011 with a corresponding energy bill of approximately $7.4 billion.

Current cooling solutions in datacenters rely on *reactive techniques*, which aim at keeping the temperature at a fixed value. Existing datacenter cooling systems control the temperature/humidity of the air based on the temperature external to the machines, i.e., the temperature inside a datacenter. Some of the cooling system control mechanisms are based on the internal temperature of the racks or blades. Irrespective of the external or internal temperature, the reactive approach has numerous disadvantages: i) it takes a corrective action after the temperature has crossed a threshold and may not be able to prevent damages in certain cases where temperature rises above the safe operating range of the internal components, ii) it is very difficult to determine the optimal threshold range as it should not be too high that a small increase above it damages the components or too low to waste energy required for cooling, and iii) if the threshold range is too small it causes cycling (or hysteresis) in the Computer Room Air-Conditioning (CRAC) unit and in turn reduces its life; conversely, if the threshold range is too high the response time of the system increases with a possible risk of damage to the internal components of the machines.

Due to the numerous disadvantages of temperature-based reactive approaches, in this work we propose a heat-imbalance estimation-based proactive approach that optimizes the cooling system operation by minimizing the cooling costs and the risk of damage to components due to overheating. The proposed proactive approach controls the cooling system before the heat imbalance can raise the temperature and cause damage to the internal components of a machine. Heat imbalance is the difference

between the heat generated and heat extracted (under ideal conditions heat imbalance should be zero). The proactive approach has numerous advantages over the reactive approach: i) there is no need for setting thresholds as needed in reactive approaches, ii) it is intrinsically predictive in nature as it estimates the heat that will be generated in the future (based on information on scheduled type/intensity of workload) and, accordingly, adjusts the operation of the CRAC unit, iii) it observes the 'cause' instead of the 'effect', i.e., it estimates the heat imbalance rather than measuring the rise in temperature caused by it.

The remainder of this chapter is organized as follows. In Sect. 3.2, we address the background and related work. In Sect. 3.3, we formulate our heat-imbalance estimation-based proactive control approach. In Sect. 3.4, we describe our mathematical model in details. In Sect. 3.5, we present the performance evaluation. Finally, in Sect. 3.6, we draw the conclusions and discuss future work.

3.2 Background and Related Work

There are two main approaches for thermal management of datacenters; one is *mechanical design based* and the other is *software based*. The former focuses on how to effectively distribute cold air by managing cooling infrastructure, while the latter focuses on how to balance or migrate jobs in such a way as to minimize heat imbalances.

Mechanical design-based approaches study the airflow models, datacenter design, and cooling system design. Datacenter design plays an important role in the efficient thermal management. The cooling systems used for current datacenters are chilled-water cooled CRAC units that supply a raised floor plenum underneath the racks with cold air. Perforated tiles are located near the racks to transfer the cool supply air to the front of the racks. The hot exhaust air from the racks is then collected from the upper section of the facility by the CRAC units, thus completing the airflow loop as shown in Fig. 3.1(a). Racks are typically arranged in rows with alternating airflow directions, forming 'hot' and 'cold' aisles [76]. This hot-and-cold aisle approach attempts to prevent mixing of the hot rack exhaust air and the cool supply air drawn into the racks with the

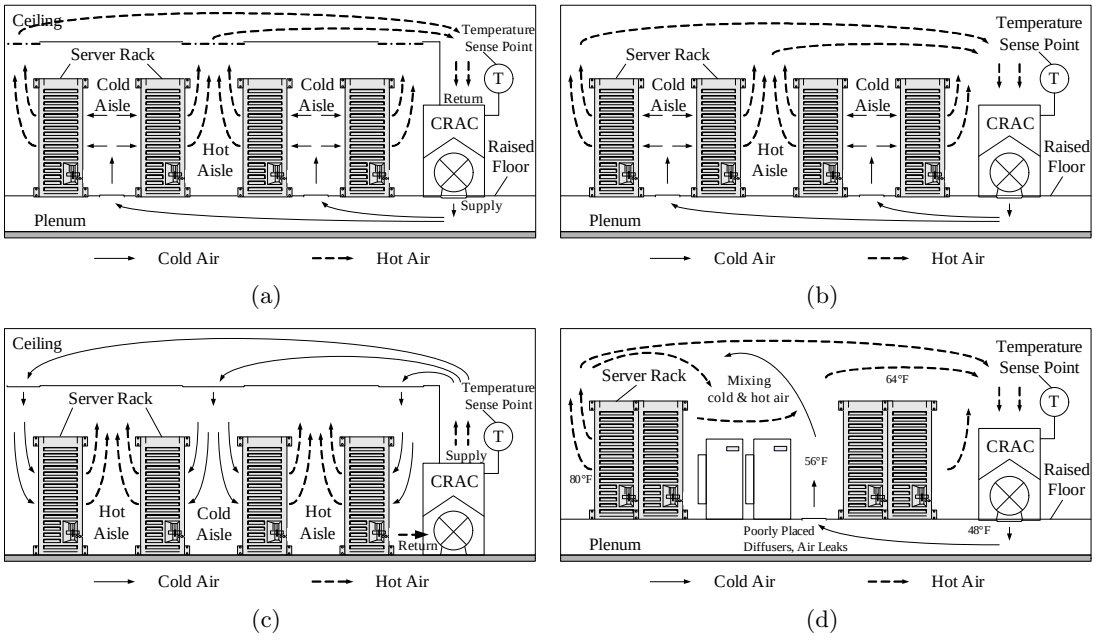

Figure 3.1: Current cooling design schemes: (a) underfloor supply and overhead return, (b) underfloor supply and horizontal return, (c) overhead supply with horizontal return, and (d) poor airflow condition.

objective of increasing the overall efficiency of the air delivery and collection from each rack in the datacenter. Different thermal efficiencies may be achieved with alternate configurations, as illustrated in Figs. 3.1(b) and 3.1(c).

Any configuration for the CRAC unit can be applied, but only certain combinations are feasible due to mechanical constraints of the CRAC units (e.g., to not introduce an excessive amount of duct work). With these constraints, achieving thermal (energy) efficiency is complicated and there is no single optimal solution. In [77], the authors made an attempt to compare cooling efficiencies among four airflow distribution systems in high heat density room: underfloor supply/overhead return, underfloor supply/horizontal return, overhead supply/underfloor return, and overhead supply/horizontal return. In [78, 79], the authors expand on the concepts proposed in [77]. Datacenter design guidelines from PG&E [80] - mainly based on [76] -, presents a poorly designed datacenter room (Fig. 3.1(d)) cooled by a raised floor system, which has often trouble maintaining an appropriate room temperature. In [43], the authors benchmark 22 datacenters according to their energy usage and conclude that energy benchmarking

using a specific metric - ranging from the energy used for IT equipment to the energy used for Heating, Ventilation, and Air Conditioning (HVAC) - is extremely helpful to understand why some datacenters perform better than others.

Another important aspect of mechanical design for cooling system involves airflow distribution. In [15, 31, 41, 81], the authors provide some insights into the airflow distribution from perforated tiles and raised floor design in datacenters. In [82, 83], the authors thermally profiled a datacenter in space and time, and analyzed trends and correlations among collected measurements. Basic mathematical modeling and parameters for modeling datacenter are proposed in [39]. In [31], the authors comment on the challenges associated with thermal management in datacenters. In [42], the authors review the existing literature on datacenter thermal modeling work. However, due to the complexity of the thermodynamics [24,25], the research done in formulating models suitable to describe the complex phenomena of heat propagation and air distribution in datacenters has been limited.

On the other hand, software-based approaches focus on minimizing the cooling cost by distributing or migrating jobs. In [34], the authors emulate the thermal behavior of server systems to manage datacenter cooling. Using the emulator, a system named Freon monitors temperature changes and, if the temperature of a machine crosses a threshold (defined thermal emergency), Freon redistributes the jobs; however, cooling cost is not considered in this study. In [3], the authors introduce the concept of power budget, which is the product of power and temperature. Higher power budget means that a machine has more capacity to accept a job in terms of temperature and power. The authors propose two scheduling algorithms based on power budget to fairly and efficiently distribute the workload. In [51], the authors propose energy savings by temporally spreading the workload and assigning it to energy-efficient computing equipment. However, these works are not coupled to a physical datacenter model and only consider scenarios that are reactive in nature. In [84], the authors propose a cooperative power-aware game theoretic solution for job scheduling in grids to minimize the energy consumption while maintaining a specific Quality of Service (QoS) level. They highlight the fact that it is not enough to minimize the total energy of grid but

there is the need to minimize energy locally at different providers in the grid. The proposed solution simultaneously minimizes the energy used by all providers so to be fair to all users. The energy usage is kept to minimum level while maintaining the desired QoS level. The authors claim that the proposed solution is robust against prediction inaccuracies. In [85], the authors study the problem of task allocation onto a computational grid and aim at simultaneously minimizing the energy consumption and the makespan (time difference between the start and finish of a sequence of jobs), subject to the constraints of deadlines and tasks' architectural requirements. The solution is proposed from cooperative game theory based on the concept of Nash Bargaining Solution (NBS). The frameworks proposed in [84, 85] do not take the datacenter design or airflow characteristics into consideration and are, hence, not suitable for optimizing the cooling system performance.

In current datacenter thermal management, the mechanical- and software-based approaches are usually independent on each other. However, there exists a strong correlation between the two; for this reason, there is the need to combine the two approaches to obtain an optimal cooling solution. The proactive approach proposed in this work can adapt itself to any type of mechanical design and it also considers the distribution and type of workload running. It combines the mechanical aspect of a datacenter with the software-based scheduling approach to optimize the performance of the cooling system.

3.3 Problem Formulation

In this section, we formulate the mathematical model for heat transfer in datacenters that our solution is based on. The heat transfer model is divided into three parts as follows: i) overall heat circulation model, ii) heat generation model, and iii) heat extraction model. We describe heat generation, extraction, and circulation based on the fundamental thermodynamic principles in physics. It is assumed that the datacenter is built on a hot-and-cold aisle based raised plenum design, supplying cold air from the raised plenum and returning hot air to the ceiling.

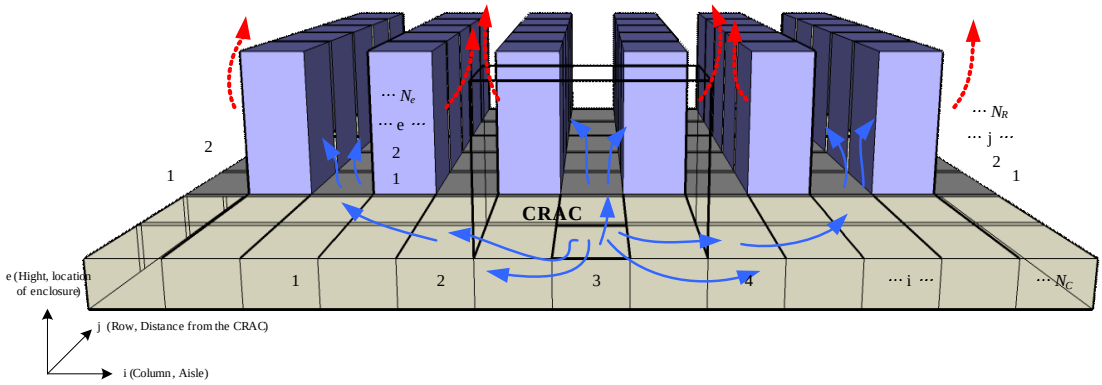

Figure 3.2: Layout of datacenter and location of a blade.

3.3.1 Air Circulation Design

The proposed datacenter model is designed in a 3-dimensional space as shown in Fig. 3.2. In Fig. 3.2, the columns represent aisles, the rows represent the distance from the CRAC, and the height represents the enclosure number from the bottom. These are referred to as 'i', 'j', and 'e', respectively. The latest rack design includes 3-4 enclosures, each containing 10-20 integrated vertical blades with an independent cooling module to cool all the blades in it. We assume that every odd numbered aisle is a cold aisle and every even numbered aisle is hot aisle. For the sake of clarity, the notation used in the model are summarized in Table 3.1.

The proposed model is based on heat imbalance equations. Heat is mainly generated by the processor and subsystems, i.e., memory and storage devices, I/O subsystem, network interface card, etc., is extracted by the fans in the enclosure and the fan in the CRAC unit. We assume that there is no external heat source and the room is thermally insulated. The notations used to describe the datacenter thermal flow model is similar to the ones used in [39].

In order to calculate the heat imbalance and flux, we assume that our datacenter is extended in a 2-dimensional space as shown in Fig. 3.3(a) and explained following airflow from the CRAC ('①' in Fig. 3.3(a)) and returning to CRAC ('⑤-⑥' in Fig. 3.3(a)). First, a cold air stream from the CRAC unit is pumped through the plenum at flow rate m_{crac}^{out} and temperature T_{crac} in '①'; this flow is evenly divided and exhausted through the plenum and perforated tiles '②' in the cold aisle ideally. However, in real

Table 3.1: Notations.

Nomenclature	
ΔI	Heat imbalance [J]
C_p	Specific heat of air at constant pressure [J/kg · K]
m	Mass airflow rate [kg/s]
M	Mass of air [kg]
T	Temperature [K]
ρ	Density of air [kg/m^3]
A	Area of the component [m^2]
N_C	Total number of columns / aisles
N_R	Total number of rows
N_A	Total number of Rack
N_E	Total number of enclosures in a rack
Subscript	
e	Enclosure
r	Rack
i	Column
j	Row
$crac$	CRAC unit
$tile$	Perforated vent tile
$room$	Room where the equipment is placed in a datacenter
Superscript	
in	Inlet
out	Outlet

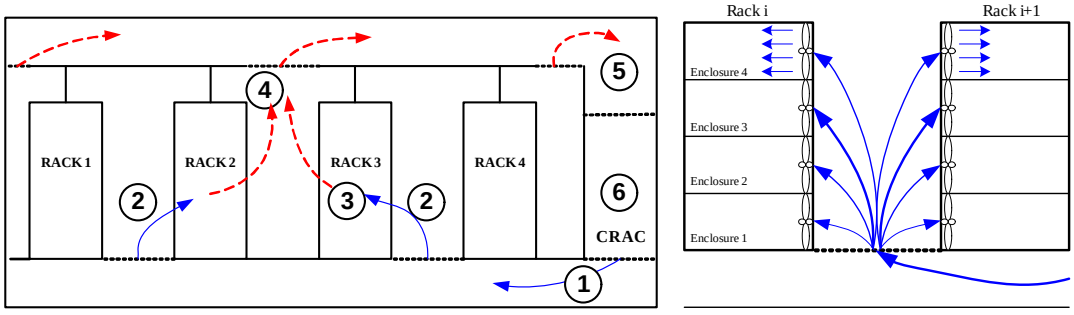

(a) (b)

Figure 3.3: (a) 6 different places to consider mathematical model; (b) Rack level airflow.

case, perforated tiles can be modeled as a lumped resistance using the relationship $\Delta P = \beta \cdot (m^{out}_{crac}/\rho)^2$, where the coefficient β can be found in standard flow resistance handbooks (e.g., [86]). Experimental values of β have been proposed in [82, 83]. Using equation for *dynamic pressure* $\Delta P = \rho \cdot v^2/2$, where v is the fluid velocity, m^{in}_{tile} and m^{out}_{tile} can be written as follows,

$$m^{in}_{tile} = m^{out}_{tile} = A_{tile} \cdot \sqrt{\beta \cdot \frac{2 \cdot (m^{out}_{crac})^2}{\rho}}. \tag{3.1}$$

Between '②' and '③', inlet airflow rate of an enclosure is proportional to inlet airflow of a rack. Ideally, if there is no air circulation, leakage or Bernoulli effect[1], and the fans on every enclosure have the same speed then, m^{in}_e is m^{in}_r/N_E and m^{in}_r is half of m^{out}_{tile} because the air flowing from a tile splits into two racks. The relations of those parameters are,

$$m^{in}_r = \sum_{e=1}^{N_E} m^{in}_e = N_E \cdot m^{in}_e = \frac{m^{out}_{tile}}{2}. \tag{3.2}$$

In ideal case, The inlet mass airflow rate in through '③' for all the racks are the same as,

$$m^r_{in} = m^l_{in}, \forall_{r,l}. \tag{3.3}$$

As these ideal assumptions are often not valid in realistic scenarios, to improve the accuracy of the model we need realtime measurements collected via a sensing infrastructure, as discussed in Sect. 3.6. Figure 3.3(b) shows the flow from the tile to the rack and Equation 3.2 shows the relation between m^{out}_{tile} and m^{in}_r, where 'r' denotes the racks between aisle i to $i+1$ (column) and j (row) cell. Cooling fans in each enclosure suck the air streams in at flow rate m^{in}_e, cool the heated components down in '③', and the air stream in the enclosure flows to the back of the racks with flow rate m^{out}_r in '④'.

We formulate ***heat-imbalance model*** in a datacenter as follows, which explains

[1]In fluid dynamics, the Bernoulli's principle states that an increase in the speed of the fluid occurs simultaneously with a decrease in pressure or a decrease in the fluid's potential energy.

the heat exchange in '③' as,

$$\Delta I_{ij}^{e} = \int_{t_1}^{t_2} (h_{ij}^{e} - q_{ij}^{e})dt = M^{e} \cdot C_p \cdot \Delta T_{[t_2,t_1]}^{e}, \tag{3.4}$$

where,

- ΔI_{ij}^{e} denotes the heat imbalance of the enclosure e in the cell (i, j), which is between i and $i + 1$, and j, during the time between t_1 and t_2;
- h_{ij}^{e} is the rate of heat generation from the enclosure e in the cell (i, j) [J/s];
- q_{ij}^{e} is the rate of heat extraction from the enclosure e in the cell (i, j) [J/s];
- If ΔI_{ij}^{e} is positive (i.e., $h_{ij}^{e} > q_{ij}^{e}$), the temperature in the enclosure e increases (hence, $\Delta T^{e} > 0$);
- If ΔI_{ij}^{e} is negative (i.e., $h_{ij}^{e} < q_{ij}^{e}$), the temperature in the enclosure e decreases (hence, $\Delta T^{e} < 0$).

Eq. (5.1) shows the difference between heat generated and heat extracted in an enclosure. If the heat difference is positive, the enclosure temperature will increase; if the imbalanced heat of the entire datacenter is set up as a function of blades of the enclosures, and enclosures of the racks, then we have,

$$\Delta I = \sum_{i=1}^{N_C} \sum_{j=1}^{N_R} \sum_{e=1}^{N_E} \Delta I_{ij}^{e} = M_{room} \cdot C_p \cdot \Delta T_{[t_2,t_1]}^{room}. \tag{3.5}$$

If the heat difference is positive, the average temperature of the datacenter will increase.

From the experiments on our test server, which has the following configuration, eight 8 core Intel Nehalem processors, 138 GB of RAM and 500 GB of storage, it was observed that the dominant power utilizing subsystems were the CPU, I/O subsystem, memory and storage subsystem, and the Network Interface Card (NIC). Out of the total power utilized, a certain percentage was dissipated as heat. The percentage of power dissipated as heat by the subsystems is denoted by α^{sub} [%] as detailed in Sect.

3.4.1. The rate of heat generation by a subsystem h_{ij}^{e} [J] is given by,

$$h_{ij}^{e} = P_{ij}^{e,cpu} \cdot \alpha^{cpu} + P_{ij}^{e,IO} \cdot \alpha^{IO} + P_{ij}^{e,mem,stg} \cdot \alpha^{mem,stg} + P_{ij}^{e,NIC} \cdot \alpha^{NIC}, \tag{3.6}$$

where P^{cpu}, P^{IO}, $P^{mem,stg}$, P^{NIC} is the power utilized by CPU, I/O subsystem, memory and storage devices, and the NIC, respectively, and α^{cpu}, α^{IO}, $\alpha^{mem,stg}$, α^{NIC} are the respective percentage power dissipation factors. Using equation 3.29, the total rate of heat generation 'H' can be calculated as,

$$H = \sum_{i=1}^{N_C} \sum_{j=1}^{N_R} \sum_{e=1}^{N_E} h_{ij}^{e}. \tag{3.7}$$

On the other hand, the heat is extracted by the inlet-air and flows out with the outlet-air, which can be calculated as,

$$q_{ij}^{e} = m_{ij,in}^{e} \cdot C_p \cdot (T_{ij,out}^{e} - T_{ij,in}^{e}). \tag{3.8}$$

By measuring inlet and outlet temperature and airflows, we can calculate how much heat is extracted from the enclosure e located in i (row) and j (column), i.e., in cell (i, j). The total rate of heat extracted can be computed as,

$$Q = \sum_{i=1}^{N_C} \sum_{j=1}^{N_R} \sum_{e=1}^{N_E} q_{ij}^{e}. \tag{3.9}$$

Temperature of the inlet airflow $T_{ij,in}^{e}$ varies depending on whether the air compressor is 'on' or 'off' in equation 3.8. In Fig. 3.4, if the air compressor is on (Q_{off}), the air from the server room with the temperature (T^{room}) can extract the heat because only the fan is working, but if the air compressor is on (Q_{on}), then the air with the temperature from the CRAC unit (T^{crac}) can extract the heat because both the compressor and fan are working. The heat generation 'H' is independent on whether air compressor is working or not because it is generated based on the workload and it's distribution. In equilibrium state, 'T^{crac}' is lower than 'T^{room}' and 'T^{room}' is lower than 'T_{out}' ($T^{crac} < T^{room} < T_{out}$).

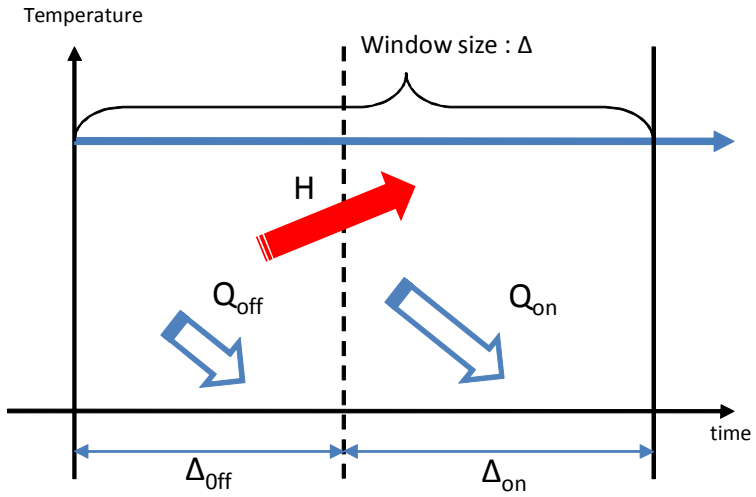

Figure 3.4: Q_{on} is the heat extraction when compressor is on, Q_{off} is the heat extraction when compressor is off which are proportional to T_{out} - T^{crac} and T_{out} - T^{room} respectively, while H is the heat generation.

Heated air streams are collected through the ceiling in '⑤' and returned to the CRAC unit. Finally, the circulated air enters the CRAC unit and is compressed and cooled again in '⑥'. In an ideal case, it holds,

$$m_{out}^{crac} = \sum_{r=1}^{N_A} m_{in}^{r} = \sum_{r=1}^{N_A} m_{out}^{r} = m_{in}^{crac}. \tag{3.10}$$

The air can be mixed for a variety of reasons, some of which are discussed in Sect. 3.2. The phenomenon and the effect of mixing cold air and hot air streams are discussed in [39]. Supply using Heat Index (SHI) is denoted as follows,

$$SHI = \frac{\sum_{i=1}^{N_C}\sum_{j=1}^{N_R}(T_{in,ij}^{r} - T_{out}^{crac})}{\sum_{i=1}^{N_C}\sum_{j=1}^{N_R}(T_{out,ij}^{r} - T_{out}^{crac})}. \tag{3.11}$$

Because the air from the CRAC unit T_{out}^{crac} is not at the same temperature as the rack T_{ij}^{r}, we can assume that there is a re-circulation and mixing of hot-cold air in the datacenter. We can apply the simple SHI configuration for our model to explain air re-circulation. However, this model using inlet and outlet airflow can be applied only if we know the inlet and outlet temperature and airflows at each blade. This requires a sensing infrastructure that measures airflows and temperature to quantify the amount of heat extracted.

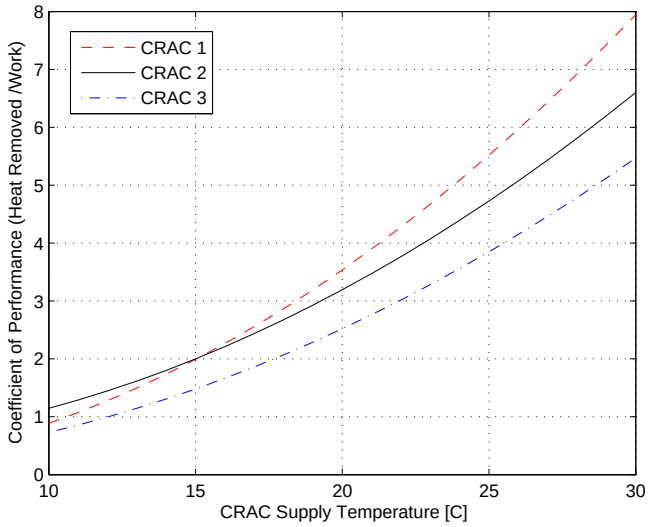

Figure 3.5: Coefficient of Performance (COP) curve for the different chilled-water CRAC units. COP can vary for different CRAC unit depending on the type of fan and compressor used. CRAC 2 shows COP of cooling system at the HP Labs Utility Datacenter [3]. This COP is also used in our simulations.

3.3.2 Air Cooling System Design

Cooling is the process where heat is transferred from a lower to a higher temperature level by doing work on a system in order to extract the heat. Most datacenters use chilled-water air conditioning system. With chilled water air conditioning, the refrigeration equipment (compressor, condenser, evaporator, etc.) does not directly cool the air; rather, it uses chilled water to cool the air, where chilled water is pumped to the cooling coils and a fan draws the air through the chilled water pipe to the coils, thus cooling the air. With chilled water air conditioning, the compressor is usually mounted on a rack or a frame, within a few feet from the evaporator that cools the chilled water. The efficiency of this cycle can be determined by several factors such as airflow and chilled water temperature, and can be quantified by the Coefficient Of Performance (COP). The COP is the ratio of total heat removed 'Q' from low-temperature level and the energy input used (W) as,

$$COP(T) = \frac{Q}{W}. \tag{3.12}$$

Since the COP is inversely proportional to the W, a higher COP means that more

heat 'Q' can be removed by doing less work (W), as given in (3.12). The COP can be calculated in each cooling cycle and through (3.12), we can calculate how much work is needed to extract a certain amount of heat. For example, a cooling cycle with COP of 2 will consume 30 kWh to extract 60 kWh of heat, while a cooling cycle of COP of 3 will consume 20 kwh to remove the same amount of heat. Figure 3.5 shows COP values for different CRAC units. As the CRAC supply temperature increases, the COP also increases (in compliance with the second principle of thermodynamics). Consequently, our cooling cost 'E_{Total}' can be calculated as,

$$E_{Total} = E_{Compressor} + E_{fan}, \tag{3.13}$$

which is the total amount of energy needed to power the air compressor, '$E_{Compressor}$', plus energy needed to run the CRAC fan, 'E_{fan}'. If the duty cycle η of the compressor is represented as,

$$\eta = \frac{\Delta t_{on}}{\Delta t_{on} + \Delta t_{off}}, \tag{3.14}$$

where the window size of the compressor cycle is represented as '$\Delta = \Delta t_{off} + \Delta t_{on}$', then the work done to extract the amount of heat for a compressor can be calculated as,

$$E_{Compressor} = P_{AC} \cdot \Delta t_{on} = \frac{Q}{COP(T)}. \tag{3.15}$$

The *'affinity law'*, which describes how the performance of a centrifugal pump is affected by a change in speed or impeller diameter is given by,

$$P = P_{ref} \cdot \frac{\omega^3}{\omega_{ref}^3}, \tag{3.16}$$

where ω is the shaft rotation speed (fan speed) and P is the shaft power correspond to the ω. Note that the shaft power is proportional to the cube of rotation speed of the fan shaft. 'E_{fan}' can be calculated using this law. Specifically, once we know the reference-point power P_{ref} of the fan and its rotational speed ω_{ref}, which both vary depending on the manufacturer and type of fan, the power required to increase the fan

speed to ω in $[t_1, t_2]$ can be computed as,

$$E_{fan} = \int_{t1}^{t2} P_{ref} \cdot \frac{\omega^3}{\omega_{ref}^3} dt. \tag{3.17}$$

Also, the mass of airflow rate injected from the CRAC, 'm_{crac}', can be computed based on fan speed ω as,

$$m_{crac} = \rho(T) \cdot K_\omega \cdot \omega, \tag{3.18}$$

where $\rho(T)$ denotes density of the air in certain temperature and K_ω denotes coefficient of amount of air through the fan, when the speed of the fan is ω.

Since there is only one fan with airflow rate 'm_{crac}' is used in this model, we cannot extract the localized heat generated in each aisle or each rack. However if we use multiple fans, then the total mass airflow rate in each corridor is given by,

$$m_{crac} = m_{fan_1} + m_{fan_2} + m_{fan_3} \ldots, \tag{3.19}$$

where $m_{fan_{num}}$ denotes the mass airflow rate of each fan. If we install multiple fans, one for each corridor or each rack, then we can have more control over the net airflow rate. However, there is an additional cost for installation and operation of these fans. Power usage by a fan ranges from hundreds to thousands Watt and the power usage by the air compressor is hundreds of kilowatt. However, the energy savings obtained by just increasing fan speed and not increasing the compressor cycle to extract localized heat offsets the additional cost. In this way we can selectively extract localized heat without much extra cost.

3.3.3 Problem Formulation to Minimize Cooling Energy

We formulate the problem to optimize fan speed (ω^*) and duty cycle of the air compressor (η^*) to minimize the energy consumed by the compressor and fan as follows.

Given (offline): $T^{crac}, T_{set}, T_{init}, P_{AC}, \omega_{ref}, P_{ref}, M_{room}, \Delta, C_p, \rho(), COP(), K_\omega, N_C, N_R, N_E$

Given (online): $h^e_{ij}, T_{out}, T^{room}$

Find: $\omega^*,\ \eta^*$

Minimize: $E_{Total} = E_{Compressor} + E_{fan} = P_{AC} \cdot \eta \cdot \Delta + P_{ref} \cdot \frac{\omega^3}{\omega^3_{ref}}$

Subject To:

$$\Delta I = \int_{t_0}^{t_0+\Delta} H dt - Q_{on} - Q_{off} = M_{room} \cdot C_p \cdot (T_{set} - T_{init}); \tag{3.20}$$

$$H = \sum_{i=1}^{N_C} \sum_{j=1}^{N_R} \sum_{e=1}^{N_E} h^e_{ij}; \tag{3.21}$$

$$Q_{on} = \rho \cdot K_\omega \cdot C_p \cdot (T_{out} - T^{crac}) \cdot \eta \cdot \Delta; \tag{3.22}$$

$$Q_{off} = \rho \cdot K_\omega \cdot C_p \cdot (T_{out} - T^{room}) \cdot \eta \cdot (1 - \Delta). \tag{3.23}$$

Constraint (3.20) forces the heat imbalance 'I' to be equal to the amount of heat to adjust the temperature to the set point from the initial temperature, so that the server room temperature remains in equilibrium with the set point. By using this constraint and on-line and off-line parameters obtained from the datacenter, the fan speed ω and compressor cycle η can be optimized. If the amount of heat generated is the same as the amount of heat extracted, then there is no heat unbalance and the temperature stays in the equilibrium point. Equation (3.21) shows the heat generated from each component of the server blade. Equation (3.22) shows heat extraction when the compressor is turned on, and Equation (3.23) shows heat extraction when the compressor is turned off.

3.4 Proposed Solution

In this section, we propose our proactive approach model and later we describe the reactive approach. The proactive approach keeps the return temperature or internal blade-temperature in a safe operating range by jointly optimizing the duty cycle of air compressor and the CRAC fan speed before the rack/enclosures are heated using knowledge of the workload. The reactive approach is based on a feedback mechanism, in which the external temperature is adjusted back to a safe operating range when the heat generated by the rack/enclosures causes heat imbalance leading the external temperature to rise above a certain safe operating threshold temperature.

3.4.1 Proactive Approach

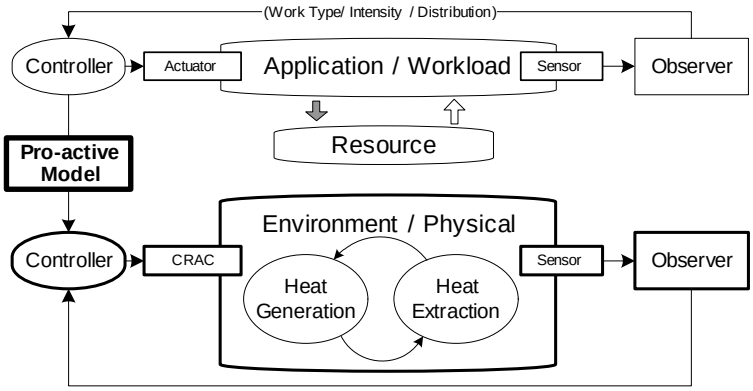

Figure 3.6: Proactive approach.

The proactive approach solves the problem at grass-root level. It is based on temperature change due to the heat imbalance between the heat generated and heat extracted. This proactive approach is *quantitative* in nature as it measures/estimates the heat imbalance. Conversely, compared to a proactive approach, a reactive approach is only *qualitative* in nature as it reacts to the changes in temperature. Using an analogy with kinematics in physics, we can say that *a proactive approach is analogous to determining the final position of an object by measuring its velocity instead of the position itself.*

The effectiveness of a proactive approach is based on having comprehensive knowledge about the behavior of the workload running and its utilization of the subsystems

in a blade. Knowing the subsystem usage pattern of the workload that is expected to run in advance and knowing the specifications of the subsystems by the application/workload layer (Fig. 3.6), the heat dissipated by each of the subsystems can be estimated. The subsystem usage pattern is obtained by observing the subsystem usage behavior of the selected workload over a certain period of time. The historical data obtained is used to generate workload patterns as shown in Fig. 3.7. These patterns can be used to estimate the power utilized and, in turn, the heat generated from the subsystems. Estimating the heat generated by subsystems and knowing the heat extracted by the cooling system, i.e., the heat imbalance, the temperature rise at a blade can be estimated. Based on this predicted temperature rise, the external temperature can be adjusted accordingly in order to maintain the internal temperature under safe operating threshold by the Environment/Physical Layer, as in the schematic in Fig. 3.6.

Based upon the subsystem usage pattern of the workload provided by the Application/Workload Layer, we propose two solutions; i) proactive approach using single fan, and ii) proactive approach using multiple fans. Single fan approach uses one fan for the CRAC unit; however, single fan approach could be too global to adjust local temperatures since only few estimations of overheating can trigger the entire CRAC unit to operate. This inefficiency may be reduced by the use of multiple fans assigned one for each aisle. For example, if the number of jobs assigned in a certain aisle is more than other aisles, we can optimize the fan speed based on the workload type and pattern. Since the energy needed for adjusting fan speed is much lower than the energy needed for adjusting the duty cycle of air compressor, we can reduce the energy consumption.

We obtain the historic subsystem usage data for some standard HPC benchmark workloads like, FFTW, HPL, NAS-benchmarks, from our test server to observe their subsystem usage pattern. The data in Fig. 3.7 shows the subsystem usage pattern of the *FFT* workload for CPU, I/O subsystem, storage device, and NIC utilization with respect to time in four subplots. Each subplot represents the magnitude of usage of a subsystem with respect to time. CPU usage is represented as percentage usage per second, I/O and storage usage is represented in the units of input/ouput instances or read/write instances per unit time, and the NIC usage is represented in units of data

bytes exchanged per unit time using the TCP protocol. The subsystem usage pattern of the workload provides us with the information about the time instances at which each subsystem is utilized and when it is idle. From the pattern in Fig. 3.7 we can estimate the power utilized by the subsystems and, in turn, the heat generated. The subsystem usage patterns are the back bone for estimating the heat generated at the blade level.

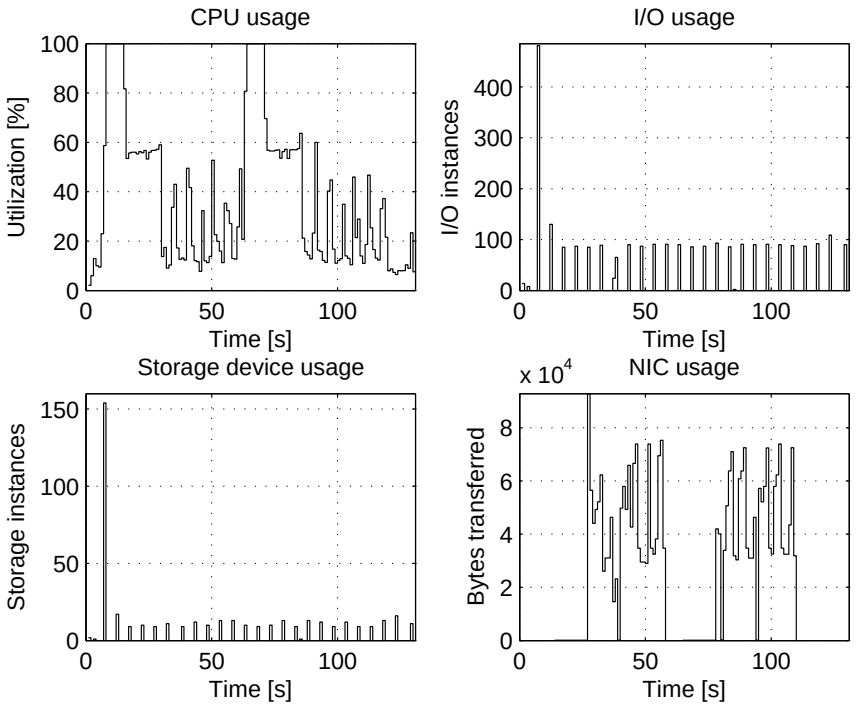

Figure 3.7: FFTW benchmark subsystem usage profile showing processor, I/O subsystem, memory subsystem, and the NIC usage.

Combining the power dissipated as heat at the CPU and other subsystems, the total power dissipation of the blade is estimated. All the subsystems are composed of semiconductor devices, hence we calculate the leakage power dissipated as heat from the formulas given in [87]. This is an approximate estimate as modeling the exact amount of heat dissipated is complicated and not absolutely necessary. In the model we are interested in estimating the maximum heat that can be generated at any of the subsystems at any time instant based on the workload pattern. Hence, we assume it is safe to use the leakage power formulas presented in [87].

The leakage power $P_{leakage}$ for a semiconductor chip as given in [87] is,

$$P_{leakage} = W_{avg} \cdot I_{leak} \cdot N_{trans} \cdot V_{dd}, \tag{3.24}$$

where W_{avg} is the average size of the transistor at the input gate, I_{leak} is the leakage current per unit width, N_{trans} is the total number of transistors in 'on' state, V_{dd} is the input voltage. In (3.24), N_{trans} is dependent on the device usage, which is obtained from the workload pattern, all the remaining parameters are obtained from the technical specifications of the motherboard and individual IC datasheets. Hence, (3.24) provides us with the direct relation between the subsystem utilization and heat dissipation. Also, some of the subsystems like the CPU have multiple sleep states also known as C-states and the value of N_{trans} varies depending on the power state in use, while other subsystems may have only two states, i.e., 'on' and 'off'.

The test server is configured to have the following specific configuration: 1)Advanced Configuration and Power Interface (ACPI) is enabled, and 2) processor dynamic frequency scaling is enabled. Since, ACPI is enabled, the CPU can transition to multiple C-states, with C_0 being the most power utilizing state or an active state, and C_n being the least power utilizing state or a deep sleep state [88]. Other subsystems, i.e., I/O subsystem, memory and storage devices, and the NIC, do not have any operating system based power management enabled and, hence, have only two states with D_0 being the 'on' state (or most power utilizing state) and D_n being the 'sleep' state (or least power utilizing state). With dynamic frequency scaling enabled, the CPU can transition to various predetermined frequency levels also known as P-states. Because the CPU has multiple sleep and power states, we calculate its power separately from other subsystems.

Power utilization of the processor in C_0 is a function of P-states or the frequency at which the processor is running. The power utilization of the processor is calculated as,

$$P^{cpu} = P^{cpu}_{C_0} + \cdots + P^{cpu}_{C_n}, \tag{3.25}$$

where $P_{C_0}^{cpu}$ is given as,

$$P_{C_0}^{cpu} = \sum_{j=0}^{k} P_{P_j}, \tag{3.26}$$

where P_{P_j} [W] is the power utilized in the P-state j assuming processor has k P-states. P-states are not relevant for sleep states (C-states other than C_0) as processor is inactive in them. Hence, the power utilized in C-states other than C_0 is given as $P_{C_n}^{cpu}$ [W], where 'n' is the C-state depth. The power utilized by other subsystems, i.e., I/O subsystem, storage devices, and the NIC, is given by,

$$P^{sub} = P_{S_0}^{sub} + P_{S_{idle}}^{sub}, \tag{3.27}$$

where P^{sub} [W] is the total power utilized by the subsystem and $P_{S_0}^{sub}$ [W] is the power utilized when the subsystem is in use or 'on' sate and $P_{S_{idle}}^{sub}$ [W] is the power consumed when the subsystem is idle or in 'sleep' state.

The percentage of power dissipated as heat $\alpha^{cpu,sub}$ [%] is given by,

$$\alpha^{cpu,sub} = \frac{P_{leakage}^{cpu,sub}}{P^{cpu,sub}} \times 100, \tag{3.28}$$

where $P^{cpu,sub}$ is the total power utilized by the CPU and subsystems, and $P_{leakage}^{cpu,sub}$ is the leakage power for the CPU and subsystem calculated from (3.24).

Total Heat generated h_{ij}^{e} [J] by the processor and subsystems in blade, over the time $t^{cpu,sub}$ is given by,

$$h = P^{cpu,sub} \cdot \alpha_{d}^{cpu,sub} \cdot t^{cpu,sub}. \tag{3.29}$$

From (3.29), we observe that h_{ij}^{e} is directly proportional to the product of power utilized by the processor and subsystems $P^{cpu,sub}$ and the time $t^{cpu,sub}$ when the processor and subsystems are 'on'.

3.4.2 Reactive Approach

The reactive approach is based on measuring the change in temperature and accordingly adjusting the duty cycle and fan speed of the CRAC unit in Fig. 3.8. Reactive approach

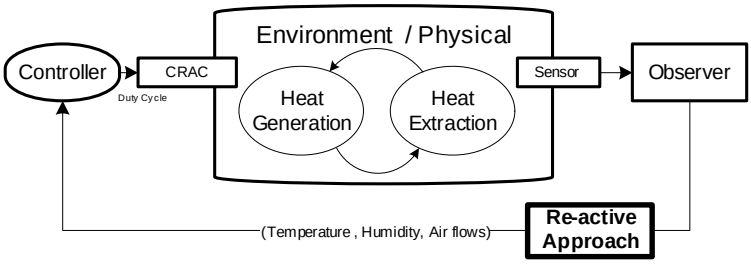

Figure 3.8: Reactive approach.

based models can be implemented in two ways depending on where the temperature change is measured: i) by measuring the change in return temperature, and ii) by measuring the change in internal temperature at each blade. There are advantages and disadvantages for both approaches. Since the former approach uses only the return air temperature, it is simple to adjust the controller. However, it could have a substantial delay depending on how fast the room air is circulated and how big the machine room size is.

The advantages of the reactive approach is that we are directly observing the temperature at the place where we want to control the temperature. If the temperature of any blade rises up to the critical point that can damage the hardware components, air conditioning system reacts to cool down the system. Controller can react faster in this approach than the former one because the source of the problem is close to the air conditioning system, but it still has delay and cooling can be more expensive as only a few overheated machines can trigger an entire cooling system. Moreover, it increases the complexity of the control and communication mechanism and this approach measures the change in temperature but does not measure the quantitative heat generated. Hence, this approach is inefficient compared to proactive approach.

The reactive approach takes corrective actions after the temperature has crossed a threshold temperature $TH_{high,low}$. One possible action is to increase the fan speed, which leads to increase the flow of the air and extract more heat from the blades. As the heat extracted depends on the inlet temperature, if the temperature of the room is not low enough to cool down the machine this action is not affective. Another possible action is to increase the compressor duty cycle so to lower the temperature of the air. Within

the reactive approach, however, a fine balance between these two actions as well as the optimal tuning of the parameters controlling them are not possible as quantitative relations between causes and effects are missing. For this reason, we proposed the proactive approach, whose performance is provided in the following section.

3.5 Performance Evaluation

In this section, we analyze the performance of the models developed in this work. The simulations are built using MATLAB®. The nonlinear optimization problem in Sect. 3.3.3 is solved using the *fmincon* solver in the Optimization Toolbox of MATLAB®. In Sect. 3.5.2, we analyze the proactive approach based mathematical model using the chosen benchmarks and compare the results with that of the reactive approach. The proactive approach using multiple fans is also simulated. In Sect. 3.5.3, we compare the energy consumption and the risk of overheating for both reactive and proactive approaches.

Table 3.2 shows the benchmarks used to validate the model. The chosen workload benchmarks are compute intensive and generate substantial amount of heat. Most of the workloads that run in high performance clusters and the datacenters are compute intensive, and are comparable to the chosen workload benchmarks. Using the chosen benchmarks and input variables in Table 3.3, we simulate the thermal behavior of a datacenter and the 'heat imbalance (Q)'. We used two types of scheduling algorithms for simulations: i) random load balancing and ii) sequential load balancing. Random load balancing selects random blades and assigns workloads, whereas sequential load balancing selects sequential blades that are closely located and assigns workloads accordingly.

Table 3.2: List of benchmarks used.

Benchmark Name	**Benchmarks Type**
FFTW	Computing discrete Fourier Transforms
NAS-SP	The NASA Parallel Benchmarks (NPB) family
HPL Linpack	Solves a linear system on distributed-memory computer

The chosen benchmark workloads were run and profiled on our test server. Their processor, I/O subsystem, memory subsystem and NIC usage with respect to time were

Table 3.3: Input variables for the simulation.

Input	
Maximum airflow from the CRAC	18000 m^3/h
Supply-temperature from the CRAC	17 °C
Number of blades	1260
Number of enclosures	90
Number of racks	30
Number of perforated tiles	15
Size of a perforated tile (A_{tile})	1 m^2
Mass of the air in the datacenter	220.5 kg
Maximum fan flow	5 m^3/s
Set point	22 °C
threshold for the risk condition	70 °C
Heat generation model	HPL Linpack, FFTW, NAS-SP
Scheduling algorithm	Random load balancing, Sequential load balancing

measured using custom scripts. We obtained the time for which each subsystem was 'on' from the profiling data, and from the data sheet we know the power utilized by the subsystems and also the power dissipated. With this information, we derived the heat dissipated per unit time based on this profile data.

3.5.1 Reactive Approach

In reactive approach, either the 'duty cycle of a compressor (η)' or 'fan speed of the CRAC (ω)' can be changed based upon the user specifications. We choose the set point to react based on the temperature recommended in [76], which ranges from 64.4 (18) to 80.6 (27). We change the fan speed or duty cycle of the compressor to adjust temperature set point within the recommended range, which is bounded by TH_{high} and TH_{low} as shown in Figs. 3.9 and 3.10. Since controller does not have knowledge about how much heat will be generated in the future, it can only adjust fan speed or duty cycle based on the temperature variations in this approach.

The reactive approach takes a corrective action after the temperature has crossed a set point temperature. Note that this set point temperature can be different based on the response time of the control system used. As a reaction, duty cycle of the compressor can vary based on the changes in returning air temperature as shown in Fig. 3.9. If returning temperature increases above the set point, then the controller increases duty cycle of the compressor accordingly to extract more heat and hence lower the temperature. Fan speed is fixed at its maximum value which is 5 m^3/s. However,

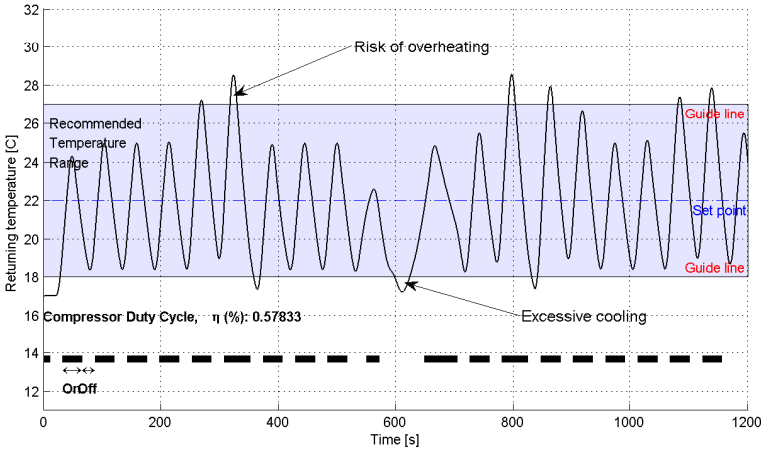

Figure 3.9: Reactive Approach: Compressor cycle change vs. returning temperature of the CRAC; Positive heat imbalance creates 'risk of overheating' and negative heat imbalance shows 'excessive cooling' when fan speed ω is fixed as 5 m^3/s.

the 'risk of overheating,' which refers to the amount of heat that can damage internal components, exists because of the non-instantaneous cooling effect on the blades. Also, because of the delayed reaction of the cooling system, excessive amount of heat may be extracted (excessive cooling) bringing the temperature below the guide line.

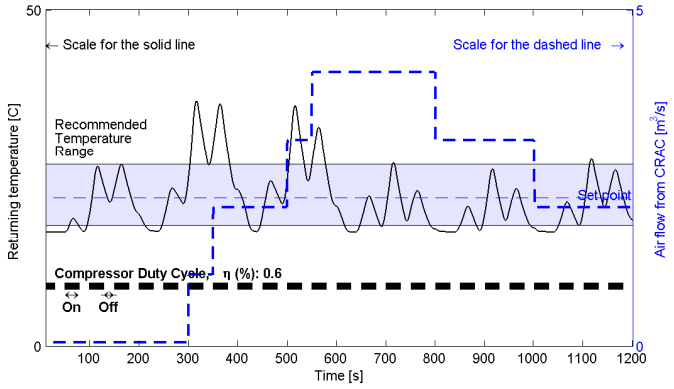

Figure 3.10: Reactive Approach: Fan speed change vs. returning temperature of the CRAC; 'risk of overheating' still exist even if the CRAC fan speed increases. Duty cycle of the air compressor η is fixed as 0.6.

Changing the fan speed to adjust the temperature is the alternative way to control the temperature in reactive approach. Figure 3.10 shows CRAC fan speed control and corresponding returning temperature when compressor cycle is fixed at 0.6 to show the

effect of fan in this case. Fan speed is controlled by the controller as the temperature increases above the set point or decreases below the set point, but the 'risk of overheating' and 'excessive cooling' still exists because of the delay.

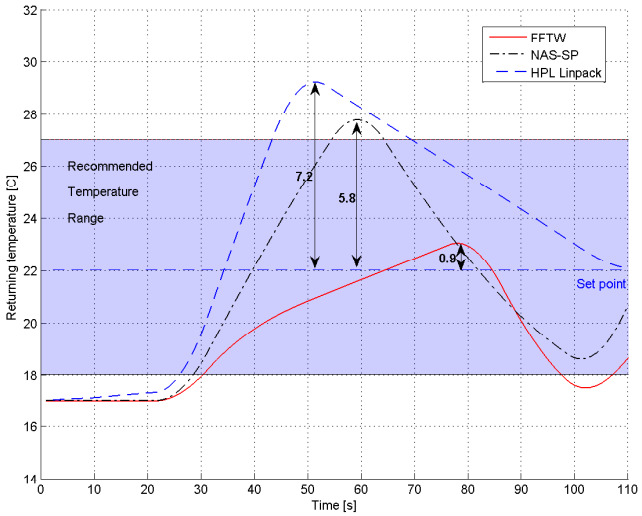

Figure 3.11: Temperature variations for FFTW, NAS-SP and HPL workloads.

We observed the temperature variations for different workloads such as FFTW, NAS-SP, and HPL. From Fig. 3.11, we see that the temperature change for different workload is different. From this, we infer that the temperature variation is dependent on the workload and each one of them has a different heat pattern. In Fig. 3.11, it is easy to control the temperature for FFTW workload as it generates less heat compare to the other two. In the case of HPL it is difficult to control the temperature because it generates more heat and the temperature rises quickly above the recommended temperature range.

In Fig. 3.12, we show the temperature variations based on the time delay which is dependent on the distance of the blade from the CRAC unit. As the heat propagation is dependent on the distance of the heat sensors from the source of the heat generation, there is a delay in detection of the rise in temperature at the CRAC unit. Hence, there is a delay in the response of the CRAC unit to control the temperature at the source of heat generation. In Fig. 3.12, we see that higher the delay, more is the chance of temperature rising above recommended temperature range.

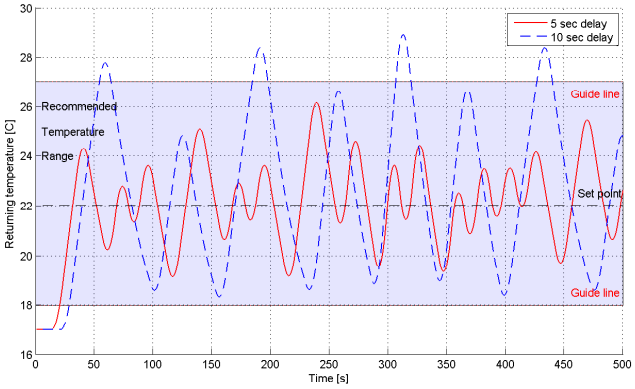

Figure 3.12: Temperature variations based on the time delay which is dependent on the distance of the blade from the CRAC unit.

3.5.2 Proactive Approach

In proactive approach, 'duty cycle of a compressor (η)' or 'fan speed of the CRAC (ω)' can be jointly optimized upon the heat estimation model provided in Section 3.4. Optimization problem is solved every 50s, which is time window size of air compressor duty cycle ($T_{ON} + T_{OFF}$) to adjust η and ω. This model is evaluated for the three chosen benchmarks. Proactive approach is intrinsically predictive in nature as it estimates the heat that will be generated in the future and adjusts CRAC unit accordingly. This way, we can prevent 'risk of overheating' and 'excessive cooling' by eliminating the delay in cooling action. In Fig. 3.13, we plot the percentage of server blade utilization with respect to time. It provides us the information about the workload intensity at a particular time instance. This server utilization rate is same for all the other workload simulations.

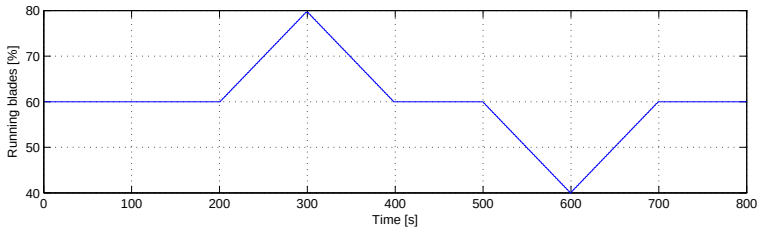

Figure 3.13: Utilization rate of server blades [%].

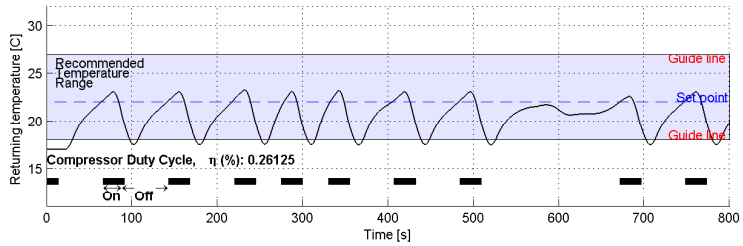

(a)

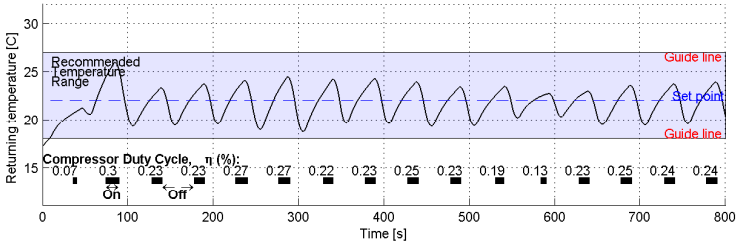

(b)

Figure 3.14: Temperature variation of datacenter, FFTW workload, Random load balancing for the job distribution (a) Reactive approach; (b) Proactive approach.

Figures 3.14(a) and 3.14(b) show the temperature changes based on FFTW workload subsystem usage profile. Workloads are assigned to the blades by random load balancing algorithm. Compressor cycle in reactive approach in Fig. 3.14(a) decreases when the temperature crosses the set point. This workload shows moderate temperature change because this workload does not generate much heat compared to the other two workloads. In this case, 'excessive cooling' appears due to non-instantaneous action of air conditioning system. Figure 3.14(b) shows that temperature remains around the 'set point,' and in turn saves energy by optimizing fan speed and compressor cycle. Energy consumption is compared in Section 3.5.3.

Figures 3.15(a) and 3.15(b) show the temperature change based on NAS-SP workload subsystem usage profile. In Fig. 3.15(a), we observe a periodic 'risk of overheating' during the time of simulation due to delay in cooling. This approach causes almost 10 second delay because of the distance of the blade from the CRAC unit. On the contrary, proactive approach estimates the heat to be generated, and optimizes fan speed and compressor cycle based on the estimation by the model proposed in Sect. 3.4. Figure 3.15(b) shows that temperature varies mostly within the recommended temperature

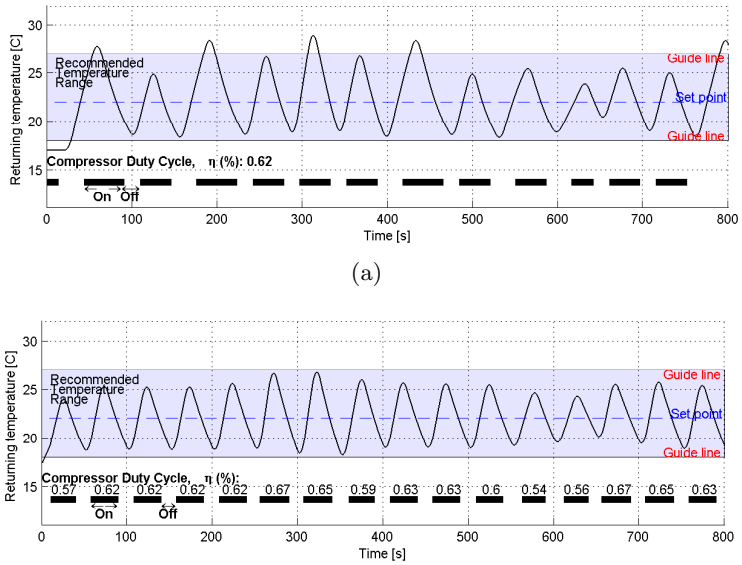

Figure 3.15: Temperature variation of datacenter, NAS-SP workload, Random load balancing for the job distribution (a) Reactive approach; (b) Proactive approach.

range.

Figures 3.16(a) and 3.16(b) show the temperature change based on HPL-Linpack workload subsystem usage profile that has a highest heat generation among the three benchmarks. Since reactive approach (Fig. 3.16(a)) does not have any knowledge about the workload, the controller increases the compressor cycle and fan speed according to increase in temperature. A high 'risk of overheating' appears between 250 to 350 seconds because heat imbalance (Q) is positive and very high. The compressor is too slow to react to the temperature rise because HPL generates more heat as compared to other workloads. Also, from Fig. 3.13, we see that the highest blades utilization occurs in this same period of time. Due to this, there is extreme load on the CRAC and hence, the slow reaction. On the contrary, proactive approach optimizes fan speed and compressor cycle based on the estimation of heat to be generated. Figure 3.16(b) shows that temperature varies mostly within the recommended temperature range.

The heat generation in a datacenter is dependent on workload distribution, with uneven distribution of the workload there is uneven heat generation. With a global control system and a single fan, we cannot partially extract the heat. Ideally one fan

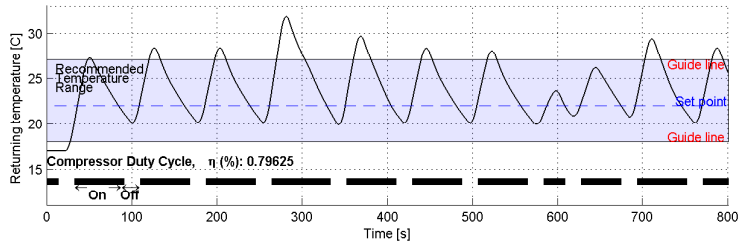

(a)

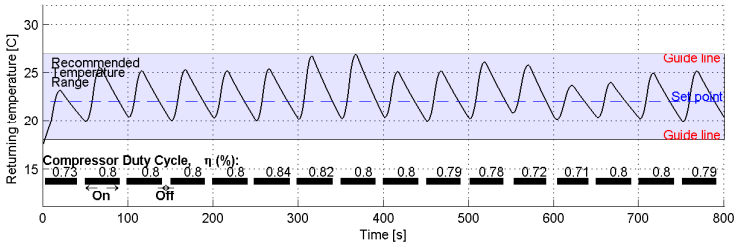

(b)

Figure 3.16: Temperature variation of datacenter, HPL Linpack workload, Random load balancing for the job distribution (a) Reactive approach; (b) Pro-active approach.

can extract the same amount of heat from each blade, but cannot extract heat from a selected blade, enclosure or a rack. Moreover, selective cooling using multiple fans is needed for an intensive workload since, it can causes a high heat imbalance and isolated hot spots. A selective air conditioning system control using multiple fans (one per 'corridor' j) is an energy efficient alternative, since we increase the fan speed only when needed depending on uneven heat generation.

3.5.3 Energy consumption and overheating risk

We estimate the 'energy consumption' for cooling systems and the 'risk' of overheating for the hardware components, which is the input to both reactive approach and proactive approach models. Energy consumption is the sum of energy consumed by the fan and the air compressor during the time for which the simulation runs. The energy is represented in the units of 'kWh', which can be directly converted into kJ, multiplying by a factor of 3600. 'Risk' refers to percentage of overheating risk of the hardware components, calculated by averaging the percentage of blades over the threshold which is set as 70elsius in this simulation. In this way, we can show what percentage of blades

Table 3.4: Risk of overheating and Energy consumption for random load balancing.

Workload	Reactive				Proactive			
	Return temperature (Global Temperature)		Internal temperature (Local Temperature)		Single fan (Global control)		Multiple fans (Local control)	
	Risk (%)	Energy (kWh)	Risk	Energy	Risk	Energy	Risk	Energy
FFTW	0.00	7.61	0.00	7.92	0.00	8.23	0.00	7.27
NAS	3.27	14.84	0.58	15.45	0.03	15.03	0.05	14.22
HPL	15.55	18.12	16.35	20.09	3.10	18.515	3.71	16.19

Table 3.5: Risk of overheating and Energy consumption for sequencial load balancing.

Workload	Reactive				Proactive			
	Return temperature (Global Temperature)		Internal temperature (Local Temperature)		Single fan (Global control)		Multiple fans (Local control)	
	Risk (%)	Energy (kWh)	Risk	Energy	Risk	Energy	Risk	Energy
FFTW	0.00	7.62	0.00	7.93	0.00	7.99	0.00	7.12
NAS	2.74	13.52	0.36	14.57	0.28	14.24	0.00	13.83
HPL	14.65	18.74	11.02	19.15	2.55	20.95	2.44	17.57

are under the state of overheating risk in a datacenter.

Tables 3.4 and 3.5 show the results for random load balancing and sequential load balancing respectively. Simulations are performed by using parameters in Table 3.3. Table 3.5 shows lower energy consumption than Table 3.4 because in Table 3.4 workloads are evenly distributed in an orderly manner so that the heat generation is uniform.

We compare the proactive and reactive approach based on their energy consumption and 'risk' factor. In reactive approach, using return temperature (global temperature measurement) for activating CRAC unit shows lower energy consumption than using internal temperature (local temperature measurement) because it only uses mixed returning air temperature that averages heat imbalance of all the blades in a datacenter. However, using the internal temperature of the blade activates CRAC unit whenever the temperature at any blade crosses threshold, and therefore prevents the 'risk' of overheating, but while avoiding this risk consumes more energy compared to the approach that uses return temperature (global temperature measurement).

The proactive approach does not have big 'risk' compared to the reactive approach since the controller can quantify and extract heat before it creates heat imbalance. However, the proactive approach is inherently based on localized estimation of heat generated at each blade and hence, it consumes more energy compare to the reactive approach using single fan in tables 3.4 and 3.5.

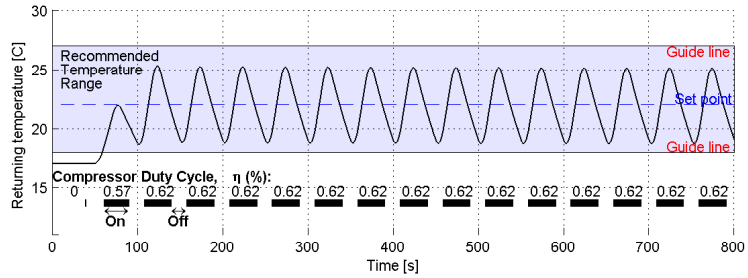

(a)

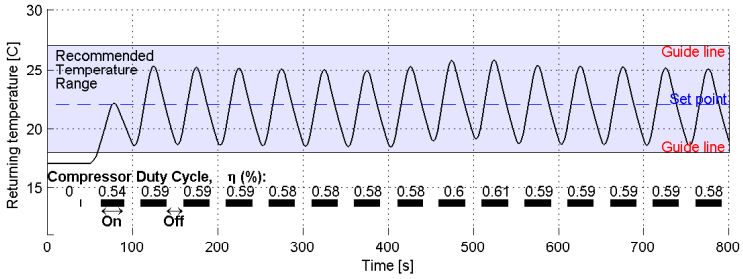

(b)

Figure 3.17: Temperature variation of the datacenter, NAS-SP workload, Random load balancing for the job distribution (a) Proactive approach with single fan; (b) Proactive approach with multiple fan.

In proactive approach, using multiple fans consumes less energy than using single fan or reactive approach. Even though multiple fans require additional power to operate which is few kilowatts, they can adjust the airflows to different aisles and selectively extract heat that causes the 'risk' of overheating efficiently. Since, proactive approach is by default based on localized estimation, multiple fans help remove this heat in a localized manner, which results in increase in the energy efficiency. Difference in energy consumption between using multiple fans and single fan is more apparent in sequential load balancing in Table 3.5. Using multiple fans (Fig. 3.17(b)) achieves lower energy consumption than using single fan (Fig. 3.17(a)), showing that temperature changes in lower range using lower compressor duty cycle.

3.6 Discussion

In this chapter we proposed a proactive approach based optimization model for cooling systems of the datacenters. The proactive approach is based on having advanced knowledge of the workload behavior and taking an appropriate action before the heat

imbalance affects the temperature. When we compared the proactive approach with the reactive approach, reactive approach was found to have many disadvantages such as delayed response, high risk of over heating, excessive cooling and recursive cycling. Proactive approach proposed in this chapter overcomes these disadvantages of the reactive approach. Proactive approach cools the system before temperature rises and prevents any occurrence of 'risk of overheating' and also prevents 'excessive cooling' as the heat imbalance is estimated based on the knowledge of subsystem usage and the workloads. For it to be effective there is a need for multiples fans under each plenum for each aisle or each corridor, which is possible with minimal or no changes to the designs of the existing datacenters. Multiple fans help in effectively controlling hot spots occurring in different locations, which cannot be eliminated by the current single fan cooling systems even by using proactive approach. The use of multiple fans in the proactive approach to control the cooling system saves approximately 4% to 10% of the energy required for cooling, depending on the workload and scheduling algorithms implemented.

Proactive approach suggested in this chapter can be optimized if it is implemented with supporting infrastructure such as external temperature and humidity sensors and airflow meters that are crucial in obtaining inlet and outlet temperatures, humidity and airflow for accurate air circulation modeling in (3.8). Thermal cameras could be another option to micro-managing heat imbalance and hot spots. Thermal cameras can be used to detect, characterize, localize, and track hot spots causing heat imbalances. We introduce implementing proactive approach using the thermal infrastructure in the next chapter.

Chapter 4

VMAP: Proactive Thermal-aware Virtual Machine Allocation in HPC Cloud Datacenters

4.1 Overview

The trend: Datacenters are a growing component of society's IT infrastructure and their energy consumption surpassed 237 billion kWh/year worldwide and 76 billion kWh/year in the US in 2010 [1]. Even though these numbers are lower than what the US Environmental Protection Agency predicted in 2007 [4], they correspond to 6% and 2% of the total electricity usage in the US. The impact of this proliferation of datacenters on the environment and society includes increase in CO_2 emissions, overload of the electricity supply grid, and rise in water usage for cooling leading to water scarcity [5]. The scale and complexity of datacenters are growing at an alarming rate and their management is rapidly exceeding human ability, making *autonomic* (self-configuration, self-optimization, self-healing, and self-protection) management approaches essential.

High-Performance Computing (HPC) applications are resource-intensive scientific workflow (in terms of data, computation, and communication) that have typically targeted Grids and conventional HPC platforms like super-computing clusters. *Clouds* – composed of one or more virtualized datacenters providing the abstraction of nearly-unlimited computing resources through the elastic use of federated resource pools – are being increasingly considered to enable traditional HPC applications. However, maximizing energy efficiency and utilization of cloud datacenter resources, avoiding undesired thermal hotspots (due to overheating of over-utilized computing equipment), and ensuring Quality of Service (QoS) guarantees for HPC applications are all conflicting objectives, which require joint consideration of multiple pairwise tradeoffs.

Need for thermal awareness: From our feasibility study and proof-of-concept

experiments conducted at our machine room in the NSF Cloud and Autonomic Computing Center (CAC), Rutgers University, we have inferred that one of the fundamental problems in HPC-cloud datacenters is the *local unevenness in heat-generation* and *heat-extraction rates*: the former can be attributed to the non-uniform distribution of workloads (of different types and intensities) among servers and to the heterogeneity of computing hardware; the latter can be attributed to the non-ideal air circulation, which depends on the layout of server racks inside the datacenter and on the placement of Computer Room Air Conditioning (CRAC) unit fans and air vents. The heat-generation and -extraction rates may differ, which over time causes *heat imbalance*. This heat imbalance will be large if the rates are significantly different from each other or if their difference prolongs over extended time periods.

A large negative heat imbalance at a particular region inside a datacenter will result in energy-inefficient overcooling and, hence, in a significant decrease in temperature. Conversely, a large positive heat imbalance will lead to a significant temperature increase, which may result in undesired thermal hotspots and server operation in the unsafe temperature range. Thus, *thermal awareness,* which is the knowledge of heat imbalance at different regions inside a datacenter, is essential to maximize energy and cooling efficiency as well as to minimize server system failure rate. Our novel concept of heat imbalance enables *proactive* datacenter management decisions (such as resource provisioning, cooling system optimization) through prediction of future temperature trends as opposed to the state-of-the-art *reactive* management decisions based on current temperature measurements.

Our contributions: In virtualized HPC datacenters, one or more Virtual Machines (VMs) are created for every application request (with one or more workloads) and each VM is provisioned with resources that satisfy the application QoS requirements, which are based on Service Level Agreements (SLAs). Once VMs are provisioned, they have to be allocated to servers. We propose a novel thermal-aware proactive VM consolidation solution referred to as VMAP. The benefit of employing VMAP is three-fold: i) energy spent on computation can be saved by turning off the unused servers after workload (or VM) consolidation; ii) the utilization of servers that are in the "better cooled"

areas of the datacenters (with high heat extraction) can be maximized; iii) heat can be extracted more efficiently (by doing a lower amount of work) by the CRAC system from the consolidated server aisles, which are hotter than non-consolidated server aisles. Note that (iii) is possible due to the fact that the efficiency of heat extraction increases with increase in return-air temperature.

Furthermore, capability to migrate VMs opens up more opportunities to save energy as we can perform better resource consolidation over time. However, migrations incur overheads in terms of delays in the applications running within the VMs and energy consumption. We propose VMAP+, which possesses all of VMAP's capabilities and in addition leverages thermal-aware VM migrations in order to achieve greater resource consolidation over time while taking into account the following key tradeoffs: energy savings vs. delay and network overhead. VM migration involves transferring the memory of the VM from the host to the destination server. The time taken and, hence, the energy consumption for VM migration depends on the network bandwidth. The higher the network bandwidth the shorter the migration time and smaller the energy footprint. Excessive migrations can cause large delays and high energy consumption compared to the no-migration case. VMAP+ takes the aforementioned factors into consideration by migrating VMs only i) if migration does not result in excessive heat generation, ii) if energy saving is greater than the migration cost, and iii) if migration delay is within the SLA.

Potential benefits: Both VMAP and VMAP+ also exploit the heterogeneity in the cloud infrastructure (federated datacenters) – in terms of electricity cost, hardware capabilities (CPU, memory, disk I/O, and network subsystems), tunable parameters of the CRAC system, and local regulations (governing CO_2 emission and water usage) – to maximize energy efficiency. Our solutions are aimed at increasing the energy and cooling efficiency and at decreasing equipment failure rates so to minimize both the impact on the environment and the Total Cost of Ownership (TCO) of datacenters. VMAP can significantly contribute to energy efficiency (9%, 9%, and 35% average reduction in energy consumption compared to the traditional temperature-based reactive thermal

management schemes: first-fit-decreasing, best-fit-decreasing, and "cool-job" [3] allocation, respectively) while not violating recommended operating temperature range. VMAP+ achieves a further 3% average reduction in energy consumption compared to VMAP. The following are the main contributions of our work.

- We introduce the novel notion of heat imbalance and validate a simple yet robust heat-imbalance model, which helps predict future temperature trends and make proactive resource provisioning and migration decisions;
- We propose a proactive thermal-aware VM consolidation solution, VMAP, which minimizes energy consumption for computation, increases resource utilization, and improves efficiency of cooling;
- We propose VMAP+, which extends VMAP's capabilities by leveraging thermal-aware VM migrations so to achieve greater resource consolidation over time while taking into account the following key tradeoffs: energy savings vs. delay and network overhead;
- We validate our proposed approach through extensive experiments – in a single-datacenter as well as in federated-datacenters (at different sites of NSF CAC, Rutgers University and University of Florida).

The remainder of this chapter is organized as follows: in Sect. 4.2, we discuss the state of the art in autonomic thermal-aware management of datacenters; in Sect. 4.3, we outline our broader vision for thermal-aware autonomic datacenter management, present details on the design and validation of our heat-imbalance model (Sect. 4.3.1), and describe VMAP and VMAP+ (Sect. 4.3.2); in Sect. 4.4, we study the performance of VMAP and VMAP+ using experiments and simulations; and finally, in Sect. 4.5, we present our conclusions.

4.2 Related Work

Prior research efforts on thermal management of datacenters [42] have focused exclusively on only one of the two fundamental approaches: management of heat extraction [31] or management of heat generation inside a datacenter [47, 50]. The first approach aims at improving cooling system efficiency by effectively distributing cold air inside the datacenter (cooling system optimization), while the second approach focuses on how to balance or migrate workloads in such a way as to avoid overheating of computing equipment. In contrast, we focus on a joint approach so to minimize the risk of overheating of servers while simultaneously maximizing the cooling efficiency.

In [43], the authors profile and benchmark the energy usage of 22 datacenters. They perform energy benchmarking using a metric that compares energy used for IT equipment to the energy used for the CRAC system and conclude that the key to energy efficiency is air circulation management (for effective and efficient cooling). As many datacenters employ raised floors with perforated tiles to distribute the chilled air to racks, researchers have tried to gain valuable insights into efficient airflow distribution strategies in such datacenter layouts [31]. Other research efforts were aimed at improving the efficiency of cooling systems through thermal profiling (knowledge of air and heat circulation) of datacenters. Basic mathematical modeling and parameters for profiling datacenter are proposed in [39]. However, capturing complex thermodynamic phenomena using complex Computational Fluid Dynamic (CFD) models [25] is prohibitive in terms of computational overhead. Measurements from scalar sensors alone [37] cannot capture the complex thermodynamic phenomena inside a datacenter. Hence, we used a heterogeneous sensing infrastructure [44] – composed of temperature and humidity scalar sensors, thermal cameras, and air flow meters – to thermally profile datacenters in space and time so to exploit that information for resource provisioning and cooling system optimization.

Several solutions that employ temperature-aware job distribution and migration have been proposed for alleviating undesired thermal behavior (higher operating temperatures) inside datacenters. Moore et al. [47] proposed thermal management solutions

that focus on scheduling workloads considering temperature measurements. They designed a machine-learning-based method to infer a model of thermal behavior of the datacenter online and to reconfigure automatically the thermal load management systems for improving cooling efficiency and energy consumption. Bash and Forman [49] developed a policy to place the workload in areas of a datacenter that are easier to cool, which results in cooling power savings. They used scalar temperature sensor measurements alone to derive two metrics that help decide whether to place workload on a server or not: the first metric, Thermal Correlation Index (TCI), gives the efficiency with which any given CRAC can provide cooling resources to any given server; while the second is Local Workload Placement Index (LWPI). Tang et al. [50] investigated the mechanism to distribute incoming tasks among the servers in order to maximize cooling efficiency while still operating within safe temperature regions. They developed a linear, low-complexity process model to predict the equipment inlet temperatures in a datacenter given a server utilization vector; they mathematically formalize the problem of minimizing the datacenter cooling cost as the problem of minimizing the maximal (peak) inlet temperature through task assignment. However, the work was validated only through simulations. In [51], the authors explore a spatio-temporal thermal-aware job scheduling as an extension to spatial thermal-aware solutions like [47, 48, 50].

Heath et al. [48] propose emulation tools ('Mercury' and 'Freon') for investigating the thermal implications of power management. In [67], the authors present 'C-Oracle', a software infrastructure that dynamically predicts the temperature and performance impact of different thermal management reactions (such as load redistribution and dynamic voltage and frequency scaling) into the future, allowing the thermal management policy to select the best reaction. However, neither of the aforementioned thermal-aware workload placement solutions explicitly take into account the direct impact of workload distribution on cooling system efficiency and vice-versa. Thermal-aware management of datacenters should strive to minimize the TCO of datacenters, i.e., to minimize the cost of running servers through energy-aware workload distribution as well as to minimize the energy spent on cooling, by thoroughly understanding the effect of one on the other. That is why we combined thermodynamic models and real-time measurements

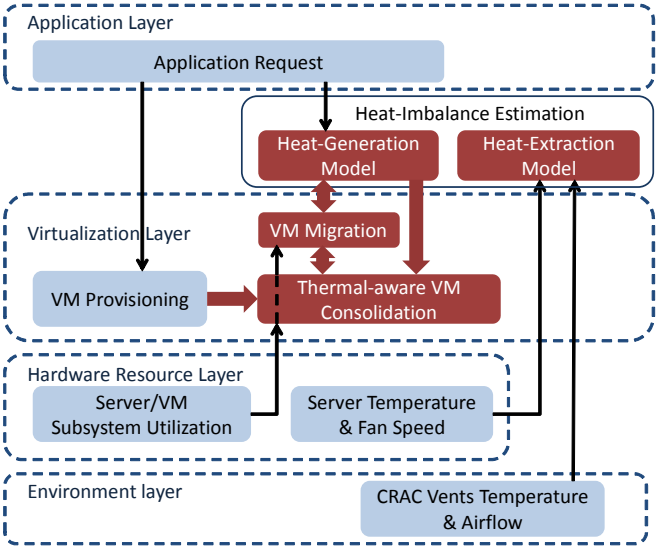

Figure 4.1: Envisioned cross-layer approach to autonomic management of virtualized HPC datacenters. The main focus of this chapter is indicated in red boxes.

(from temperature and humidity scalar sensors as well as air flow meters) to capture the complex thermodynamic phenomena of heat generation (due to specific workload distribution) and heat extraction (due to cooling system parameters and characteristics), in order to predict the future temperature map of the datacenter for enabling proactive thermal-aware datacenter management decisions.

4.3 Proposed Approach

We propose a proactive *cross-layer* approach to autonomic datacenter management, which is *information centric* and requires continuous processing and analysis of real-time feedback from multiple layers of abstraction (as depicted in Fig. 4.1). The *application layer* provides information regarding the applications' (and, hence, the workloads') characteristics such as their computing resource requirements, energy consumption, and performance on different hardware platforms. Modern blade servers (*hardware resource layer*) are equipped with a number of internal sensors that provide information about server fan speed and subsystem operating temperatures as well as utilization. However, information extracted from the application and hardware resource layers alone cannot capture the complex thermodynamic phenomena of heat and air circulation inside a datacenter.

Information from the *environment layer,* comprising of an heterogeneous sensing infrastructure (with scalar temperature and humidity sensors, thermal cameras, and airflow meters) is key to characterize the thermal behavior of a datacenter under a given load (information from the *application layer*) [44]. As mentioned earlier, the estimation of heat imbalance requires estimation of the heat-generation and heat-extraction rates. The heat-generation model exploits the information provided by the application layer while the heat-extraction model leverages information provided by the environment as well as hardware resource layers. The *virtualization layer* – which provisions, allocates, and manages VMs (created based on application requests) – exploits the knowledge of heat imbalance to predict future temperature trends for optimal resource allocation in datacenters. In this chapter, we focus on the design and validation of the heat-imbalance model and on how the knowledge of heat imbalance can be exploited to perform energy-efficient proactive VM consolidation in datacenters (shown in red boxes in Fig. 4.1).

While proactive VM consolidation (involving allocation and migration) has several clear advantages, namely, reduced energy cost for computation (through high utilization of fewer computing resources) as well as for cooling (through better heat extraction at higher operating temperatures), it has certain drawbacks. Increased utilization of servers results in continuous operation of computing hardware at temperatures close to the upper bound of the recommended operating temperature range. This, however, is not a major concern due to the following reasons: i) manufacturers usually provide a conservative upper bound for the recommended operating temperature range; ii) our consolidation solution is thermal-aware and does not let the operating temperatures go beyond the recommended range (referred to as *thermal violation*) unlike other temperature-agnostic solutions; iii) the frequency of equipment upgrades (due to tremendous rate of innovation in computing hardware) is much higher than the rate of replacement due to failures.

Another drawback of traditional server consolidation is violation in SLAs (in terms of application runtime) due to greater resource contention at higher utilization levels. However, this is not a concern for virtualized HPC clouds as i) users are guaranteed

the resources they specifically ask for, ii) VMs are isolated from each other, and iii) we do not multiplex resources, i.e., the total subsystem utilization of all VMs in a server will not exceed the total subsystem capacity of that server. In our prior work [12, 66], we have shown through simulations that heat-imbalance-based proactive datacenter management (cooling system optimization) is superior in terms of energy efficiency and minimization of risk of equipment failures compared to its conventional temperature-measurement-based reactive counterpart. Our envisioned approach represents a transformative shift towards cross-layer autonomics for datacenter management problems, which have so far been considered mostly in terms of individual layers. In the following, we first focus on our novel heat-imbalance model, which incorporates information from the application, hardware resource, and environment layers. We then present our heat-imbalance-based proactive VM allocation and migration solutions, which resides in the virtualization layer.

4.3.1 Heat-Imbalance Model

A VM is created for every application request and is provisioned with resources (CPUs, memory, disk, and network capacity) that satisfy the application's QoS (usually deadline) requirements. Without any loss of generality, we assume that this provisioning has already been performed using techniques such as the ones described in [89]. The provisioned VMs now have to be allocated to physical servers housed within racks in datacenters. Let $\mathcal{M}$ be the set of VMs to be allocated and $\mathcal{N}$ be the set of servers. An associativity binary matrix $\mathbf{A} = \{a_{mn}\}$ (with $a_{mn} \in \{0, 1\}$) specifies whether VM m is hosted at server n or not. A VM m is specified as a vector $\mathbf{\Gamma}_m = \{\gamma_m^s\}$, where $s \in \mathcal{S} = \{CPU, MEM, IO, NET\}$ refers to the server subsystems and γ_m^s's are the VM subsystem requirements (e.g., CPU cores, amount of volatile memory [MB], disk storage space [MB], network capacity [Mbps]).

Representation (or *mapping*) of a VM's subsystem requirement (γ_m^s) as a factor of physical server subsystem capacity is straightforward if all the servers of the datacenter are assumed to be homogeneous. For example, a VM m requiring 4 virtual CPUs, 2 GB of RAM, 64 GB of hard-disk space, and 100 Mbps network capacity can be

represented as $\mathbf{\Gamma_m} = \{0.25, 0.125, 0.125, 0.1\}$ if all the servers in a datacenter have 16 CPU cores, 16 GB of RAM, 512 GB of local hard-disk space, and a gigabyte ethernet interface. However, homogeneity is rarely the case as datacenters usually have a few different generations of each subsystem, for example, CPUs with different clock rates and number of cores (1.6/2.0/2.4 GHz and 4/8/16/32 cores), different generations or sizes of RAMs (SDR/DDR SDRAM or sizes ranging from 4 to 32 GB), network switches of varying capacities (0.1, 1, or 10 Gbps), etc. The mapping problem becomes non trivial in an heterogeneous environment. However, assuming that only a small finite number of generations of each subsystem are present in the datacenter, we create such a mapping for each generation of every subsystem.

Estimation of heat imbalance: We formulate the heat-imbalance model in a datacenter based on heat-generation and heat-extraction rates as follows,

$$\Delta I_n = \int_{t_0}^{t_0+\delta} (h_n - q_n)dt = M_n \cdot C \cdot \Delta T^n_{[t_0,t_0+\delta]}, \tag{4.1}$$

where ΔI_n [J] denotes the heat imbalance of CPU inside server n during the time between t_0 and $t_0 + \delta$, and M_n and C denote the mass and specific heat capacity, respectively, of the CPU. Note that if ΔI_n is positive (i.e., $h_n > q_n$), the temperature of the CPU at server n increases in the time interval $[t_0, t_0 + \delta]$ (hence, $\Delta T^n > 0$); conversely, if ΔI_n is negative (i.e., $h_n < q_n$), the temperature of the CPU at server n decreases (hence, $\Delta T^n < 0$).

This estimated heat imbalance helps us predict the increase or decrease in temperature, given by ΔT^n, to take management decisions such as VM placement, VM migration, and cooling system optimization.

Validation of the proposed models: Certain parameters in the proposed heat-imbalance model are determined empirically as they cannot be obtained directly (e.g., from server specification documents). The heat dissipation factor α in (1) is one of the key parameters that is determined empirically. Similarly, the server outlet temperature T_{out} in (2) varies with time and is a function of CPU temperature, which is what the heat-imbalance model is designed to estimate. Hence, the relationship between

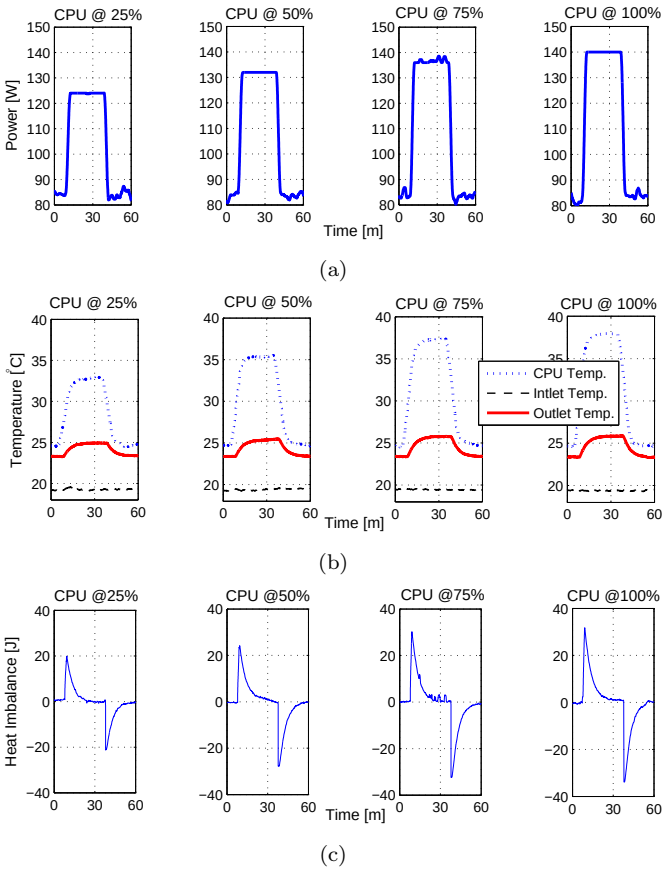

Figure 4.2: Empirical data collected from servers at RU – (a) power consumption, (b) CPU and external (inlet and outlet) temperatures, and (c) calculated heat imbalance – when a representative CPU-intensive workload is run at different CPU utilization levels.

T^{out} and ΔT is determined empirically (assuming T^{in} is known and is constant in the time interval $[t_0, t_0 + \delta]$) and is substituted in the heat-imbalance model so to eliminate an extra unknown. We performed simple experiments (measurements shown in Fig. 4.2) to obtain α, to derive the relationship between T^{out} and ΔT, and to validate the resulting heat-imbalance model by comparing its output (predicted increase in the CPU temperature ΔT at a server) with actual observation (shown in Fig. 4.4).

We started from an initial idle condition, with 0% CPU utilization and a corresponding zero heat imbalance, and increased the CPU utilization from 0% to 25%, 50%, 75%, and 100% progressively as shown in Fig. 4.2. The CPU was subject to each of the aforementioned load levels for around 60 minutes so to allow the CPU temperature to reach steady state. To increase the CPU utilization we used *Lookbusy* (a synthetic load

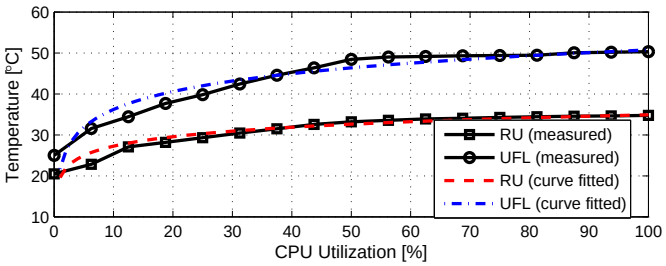

Figure 4.3: Relationship between ΔT and CPU utilization using data from both RU and UFL servers.

generator for Linux systems), which keep the CPU(s) at the chosen utilization level by adjusting its own load up or down to compensate for other loads on the system. We measured the corresponding increase in power consumption (Fig. 4.2(a)) as well as CPU and server outlet temperatures (Fig. 4.2(b)), and also calculated the variation in heat imbalance over time (Fig. 4.2(c)). Obtaining the value of α using (3) is now straightforward as the heat imbalance, heat extraction, and power consumption are known. On the contrary, deriving the relationship between T^{out} and ΔT is non-trivial.

First, we use logarithmic regression equations to model the relationship between CPU utilization ($u_n\%$) and the increase in CPU temperature (ΔT_n $^\circ C$, shown in Fig. 4.3), i.e., $\Delta T_n = \alpha \ln(u_n) + \beta$. Then, based on this knowledge and our observation from Fig. 4.2(c), we derive a simple linear regression model that represents the relationship between ΔT_n and T_n^{out} (for a fixed server inlet temperature and airflow rate) for use in our heat-imbalance model. We verify the accuracy of the logarithmic regression equations with the empirically determined coefficients (α and β) as well as the linear regression model by repeating the aforementioned experiment again and comparing the predicted CPU temperatures over time with the actual CPU operating temperature as shown in Fig. 4.4. Prediction of future CPU operating temperatures using our heat-imbalance model is sensitive to the variable heat and air circulation patterns (thermodynamic phenomena) at different regions inside a datacenter.

4.3.2 Thermal-aware VM Consolidation

For a given set of VMs, minimizing the number of servers that are in operation (consolidation) will help reduce the energy overhead and, hence, the total energy consumption.

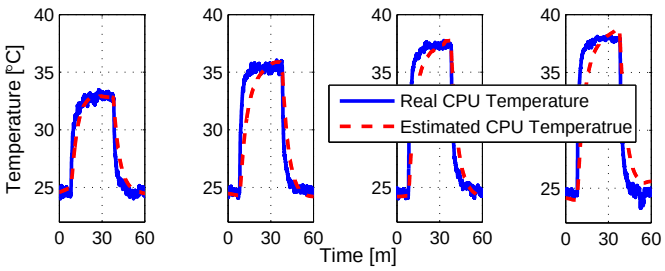

Figure 4.4: CPU temperature – measured and estimated (using the heat-imbalance model) – when a representative CPU-intensive workload is run at different CPU utilization levels.

In addition to saving the energy spent on computation, thermal-aware VM consolidation also helps achieve a higher COP of cooling. In this section, we first formulate the VM allocation problem as an optimization problem, which employs our heat-imbalance model. As this optimization is NP-hard, we then present our heuristic solution, VMAP (thermal-aware proactive VM mapping solution). The motivation for formulating the optimization problem is to gain insight and make key design decisions for our heuristic solution.

Optimization Problem: The total energy consumption in a datacenter can be split into energy consumption for *computing* (E^{comp} [kWh]), i.e., for running the workloads (or VMs) on servers, and energy consumption for *cooling* (E^{cool} [kWh]). We assume that the cooling system parameters (fan speed and compressor duty cycle of the CRAC) are fixed, i.e., the energy spent on cooling is fixed ($E^{cool} = const$) for the duration δ. Note that E^{cool} can be optimized independently at a periodicity $\Delta \gg \delta$. The goal is to find an *optimal* mapping of VMs to physical servers (represented by the binary associativity matrix $\mathbf{A}$) so to minimize E^{comp} while simultaneously increasing COP of cooling. The known (given as well as measured) parameters and optimization variables of the optimization problem can be summarized as,

$$\begin{aligned}
&\textbf{Given (offline)} && : \mathcal{N}, T^{reco}, \delta, M_n, C_p; \\
&\textbf{Given (online)} && : \mathcal{M}, \mathbf{\Gamma}_m \ \forall m \in \mathcal{M}; \\
&\textbf{Measured (online)} && : T_n^{t_0}, m_n^{in}, T_n^{in}, \mathbf{\Lambda}_n \ \forall n \in \mathcal{N}; \\
&\textbf{Find} && : \mathbf{A} = \{a_{mn}\}, \ m \in \mathcal{M}, n \in \mathcal{N}.
\end{aligned} \tag{4.2}$$

Here, $T_n^{t_0}$ and $\mathbf{\Lambda}_n = \{\lambda_n^s\}$ represent the current CPU temperature and the maximum residual capacity of each subsystem s at server n, respectively. The objective of the optimization problem is,

$$\textbf{Minimize}: \; E^{comp} = \sum_{n \in \mathcal{N}} E_n^{comp}, \tag{4.3}$$

$$E_n^{comp} = \sum_{s \in \mathcal{S}} \left(P_n^{s,on} \cdot t_n^{s,on} + P_n^{s,idle} \cdot t_n^{s,idle} \right) \cdot \alpha_n^s; \tag{4.4}$$

$$\textbf{Subject to}: \; \textbf{C1, C2, C3.}$$

The first constraint (**C1**) ensures that a VM is allocated to *one* and *only one* server, i.e.,

$$\textbf{C1:} \sum_{n \in \mathcal{N}} a_{mn} = 1, \forall m \in \mathcal{M}. \tag{4.5}$$

The second constraint (**C2**) ensures that the resource requirements of all VMs allocated to one server do not exceed the maximum capacity of a server subsystem and is given by,

$$\textbf{C2:} \sum_{m \in \mathcal{M}} a_{mn} \cdot \gamma_m^s \leq \lambda_n^s, \forall n \in \mathcal{N}, \forall s \in \mathcal{S}. \tag{4.6}$$

The third constraint (**C3**) ensures that the predicted CPU temperature – sum of the current CPU temperature $T_n^{t_0}$ and the predicted temperature increase $\Delta T_{[t_0, t_0+\delta]}^n$ calculated using (5.1) – is always below the recommended maximum operating temperature (T^{reco}) and is represented as,

$$\textbf{C3:} \; T_n^{t_0} + \Delta T_{[t_0, t_0+\delta]}^n \leq T^{reco}, \forall n \in \mathcal{N}. \tag{4.7}$$

The optimization problem presented here naturally forces VM consolidation. As heat generation increases logarithmically with increase in CPU utilization (shown in Fig. 4.5), the optimization prefers already loaded active servers for VM allocation when all the constraints (**C1**, **C2**, and **C3**) are met. This is because the additional cost of

placing a VM in an already loaded server (in terms of increase in temperature) decreases as the load increases. Also, constraint **C3** ensures that more VMs are allocated to servers in better-cooled areas of the datacenter. Such thermal-aware VM consolidation leads to better utilization of resources. In addition, consolidation increases the return air temperature in the consolidated server aisles thus increasing the efficiency of cooling. This can be attributed to the fact that higher the CRAC return air temperature the higher the COP of cooling.

VMAP: We characterize the aforementioned optimization problem as a *variable-size multi-dimensional bin-packing problem* [90, 91]. This is a generalized version of the traditional fixed-size one-dimensional bin-packing problem as the *bins (servers)* and *objects (VMs)* are represented as "hypercuboids" with multiple dimensions d (5 in our problem) and all the bins need not have the same capacity along each dimension. The size of each VM along the five different dimensions are its four normalized subsystem utilization requirements and the heat-generation rate. The size of a server along the five different dimensions are the normalized residual capacity (or availability) of each of the four subsystems and the heat extraction rate. The first four dimensions corresponding to VM subsystem requirements (in the object definition) and server subsystem residual capacities (in the bin definition) are straightforward to interpret and incorporated into a bin-packing problem. However, the relationship between the heat-generation (in the object definition) and heat-extraction (in the bin definition) rates is more involved. The bin capacity along the fifth dimension is actually the difference between current CPU temperature (T^{t_0}) and the upper bound of the recommended temperature range T^{reco}.

We use a multi-dimensional best-fit-like algorithm [92] to allocate a set of VMs ($\mathcal{M}$) that have arrived in a time window to a set of physical servers ($\mathcal{N}$). First, the VMs are sorted in decreasing order of their deadlines (or running time). Note that this is a shift from the traditional method of sorting based on one of the dimensions. This is because, in HPC clouds, the subsystem requirements of VMs are comparable and, hence, their durations play a pivotal role in determining energy consumption. It is desirable to pack longer duration VMs together so that server that host smaller duration VMs can be switched off at the completion of workload tasks so to save energy. Once the VMs are

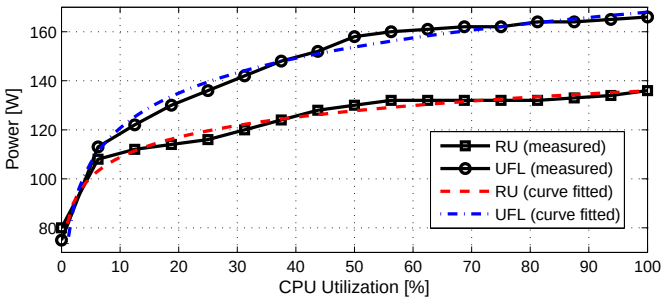

Figure 4.5: Relationship between power consumption and CPU utilization in multi-core multi-threaded systems at RU and UFL servers.

sorted according to their deadline, each VM $m \in \mathcal{M}$ is allocated a server $n \in \mathcal{N}$ whose residual volume (of the hypercuboid) is the lowest of all servers' after assignment. The time complexity of the aforementioned heuristic is $\mathcal{O}(|\mathcal{M}| \cdot \log |\mathcal{M}| + d \cdot |\mathcal{M}| \cdot |\mathcal{N}|)$, where the first and second components correspond to the sorting step and the assignment steps, respectively.

The objective of bin packing (minimize the number of bins used) is in line with the objective of the optimization problem, i.e., the fewer the active physical servers, the lower the energy consumption. This is also made possible due to the logarithmic behavior (as shown in Fig. 4.5) of CPU temperature as well as energy consumption with respect to CPU utilization in multi-core multi-threaded systems (which are the most common computing equipment configuration in cloud datacenters). In addition, bin-packing heuristics require that the objects are not further manipulated (i.e., divided or rotated) and do not overlap inside the bins (similar to constraint **C1**), the total volume of all the object inside a bin cannot exceed the bin's volume (similar to constraints **C2** and **C3**).

VMAP+: We extend the capabilities of VMAP to migrate VMs between servers (based on the utilization of the servers) for greater resource consolidation and energy savings while not exceeding the recommended operating thermal conditions. As mentioned earlier VM consolidation (i.e., allocation using VMAP) is performed every δ. The goodness of consolidation reduces over time as some VMs terminate their operation. In other words, VMAP as it is may lose the opportunity to further consolidate and save energy over time, especially in the case when δ is large. Therefore, we propose

VMAP+, an adaptive solution that builds on VMAP's capabilities. VMAP+ uses the information about VM terminations and the resource utilization (CPU, memory, disk I/O, and network subsystems) at the corresponding host servers to make migration decisions. VMAP+ determines the subset of "lightly-loaded" servers ($\mathcal{L} \subseteq \mathcal{N}$) and migrates the remaining VMs in those servers to already moderately- or highly-loaded servers. We know from Fig. 4.5 that the additional cost of running a VM in a moderately- or highly-loaded server is lower than running it in a lightly-loaded server.

VMAP+ also takes the migration cost (in terms of service delay, energy consumption, and heat generation) into consideration while choosing the candidate VMs for migration from the lightly-loaded servers. The process of VM migration is in itself a network-intensive workload. Network-intensive workloads also result in an increase in the CPU utilization thus resulting in heat generation at the host and destination servers of the candidate VM under consideration [93, 94]. VMAP+ takes into account the following key tradeoffs: energy savings vs. delay and network overhead. As VM migration involves transferring the memory of the VM from the host to the destination server, the time taken and, hence, the energy consumption depends on the network bandwidth. The higher the network bandwidth the shorter the migration time and smaller the energy footprint. Excessive migrations can cause large delays and high energy consumption compared to the no-migration case. VMAP+ migrates VMs only i) if migration does not result in excessive heat generation, ii) if energy saving is greater than the migration cost, and iii) if migration delay is within the SLA.

Let us consider a toy example in which we use only one of the dimensions (CPU) for ease of illustration. Let us define a highly-loaded server as one that uses >75% of CPU, a moderately-loaded server as one that uses 50-75% of CPU and a lightly-loaded server as one that uses <50% of CPU. Once VMs are allocated using VMAP at time t, the servers 1 and 2 are highly loaded as shown in Fig. 5.7(a). As the time elapses, one or more VMs terminate their operation and one or more highly-loaded servers become moderately-loaded and the moderately-loaded ones may become lightly-loaded. At time t' Server 2 becomes lightly-loaded and at time t'', Server 1 becomes moderately-loaded. If only VMAP were to be employed at $t + \delta$, then the servers would continue operation

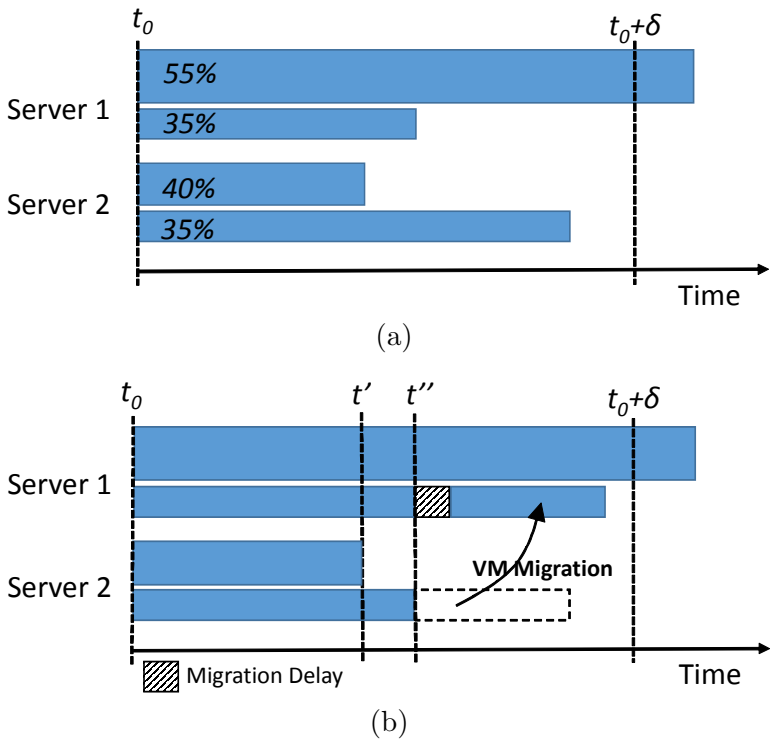

Figure 4.6: Placement and progress of VMs between two allocation instants (duration δ) (a) when only VMAP is employed; (b) when VMAP+ is employed. VMAP+ relieves Server 2 of its load and switches it off after the migrating the remaining VM thus saving on energy required for running the different server subsystems.

under the current load conditions. In general, if the rate of VM termination is greater than the rate of VM arrival, there will be a high number of light and moderately-loaded servers.

However, VMAP+ keeps track of lightly-loaded and the moderately-loaded servers and relieves lightly-loaded servers as shown at time t'' in Fig. 5.7(b). Therefore, Server can be shut down after migrating the remaining VM to Server 1 making it highly-loaded. However, the migrated VM experiences a temporary suspension in service referred to as migration delay (overhead). This delay is kept to pre-specified limits specified in the SLA. VMAP+ strives to maximize the number of highly-loaded and moderately-loaded servers and minimize the number of lightly-loaded servers. Algorithm 1 describes the migration strategy in VMAP+.

The datacenter manger has the capability to modify the definition of a lightly-loaded server. If the definition is too strict, the migration delay is small as less amount

Algorithm 1: VMAP+ migration strategy

Input: $\mathcal{L}, T^{reco}, T_n^t, \mathbf{\Lambda}_n, \mathbf{\Lambda}_n^{max}$
Output: Migration source server l and target server n
  if a VM (m) in server n finished its operation at time t **then**
    Update $\mathcal{L}$ based on the newly available resources
    for all $l \in \mathcal{L}$ **do**
      if $T_n^t + \Delta T_n \leq T^{reco}$ & $\mathbf{\Lambda}_l^{max} - \mathbf{\Lambda}_l \leq \mathbf{\Lambda}_n$ **then**
        Migrate VMs in server l to n.
        Break
      end if
    end for
  end if

of memory should be copied to the migration target (destination server). In general, servers and their VMs that consume small amount of resources can be easily migrated. However, if the definition of the lightly-loaded server is relaxed, the migration delay will increase as there will be more candidate VMs (also big in terms of resource requirements) for migration. VMAP+ is suited for HPC applications in the cloud as these applications are not elastic in nature (i.e., fixed resource requirements and predictable performance). This non-elasticity allows VMAP+ to estimate and control the overhead in terms of time (service delay due to VM suspension), energy consumption (as the migration process is modeled as a workload in itself), and heat generation.

VMAP (and hence, VMAP+) also has the ability to optimize resource allocation across a network of heterogeneous yet federated datacenters. Heterogeneity here refers to the difference in characteristics and capabilities of computing (e.g., heat-generation rate of servers, processing power, network capacity, etc.) and cooling (e.g., COP of air-chilled vs. water-chilled cooling) equipment, sources of energy for operation and cooling (e.g., renewable or non-renewable), and environmental regulations in the respective geographical region (e.g., cap on CO_2 footprint or cap on water temperature increase caused by cooling systems). We follow a two-step approach in which the problem of deciding which datacenter should handle the VM and which physical server should host the VM are determined sequentially. For example, if reducing the CO_2 footprint and the aggregate TCO are the goals, the solution will load datacenters that rely on renewable sources of energy as long as the following conditions are met: high COP of cooling,

compliance with requirements of VMs/workloads and with environmental regulations such as cap on water consumption and cap on water temperature increase caused by the cooling system. As mentioned earlier, we have a testbed of geographically separated yet federated datacenters to validate our solutions.

4.4 Performance Evaluation

We evaluated the performance of VMAP and VMAP+ via experiments on a small-scale testbed and via trace-driven simulations. The system model used in our simulations has the same characteristics of our real testbed. First, we provide details on our testbed and experiment methodology (workload traces, performance metrics, and competing approaches). Then, we elaborate on the experiment and simulation scenarios aimed at highlighting the benefits of thermal-aware VM consolidation using VMAP (only allocation) and VMAP+ (only allocation as well as migration). As the performance of VMAP+ is affected by the definition of a lightly-loaded server, a separate sensitivity analysis is presented at the end of this section.

4.4.1 Testbed and Experiment Methodology

Testbed: We have fully equipped machine rooms at two sites of NSF CAC – Rutgers University (RU) and University of Florida (UFL) – with state-of-the-art computing equipment (modern blade servers in enclosures) and fully controllable CRAC systems. The blade servers at both sites are equipped with a host of internal sensors that provide information about server subsystem operating temperatures and utilization. In addition, the machine room at RU is instrumented with an external heterogeneous sensing infrastructure [44] to capture the complex thermodynamic phenomena of heat generation and extraction at various regions inside the machine room. The sensing infrastructure comprises of scalar temperature and humidity sensors placed at the server inlet (cold aisle) and outlet (hot aisle), airflow meters at the server outlet, and thermal cameras in the hot aisle.

The computing equipment configuration at RU is two Dell M1000E modular blade

enclosures. Each enclosure is maximally configured with sixteen blades, each blade having two Intel Xeon E5504 Nehalem family quad-core processors at 2.0 GHz, forming an eight core node. Each blade has 6 GB RAM and 80 GB of local disk storage. The cluster system consists of 32 nodes, 256 cores, 80 GB memory and 2.5 TB disk capacity. The cooling equipment at RU is a fully controllable Liebert 22-Ton Upflow CRAC system. The computing equipment configuration at UFL is two IBM Blade Center with sixteen blades in each, each blade having two Intel Xeon E5504 Nehalem family quad-core processors at 2.0 GHz, forming an eight core node. Each blade has 24 GB RAM and 80 GB of local disk storage. The cluster system consists of 32 nodes, 256 cores, 768 GB memory and 2.5 TB disk capacity. The cooling equipment at UFL consists of two fully controllable Liebert 14- and 9-Ton CRAC system (Model FH302C-CA00 and FH147C-CAEI) with humidifier and reheating capacity.

Workloads: We used real HPC production workload traces from the RIKEN Integrated Cluster of Clusters (RICC) [95]. The trace included data from a massively parallel cluster, which has 1024 nodes each with 12 GB of memory and two 4-core CPUs. As the RICC is a large-scale distributed system composed of a large number of nodes, we scaled and adapted the job requests to the characteristics of our system model. First, we converted the input traces to the Standard Workload Format (SWF) [96]. Then, we eliminated failed and canceled jobs as well as anomalies. As the traces did not provide all the information needed for our analysis, we needed to complete them using a model based on [97].

The entire trace consists of 400,000 requests spread over 6 months. We extracted three versions out of this long trace, one for use in the small-scale experiments (with tens of servers) and two for use in medium-scale (hundreds of servers) simulations. The trace used in our experiments have 100 requests over the course of one day. The two other traces used in our simulations, however, have 5,200 requests spread over 2 days and 10,000 requests spread over 3 days. We assigned one of four benchmark profiles (based on *Sysbench* for CPU-intensive and *TauBench* for CPU-plus-memory-intensive workloads) to each request in the input trace, following a uniform distribution by bursts. The bursts of job requests were sized (randomly) from 1 to 5 requests.

Competing Strategies: We compared the performance of VMAP against six strategies, namely, Round Robin (RR), First-Fit-Decreasing (FFD), Best-Fit-Decreasing (BFD), FFD Reactive (FFD_R), BFD Reactive (BFD_R), and Cool-Job (CJ) [3] allocation. Of these six strategies, RR, FFD, and BFD are thermal-unaware while FFD_R, BFD_R, and CJ make reactive allocation decisions based on current temperature measurements.

- In RR, the VMs are allocated sequentially to servers.
- In FFD, the VMs corresponding to the requests that have arrived in the previous time window (of duration δ [s]) are first sorted in the decreasing order of volumes of the hypercuboids representing the VMs. Then, each VM is allocated to the first server (w.r.t. server ID) that satisfies all the four subsystem utilization requirements.
- In BFD, the VMs are again sorted according to volume as in FFD. Then, each VM is allocated to the first physical server (w.r.t. server ID), which not only satisfies all the four subsystem utilization requirements but also has the least residual volume after packing that VM.
- In FFD_R, VMs are first placed following the FFD policy. Then, VMs in overheated servers are relocated to cooler servers again based on the FFD principle.
- In BFD_R, VMs are first placed following the BFD policy. Then, VMs in overheated servers are relocated to cooler servers again based on the BFD principle.
- In CJ, each VM (that has arrived in the previous δ [s]) is allocated to the first "coolest" physical server, which satisfies all the four subsystem utilization requirements. Similar to FFD and BFD, the VMs are sorted in the decreasing order of their normalized volume. Note that CJ does not predict future temperatures like VMAP does.

Metrics: We evaluate the impact of our approach in terms of the following metrics: *energy consumption* (in kilo- Watt-hour [kWh]), and *thermal violation* (duration

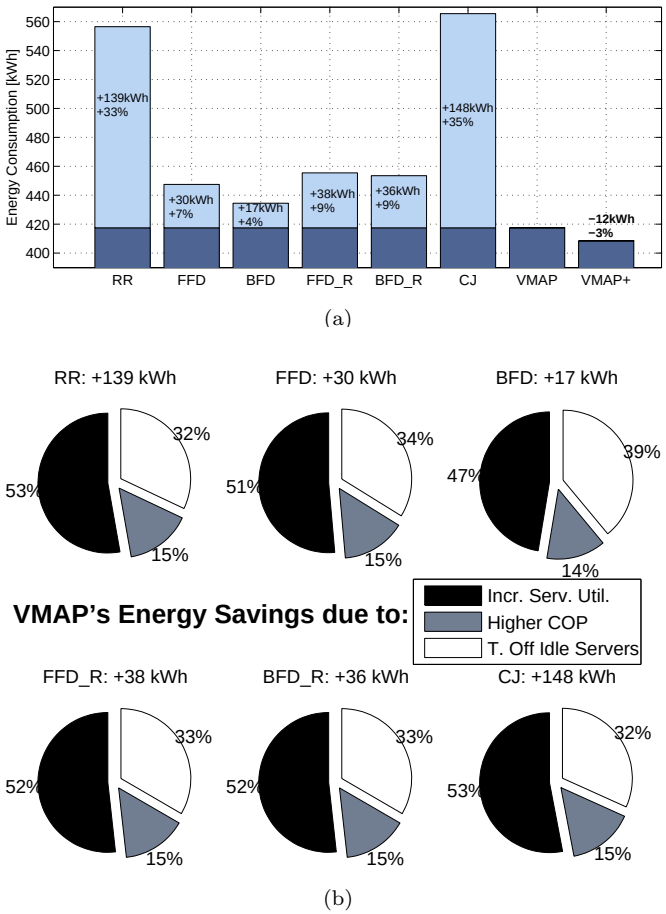

Figure 4.7: (a) VMAP's overall energy savings [kWh] in comparison to the competing algorithms; (b) Components (and their percentage) of energy savings due to: 1) increased server utilization rate, 2) efficient cooling because of the higher COP, and 3) turning off idle servers.

in second per day[s/day]). The thermal violation was calculated by monitoring the average time the servers were operating in the unsafe temperature region in a day (24 hours). Unsafe temperature region here refers to temperatures greater than the upper bound of the recommended range specified by equipment manufacturers. A higher percentage of thermal violation results in greater risk of equipment failure and/or drop in performance.

4.4.2 Energy savings

Non-consolidation vs consolidation: We performed trace-driven simulations to quantify the energy savings achieved by VMAP in a large-scale setting (180 servers and

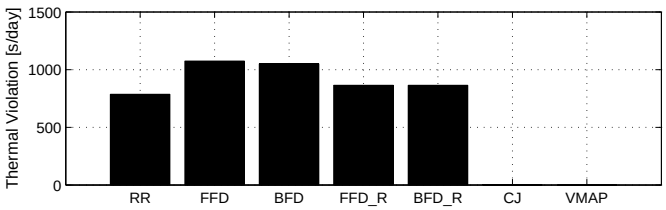

Figure 4.8: Thermally violated duration of CPU temperature per server in a day.

10,000 VM requests spread over 3 days). Figure 4.7(a) shows VMAP's energy savings in comparison to each competing algorithm. RR and CJ are the least energy efficient in comparison to VMAP as they spread the workload (VMs) over the entire datacenter (to balance the load in the case of RR and in search for the coolest server in the case of CJ). The other four schemes consolidate VMs like VMAP does, however, they consume more energy than VMAP. In Fig. 4.7(b), we analyzed different components (and their percentage of the total) of VMAP's energy savings. The main reasons for VMAP's superior energy performance are savings due to 1) increased server utilization, 2) efficient cooling because of the higher COP, and 3) turning off idle servers. Even though the actual amount of energy savings ranges from 17 (in comparison to BFD) to 148kWh (in comparison to CJ), the ratio of the three components of savings does not fluctuate significantly. It can be clearly observed that increased server utilization is the largest contributor to energy efficiency followed by shutdown of idle servers.

Non-thermal-aware vs thermal-aware: Figure 4.8 shows thermal violation of the same simulation performed above. Thermal-aware algorithms (FFD_R, BFD_R, CJ, VMAP) exhibit a smaller degree of violation in comparison with non-thermal-aware algorithms (FFD, BFD). FFD_R and BFD_R perform better in comparison to FFD and BFD because VMs from overheated servers are reallocated in reaction to thermal violation alarms. However, due to the reactive nature of these techniques, undesired equipment overheating is still an issue. VMAP and CJ avoid thermal violations. However, CJ's performance in terms of this metric is similar to VMAP's, it comes at a very high energy cost as shown in Fig. 4.7(a).

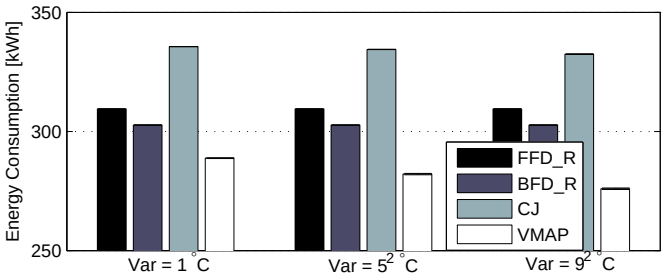

Figure 4.9: Energy consumption [kWh] under different degrees of unevenness in heat extraction (induced by difference in server inlet temperatures).

4.4.3 Consolidation in "better-cooled" areas

We performed trace-driven simulations to show how VMAP can exploit unevenness in heat imbalance inside a datacenter (with homogeneous computing equipment) caused by unevenness heat-extraction rates due to difference in server inlet temperatures. Evaluation was carried out in a small-scale setting (180 servers and 5,200 VM requests spread over 2 days). We studied the performance of the four thermal-aware techniques (FFD_R, BFD_R, CJ, and VMAP) under different degrees of Gaussian variation in the server inlet temperature; $\mathcal{N}(25,1)^{\circ}\mathbf{C}$, $\mathcal{N}(25,5^2)^{\circ}\mathbf{C}$, and $\mathcal{N}(25,9^2)^{\circ}\mathbf{C}$. Unevenness of inlet temperature of each server can be attributed to the non-ideal air circulation, which depends on the layout of server racks inside the datacenter and on the placement of CRAC unit fans and air vents. Figure 4.9 shows that the total energy consumption (for computation as well as cooling) of VMAP decreases as the degree of unevenness increases. This is because VMAP consolidates VMs in better-cooled areas where a higher heat-extraction rate leads to a lower increase in CPU temperature (with the same heat-generation rate).

4.4.4 Performance Under High COP

We performed trace-driven simulations to study VMAP's performance under varying COP. First, based on the system model of the infrastructure at RU and UFL, we carried out evaluations in a large-scale setting (180 servers and 10,000 VM requests spread over 3 days at each site). The CRAC outlet temperature in the UFL system model was set to a higher value (30 $^{\circ}C$) compared to the 25 $^{\circ}C$ in RU system model. It can be seen

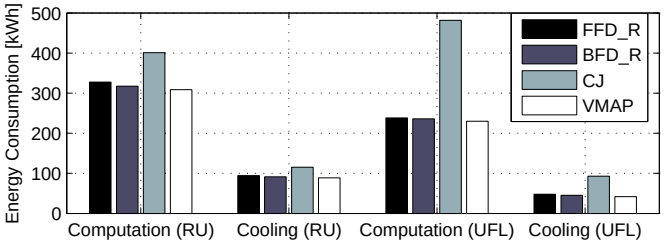

Figure 4.10: Energy consumption for computation and cooling at different datacenters – RU and UFL

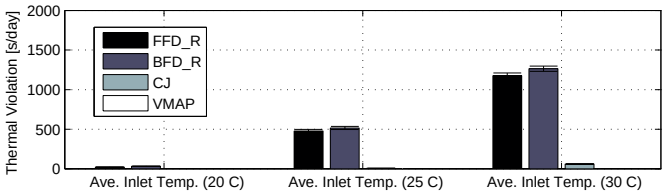

Figure 4.11: Thermal violations under different average server inlet temperatures.

in Fig. 4.10 that the energy consumption for cooling at UFL is lower than the one at RU because the COP of the CRAC system model at UFL is higher than the one at RU. COP of a CRAC unit increases with increase in the outlet temperature [3] as the work that needs to be done to reduce the hot-air temperature to 30 $^{\circ}C$ is lower than the work that needs to be done to reduce it to 25 $^{\circ}C$. We then studied the performance of VMAP and the other thermal-aware techniques using one system model (RU's) with different CRAC COPs. In Fig. 4.11, we also show that VMAP does not incur thermal violation while others do even for the servers in higher temperature.

4.4.5 Impact of Decision Window (δ)

We studied the impact of the periodicity (δ) on the performance of VMAP. Evaluation was carried out in a large-scale setting (180 servers and 10,000 VM requests spread over 3 days). If δ is big, the complexity increases because the number of VM requests ($|\mathcal{M}|$) increases. If δ is small, the complexity decreases because $|\mathcal{M}|$ decreases, but it is less efficient as only fewer VM requests can be optimized. Generally, VMAP can do better packing and save energy when δ is big but δ cannot exceed certain time bound because the extra delay incurred may violate SLA. Figure 4.12 shows high energy consumption for small δ but lower energy consumption for large δ. VMAP outperforms the other thermal-aware strategies for any δ. The best choice of δ is, however, dependent on the

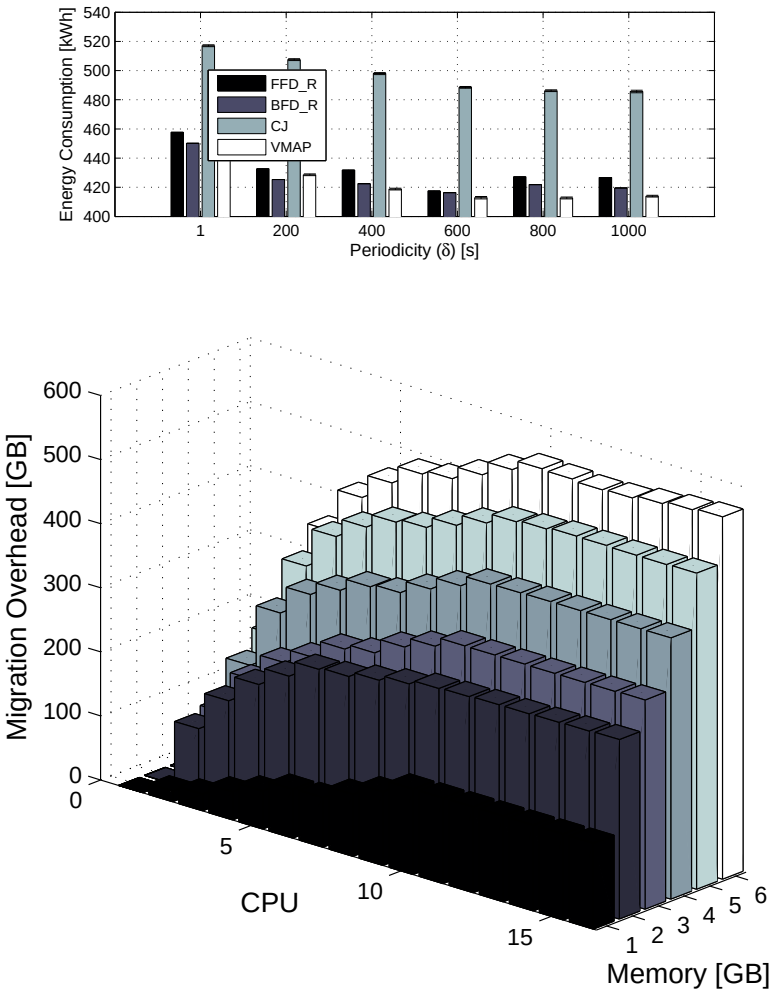

Figure 4.13: Definition of lightly loaded server (CPU-memory) and migration overhead [GB] (data), which should be migrated over the network.

workload pattern and its statistics.

4.4.6 Impact of Definition of "Lightly-loaded server"

We performed simulations to study the impact of the definition of "lightly-loaded server" on the performance of VMAP+. Evaluation was carried out in a small-scale setting (180 servers and 5,200 VM requests spread over 2 days). It is clear that the migration overhead increases when there is an increase in the amount of memory (in bytes) to be migrated from source to destination servers. Figure 5.12 shows the migration overhead in GB along the z-axis and the parameters for the definition of a lightly-loaded server along the x- and y-axes. If the definition of a lightly-loaded server is strict (i.e., those

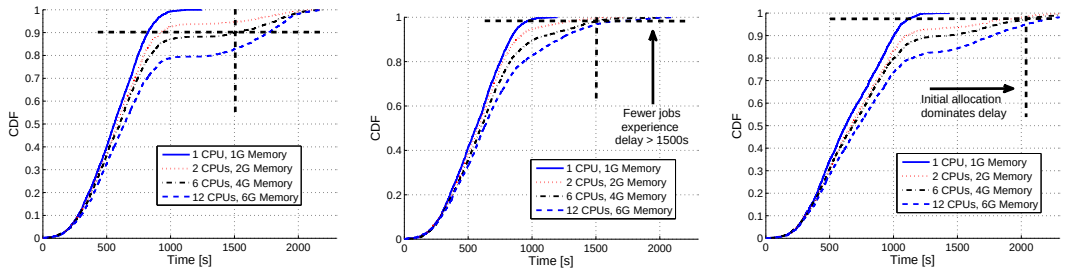

Figure 4.14: CDF of VM execution delay (time difference between requested VM termination and actual VM termination) when (a) the network capacity is 100 Mbit/s and δ is set 800s; (b) the network capacity is 1000 Mbit/s and δ is set 800s; (c) the network capacity is 1000 Mbit/s and δ is set 1200s.

servers in which a small number of CPUs and small amount of memory are being utilized), the migration overhead is small because there are small number of candidate VMs for migration. The overhead increases when we relax the definition of lightly-loaded server and the overhead saturates when the definition is extreme (i.e., anything greater than 10 active CPUs and 6GB utilized memory) as we cannot find candidate destinations to migrate a large number of VMs.

Figures 4.14(a)-(c) show the cumulative distribution function (CDF) of VM execution delays for all the jobs in the workload trace when the four different definitions of lightly-loaded server are employed. It is clear that when the definition of lightly-loaded server is relaxed (higher number of utilized CPUs and larger amount of utilized memory) the migration overhead increases (due to increase in number of candidate VMs for migration) and the service delay of a greater percentage of VM (i.e., jobs) increases. Figures 4.14(a) and (b) shows the CDF of VM execution delay when the average network bandwidth in the datacenter is 100 Mbit/s and δ is set to 800s, and when the average network bandwidth is 1000 Mbit/s and δ is set 800s, respectively. A greater percentage of jobs experience longer delays in Fig. 4.14(a) than in Fig. 4.14(b) due to the smaller network bandwidth. For example, when the definition is set to 6 CPUs and 4GB memory, nearly 10% of the jobs experience a delay of more than 1500s when the average network bandwidth is 100 Mbit/s while only less than 2% of jobs experience similar delays when the network bandwidth is 1000 Mbit/s. The larger network bandwidths a datacenter infrastructure supports, the shorter time needed to migrate VMs,

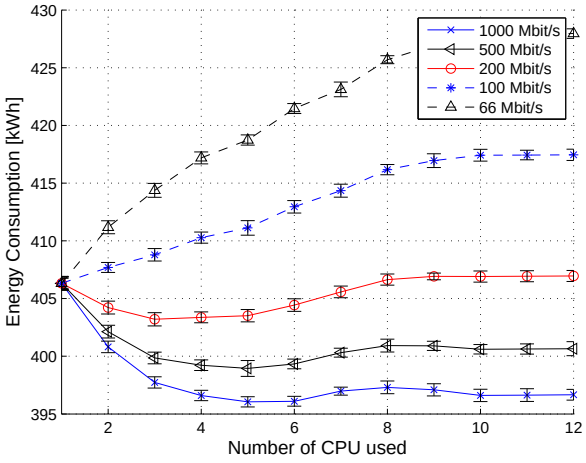

Figure 4.15: Energy consumption of VMAP+ with various network speed.

and hence, the smaller overhead (in terms of time) incurred. Figure 4.14(c) shows the CDF of VM execution delay when the network capacity is 1000 Mbit/s and δ is set 1200s. This is a case where the service delay is dominated by $\delta = 1200s$ as compared to Fig. 4.14(b) in which $\delta = 800s$.

Figure 5.13 shows the overall energy consumption under different average datacenter network bandwidths. It can be observed that the overall energy consumption for VMAP+ decreases when the average datacenter network bandwidth is high as it decreases migration delay and hence, the energy consumption at the host and destination servers. Our inference from this study is that the migration strategy is beneficial only if the network speeds are high so that the benefits of further consolidation outweigh the cost for migration (in terms of service delays and energy consumption).

4.5 Conclusions and future work

We first introduced and validated the novel concept of *heat imbalance*, which captures the unevenness in heat generation and extraction, at different regions inside a HPC cloud datacenter. We then proposed thermal-aware (knowledge of heat imbalance) proactive Virtual Machine (VM) mapping (consolidation) solution, VMAP. Our

solution maximizes computing resource utilization, minimizes datacenter energy consumption for computing, and improves the efficiency of heat extraction, while not violating recommended maximum operating temperature. We proposed VMAP+, which extends the capabilities of VMAP by performing VM migrations from lightly-loaded to moderately- and heavily-loaded servers in between two allocation instants. We verified the effectiveness of VMAP and VMAP+ through experimental evaluations with HPC workload traces at Rutgers University and University of Florida machine rooms. We observed that the VMAP is 9%, and 35% more energy efficient than best-fit and "cool job", respectively, two state-of-the-art reactive thermal-aware solutions. VMAP+ achieves a further 3% average reduction in energy consumption compared to VMAP due to its ability to further consolidate resources in between two allocation instants.

Chapter 5

Model-based Thermal Anomaly Detection in Cloud Datacenters using Thermal Imaging

5.1 Overview

High-Performance Computing (HPC) systems housed in datacenters are a key component of the society's IT infrastructure. Due to their growing importance, datacenters are strategic targets [17] for *denial-of-service attacks* (running illegitimate workloads) and *cooling system attacks* aimed at causing thermal runaways and, hence, costly outages, which can potentially cripple critical health, banking and commerce, defense, scientific research, and educational infrastructures. Furthermore, due to their large scale and high server density, the probability of computing and cooling system misconfigurations as well as of cooling equipment and server fan failures is high [18]. Such unpredictable events may result in unexpected high temperature areas/regions (hotspots) or excessively cooled low temperature areas/regions (coldspots) also referred to as *thermal anomalies*.

Local unevenness in heat-generation (by computing and communication equipment) and heat-extraction (by cooling equipment) rates determines the temperature distribution inside a datacenter. The heat-generation and -extraction rates may differ, which over time leads to what we call *heat imbalance* [66]. Unexpected changes in the local heat-generation and -extraction rates due to 1) attacks (on the computing or cooling infrastructure), 2) Computer Room Air Conditioning (CRAC) unit and server fan failures, and/or 3) computing and cooling system misconfigurations may over time cause an unexpected large positive heat imbalance resulting in a significant temperature increase and, hence, in unexpected thermal hotspots. Such hotspots may also result in thermal *fugues*, which are characterized by a continuous increase in the rate of temperature rise.

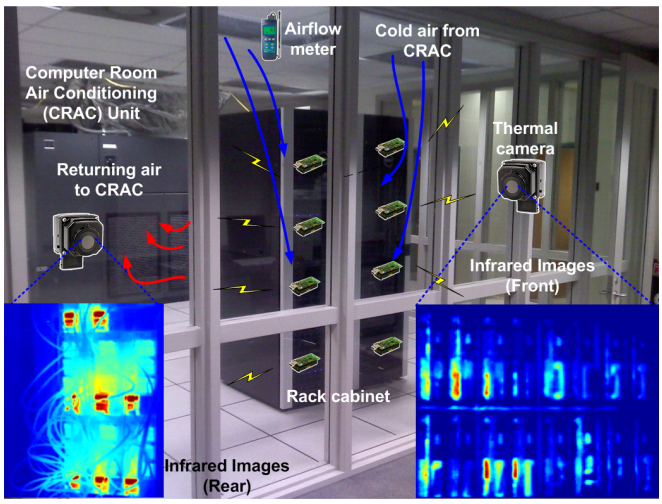

Figure 5.1: Multi-tier sensing infrastructure composed of temperature and humidity sensors, airflow meters, and thermal cameras for thermal-aware datacenter management.

Thermal anomalies such as unexpected hotspots and fugues lead to system operation in unsafe temperature regions [16], which will increase the server failure rate. Computing as well as cooling system misconfigurations may also cause energy-inefficient overcooling resulting in unexpected coldspots. In summary, thermal anomalies, i.e., hotspots, fugues, and coldspots, significantly impact the Total Cost of Ownership (TCO) of datacenters.

Thermal awareness, i.e., the knowledge of heat imbalance (for a given distribution of workloads at different regions inside a datacenter), is essential for timely detection and classification of thermal anomalies so to minimize their effect on the efficiency, availability, security, and lifetime of mission-critical HPC systems. In [98], we proposed a multi-tier sensing infrastructure (composed of temperature and humidity sensors, airflow meters, and thermal cameras as shown in Fig. 5.1) for autonomic management of datacenters (i.e., self-organizing, self-optimizing, and self-healing), and in [66], we designed and validated a simple yet robust heat-imbalance model, which exploits the data from the sensing infrastructure. The notion of heat imbalance allows us to predict future temperature maps of the datacenter and take proactive management decisions such as workload placement [66] and cooling system system optimization [26].

We propose an efficient method for online (real-time) processing and interpretation of infrared images (also called thermograms) with the knowledge of heat imbalance

at various regions inside a datacenter for thermal anomaly detection. State of the art in thermal anomaly detection in datacenters involves complex offline processing of thermograms in order to construct 2-D (two dimensional) or 3-D reconstruction of thermal maps of datacenters for detailed visual inspection by domain experts [99–101] and to identify electrical, mechanical, computing, and/or cooling system faults. On the contrary, we decompose thermograms from thermal cameras with a large field of view in to localized sub-images by exploiting our prior knowledge of the *aisle-rack-enclosure-blade layout* in order to focus on specific Regions of Interest (ROIs). Our anomaly detection scheme is *model-based*, i.e., it involves comparison of *expected* (generated using the heat-imbalance model) and *observed* (obtained from thermograms) thermal maps.

In addition to model-based thermal anomaly detection, we propose a Thermal Anomaly-aware Resource Allocation (TARA) solution, which exploits the knowledge of heat imbalance. Our idea is to create a time-varying thermal fingerprint (thermal map) of the datacenter so that the intensity of an unexpected hotspot is sufficiently high for the model-based detection to notice (even when a very low detection threshold is used). TARA allocates workloads or Virtual Machines (VMs, in case of a virtualized datacenter) to servers where the heat imbalance due to the workload is high in order to maximize the temperature difference between two consecutive thermal maps. This strategy allows for early detection of failures (as the solution tries to utilize as many servers as possible) and easy detection of attacks (such as illegitimate workloads).

Creation of consecutive thermal maps that vary significantly from each other (for early and easy detection of anomalies) is possible when there is no cap on power usage. No cap on the power budget implies unlimited access to additional servers and no restriction on costly workload migrations, which are unrealistic. Hence, we assume that TARA operates under strict power budgets and it is designed to factor in the costs of operating additional servers and of workload migrations while maximizing the detection accuracy for a given power budget. TARA can significantly contribute to the thermal anomaly detection (7%, 15%, and 31% average improvement in thermal anomaly detection with only 10% false positive rate) compared to the traditional scheduling

algorithms: random, round robin, and best-fit-decreasing.

In summary, our contributions in this chapter include:

- We propose a model-based thermal anomaly detection solution in HPC cloud datacenters that involves comparison of expected and observed thermal maps.
- We propose TARA, an anomaly-aware resource allocation (VM placement and migration) solution for virtualized datacenters, which significantly improves the accuracy of our model-based anomaly detection technique.
- We propose a solution to employ thermal imaging for anomaly detection in cloud datacenters.
- We introduce thermal image processing solutions to capture thermal features. Histogram analysis and high order statistics are employed to understand the hotspots and coldspots.

The remainder of this chapter is organized as follows: in Sect. 5.2, we present the state of the art in thermal anomaly detection in cloud datacenters; in Sect. 5.3, we discuss our proposed solution for thermal anomaly detection along with a anomaly- and energy-aware virtual machine allocation solution, which significantly improves the detection accuracy; in Sect. 5.4, we explain our evaluation methodology and the observations from our simulation study; finally, we conclude in Sect. 5.5.

5.2 State of the Art

Cloud deployments employ a Intrusion Detection Systems (IDS) - or an anomaly detection system - for a cyber-threat defense mechanism, and there are efforts to integrate the anomaly detection system into a cloud environment. An open source and distributed strategy is proposed in [102], and IDS as a service is proposed [103]. Self-healing and rejuvenation techniques in virtualized cloud environment are also proposed in [104] and evaluated in [105]. However, those methods, which rely on aggregation (at a monitor node) and on continuous online analysis of huge amounts of data (e.g., VM or job requests and distribution, server utilization, network traffic, internal sensor values,

etc.), are prohibitive in a large datacenter because of limitations in terms of network bandwidth and computational overhead. In addition, such methods are not capable of detecting thermal anomalies as they do not capture the complex thermodynamic phenomena inside a datacenter.

Hence, we propose a novel method to detect anomalies in thermal domain. We extract information from the raw measured data (using an external sensing infrastructure as shown in Fig. 5.1) and to create knowledge about the heat imbalance and impart self-protection capabilities on large-scale HPC systems. Note that our proposed thermal anomaly detection solution can be complimentary to any IDS as it operates in thermal domain, which is totally different from the traditional solution. We introduce the use of thermal cameras, which have a large field of view and can provide temperature distribution information at a greater granularity than scalar point-source temperature sensors.

In [106], the authors propose four methods to perform 'prediction' or 'early detection' of thermal anomalies so to enable proactive thermal management decisions. The first three methods are variants of a simple temperature-threshold-based approach, while the fourth method employs a Bayesian classifier to 'predict' thermal anomalies. The threshold-based methods rely heavily on the choice of the threshold and the time window used for classification of events as either normal or anomalous. Of the four proposed methods, the Bayesian classifier method performs best by predicting thermal anomalies earlier than the rest while also minimizing false positives. However, this Bayesian method takes only scalar temperature measurements as inputs, involves an offline training phase, does not use models, and finally does not provide insights into the causes and locations of the anomalies: for these reasons appropriate preemptive steps cannot be taken immediately. In contrast, we incorporate data from a heterogeneous sensing infrastructure [14, 41, 98] into models to profile thermally a datacenter in space and time, and then exploit this information for early thermal-anomaly detection by comparing it against appropriate features extracted from raw thermal images.

Recently, datacenter managers have started using infrared thermography to locate

and diagnose problems such as short cycling of the air conditioning system, loose electrical connections, worn out wires, and fan failures [99–101]. These solutions rely heavily on manual inspection of thermal images to detect anomalies. On the contrary, our thermal anomaly detection solution employs knowledge of workload distribution and multi-modal sensor data in the heat-imbalance model as well as infrared thermograms to automatically detect thermal anomalies through comparison of expected and observed thermal maps. Prior work on thermal management of datacenters target either the *uptime-conscious datacenter managers*, who are interested in eliminating hotspots (for security, reliability, and availability), or the *energy-conscious datacenter managers,* who are interested in eliminating coldspots (for energy efficiency). In contrast, our solution for thermal-aware resource allocation, TARA, is bestowed with controllable parameters that allow the datacenter managers to exploit the energy vs. anomaly-detection accuracy tradeoff.

5.3 Proposed Solution

The overall scope of our work is shown in Fig. 5.2. Our solution is composed of four main modules: VM placement (resource allocation), prediction (heat-imbalance estimation), image/data processing (thermography), and anomaly detection (comparison) modules. The VM placement module includes TARA, which allocates physical resources to the VMs in such a way that the detection accuracy is maximized for a given power budget. The prediction module uses the VM placement information and other sensor information to estimate the thermal map. Then, the model-based anomaly detection module compares the expected thermal map with the observed thermal map (processed by image/data processing module) to detect anomalies. Firstly, we elaborate on our heat-imbalance model for the estimation of expected thermal map. Secondly, we present how we extract observed thermal maps from thermograms in our solution. Thirdly, we elaborate solutions to compare the expected thermal map with observed thermal map, and detect anomalies. Finally, we formulate thermal anomaly-aware resource allocation as an optimization problem aimed at maximizing the anomaly detection probability, and propose a heuristic that balances the energy-accuracy tradeoff.

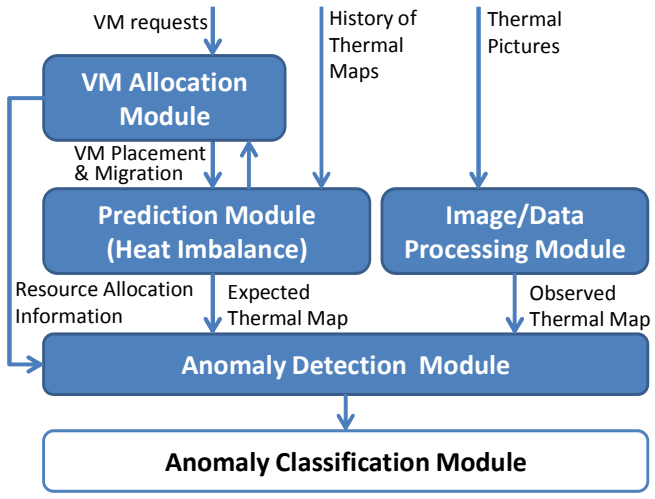

Figure 5.2: The different modules of our thermal anomaly detection solution (the focus of this chapter is indicated in blue boxes).

5.3.1 Expected Thermal Map: Model-based

A VM is created for every application request and is provisioned with resources (CPUs, memory, disk, and network capacity) that satisfy the application requirements (usually deadline). Without any loss of generality, we assume that this provisioning has already been performed using techniques such as the ones described in [89]. The provisioned VMs now have to be allocated to physical servers housed within racks in datacenters. Let $\mathcal{M}$ be the set of VMs to be allocated and $\mathcal{N}$ be the set of servers. An associativity binary matrix $\mathbf{A} = \{a_{mn}\}$ (with $a_{mn} \in \{0, 1\}$) specifies whether VM m is hosted at server n or not. A VM m is specified as a vector $\mathbf{\Gamma}_m = \{\gamma_m^s\}$, where $s \in \mathcal{S} = \{CPU, MEM, IO, NET\}$ refers to the server subsystems and γ_m^s's are the VM subsystem requirements (e.g., CPU cores, amount of volatile memory [MB], disk storage space [MB], network capacity [Mbps]).

Representation (or *mapping*) of a VM's subsystem requirement (γ_m^s) as a factor of physical server subsystem capacity is straightforward if all the servers of the datacenter are assumed to be homogeneous. For example, a VM m requiring 4 virtual CPUs, 2GB of RAM, 64GB of hard-disk space, and 100Mbps network capacity can be represented as $\mathbf{\Gamma_m} = \{0.25, 0.125, 0.125, 0.1\}$ if all the servers in a datacenter have 16 CPU cores, 16GB of RAM, 512GB of local hard-disk space, and a gigabit ethernet interface. The

mapping problem becomes non trivial in an heterogeneous environment. However, assuming that only a small finite number of generations of each subsystem are present in the datacenter, we create such a mapping for each generation of every subsystem.

Heat imbalance: In our previous work [66], we formulated the heat-imbalance model in a datacenter based on heat-generation (h_n) and heat-extraction (q_n) rates as follows,

$$\Delta I_n = \int_{t_0}^{t_0+\delta} (h_n - q_n)dt = M_n \cdot C \cdot \Delta T^n_{[t_0,t_0+\delta]}, \tag{5.1}$$

where ΔI_n [J] denotes the heat imbalance of CPU inside server n during the time between t_0 and $t_0 + \delta$, and M_n and C denote the mass and specific heat capacity, respectively, of the CPU. Note that if ΔI_n is positive (i.e., $h_n > q_n$), the temperature of the CPU at server n increases in the time interval $[t_0, t_0 + \delta]$ (hence, $\Delta T^n > 0$); conversely, if ΔI_n is negative (i.e., $h_n < q_n$), the temperature of the CPU at server n decreases (hence, $\Delta T^n < 0$). This estimated heat imbalance helps us generate the expected thermal map of the datacenter as long as the information about the different workloads and their potential location (physical server) as well as real-time temperature and airflow measurement are known.

5.3.2 Observed Thermal Map

Infrared thermogram is used to acquire accurate thermal maps of the datacenter. As the thermogram gives higher resolution than scalar sensors, the it provides more information about the object. Histogram analysis and high order statistics are employed to understand the hotspots (i.e., size, intensity, degree) and their behavior over time. However, as the thermogram includes denser and wide-view information (each pixel indicates the temperature of certain location) than a scalar sensor measurement, the observed thermal map could be readily over fitted giving wrong observation result. Hence, the ROIs of the thermogram and the image features should properly selected.

Regions of Interest (ROI): We decompose thermograms from thermal cameras with a large field of view to localized sub-images in order to focus on specific ROIs (e.g., racks, enclosures, servers, and fans). This decomposition allows for fast diagnosis

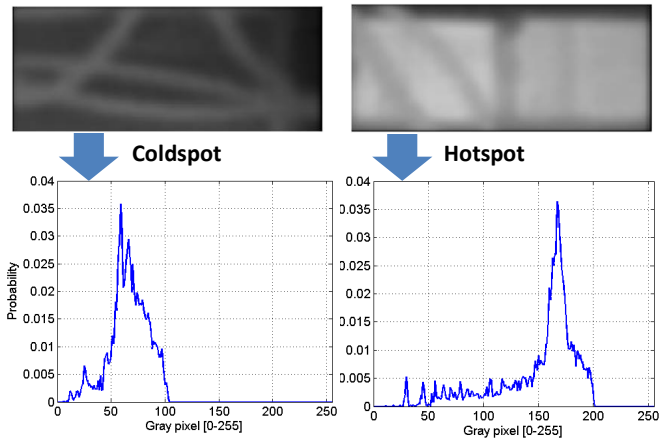

Figure 5.3: Histogram representation of thermal images capturing a coldspot on an idle server (left) and a hotspot on a 70% utilized server (right).

of thermal anomalies through comparison of localized thermal maps of ROIs with the estimated thermal maps from the heat-imbalance model. Our proposed solution does not require accurate 2-D or 3-D thermal maps of the entire datacenters. Thermogram decomposition and the ROI selection can be a simple, yet powerful solution, as the servers, enclosures, racks and vents has clear edges and the shapes are rectangular and static.

Histogram Analysis: The output of a thermal camera is a gray scale image formed using infrared radiation from an object. Image pixels indicate temperature measures of different points on the object's surface. To extract the temperature of the observed thermal map, one specific point of a pixel value (i.e., specific location, or mean, min, max value of ROI) can be used to represent the object's temperature, but the more useful information can be extracted with the help of image processing technique. We employ higher-order histogram statistics as they convey information about not only the intensity of the image pixels (temperature) but also their distribution. Histogram is a graphical representation of the 'distribution' of data (pixel values). Figure 5.3 shows the pdf of pixel intensities in an image of server fan vents under two different settings, idle and 70% CPU utilization, where the pixel value 0 is calibrated to $15°C$ with the pixel step difference is $0.075°C$.

Higher-order Histogram Statistics: The statistics are obtained by first getting the histogram of pixel values from an $N \times M$ sub-image matrix I, which has been processed

using a median blurring filter as well as a Gaussian blurring filter, to remove the noise. We then calculate an empirical probability density function $p(x)$ (where $0 \leq x \leq 255$) from the relative frequencies of each pixel in the sub-image. Using the standard notation to denote p's moments-about-the-mean (μ_i) and its standard deviation (σ), histogram statistics can be calculated, where N denotes the number of histogram bins. Followings are the 5 histogram statistic feature set ($\mathcal{F}$) considered in this chapter, where $f \in \mathcal{F} = \{energy, Hsize, skewness, kurtosis, entropy\}$.

- **Energy**: Measures the energy of the histogram distribution, in the usual manner:

$$energy = \sum_{x=0}^{N} p(x)^2$$

- **Hotspot size**: Count the 'number of pixels,' which are greater than the threshold bin number:

$$Hsize = \sum_{x=threshold}^{N} p(x)$$

- **Skewness**: Formally, the 'third standardized moment'; informally, measures asymmetry of the histogram distribution:

$$skewness = \frac{\mu_3^3}{\sigma^3}$$

- **Kurtosis**: Formally, the 'fourth standardized moment'; informally, measures the 'peakedness' of the histogram distribution:

$$kurtosis = \frac{\mu_4^4}{\sigma^4} - 3$$

- **Entropy**: Measures the 'uncertainty' or randomness inherent in the histogram distribution; also, a measure of the complexity:

$$entropy = -\sum_{x=0}^{N} p(x) \log_2 p(x)$$

Extraction of energy, hotspot size, skewness, kurtosis, and entropy allows us to

determine how to interpret the pixel values to get the best representation of the temperature of the server. *Energy* metric helps us to understand how the overall trend of server temperature is, not the specific point of temperature, and the *Hsize* metric helps us to measure the heat influx if it is measured in time because it measures the size of the hotspot. *Skewness* and *kurtosis* help us understand whether a thermal image has a hotspot or a coldspot and how big it is, respectively. Skewness gives information regarding the asymmetry of the histogram. Positive skewness indicates presence of a coldspot as there are more pixels in low temperature (Fig. 5.3 left), and negative skewness indicates presence of a hotspot as there are more pixels in high temperature (Fig. 5.3 right). Thus, more accurate representation of the temperature can be extracted by jointly employing the features of thermograms.

In addition to the higher-order histogram statistics, we have investigated *Haralick Features*, which were first introduced in [107]. These features are indirect features as they are not directly extracted from the raw image, but instead, extracted from *coocurrence matrix* [108], which is originally known as *gray-tone spatial dependency matrices* [107]). It is not discussed in this chapter in detail because we focused more on the higher-order histogram statistics, but a full discussion of the relative benefits and information inherent in these features can be found in [107].

5.3.3 Anomaly Detection Module

The abstracted features of hotspots and coldspots are determined through the histogram analysis features of their thermograms. These features are then jointly selected, and translated to the observed thermal map in order to detect anomalies by comparing it with the estimated thermal map. However, the estimated thermal map may not be directly compared with those observed thermal features like *energy* and *Hsize* as their unites are different. Hence, we proposed a simple linear-scaled metric to translate the features into temperature as follows,

$$T_f = T_{idle} + \alpha_f (f - f_{idle}). \tag{5.2}$$

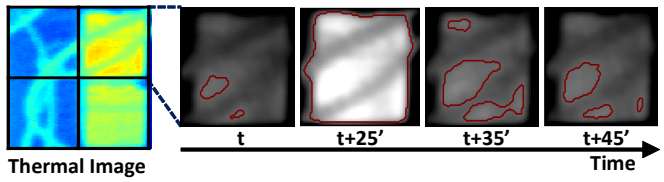

(a)

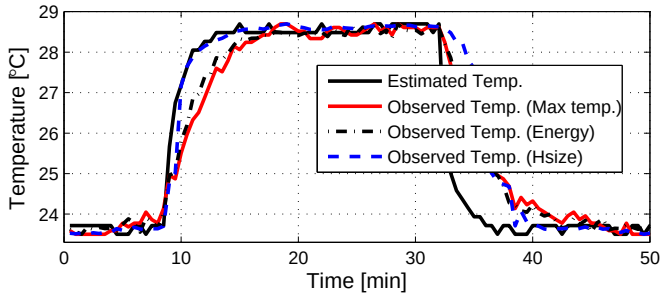

(b)

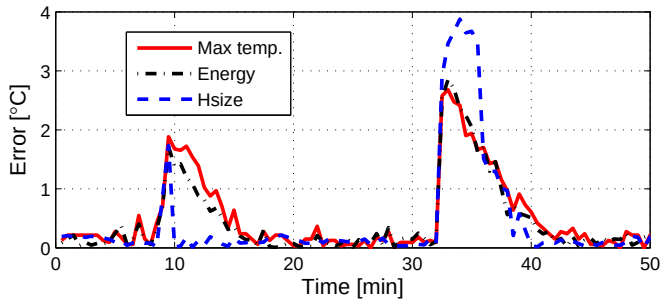

(c)

Figure 5.4: (a) Original thermal image (left), and ROI selection and hot spot extraction (right); (b) Real temperature estimation using different features of hotspots; (c) Error (difference in temperature against actual CPU Temperature) in time.

It shows the linearly scaled temperature T_f from the feature f with the factor of α_f, where T_{idle} and f_{idle} indicates the temperature and the feature value when the server is idle, respectively. The linear scale parameter α_f needs to be tuned apriori as it may differ depending on the field of view (number of total pixels that covers the area in the ROI), and the camera specifications.

We have performed a small-scale experiment to tune the α_f and show the performance of thermal map extraction using different features. Our experiment had performed by running a server 35 min with 100% CPU utilization rate and making hotspots until it reaches high steady state, and cooling it down until it reaches low steady state when the server is idle. Figure 5.4 (a) shows a original thermal image

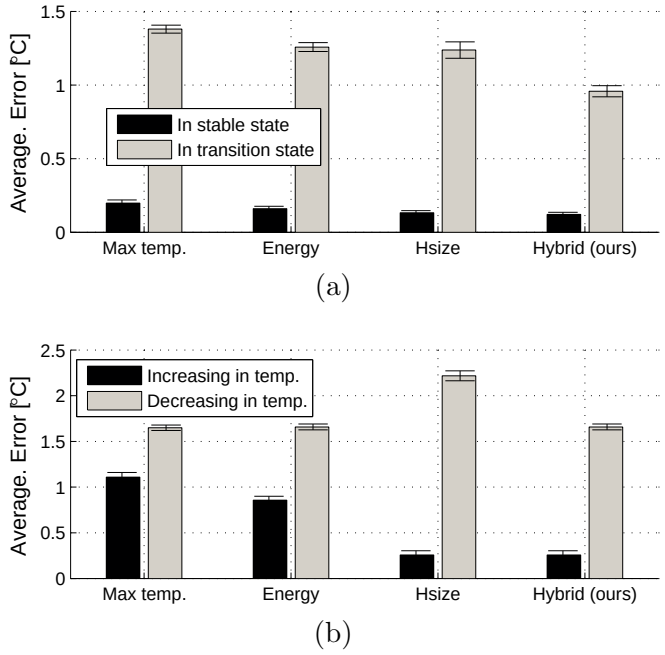

Figure 5.5: (a) Average error in different states; (b) Average error in different temperature changing trend.

used for this experiment and ROI selection, and also shows changes in thermogram over time. Hotspots are indicated in dark-red line. Figure 5.4 (b) and (C) shows the observed temperature using different features $\mathcal{F}$, and their error values compared with the estimated temperature using 5.1 [66]. We have used an energy and Hsize features for this experiment as we observe that they reflects the temperature more accurate than the other features. We have also used the maximum temperature of ROI as for a comparison assuming it represents hotspots in ROI. Note that the error incurs due to the delay of the heat transfer from inside to outside of the servers.

We observe that using the Hsize feature gives less error in increasing temperature but higher error in decreasing temperature than using other features. We imply that the hot air heated by the internal computing components increases the size of hot spot immediately when the server is running, but once the metal part of the vent get heated up, it takes time in decreasing temperature showing delay in reducing hotspot size. In contrast, using max temperature or energy feature gives less error in decreasing temperature. Thus, we propose to use hybrid feature, which uses the Hsize feature when the temperature is in increasing trend and uses energy feature when the temperature is in

decreasing trend to minimize the error. Figures 5.5 (a) and (b) shows the average error of our Hybrid solution compared with other feature extraction solution in different state and different temperature trend, respectively. Our Hybrid solution achieve minimum error for temperature estimation by jointly employing the features of thermograms.

5.3.4 Thermal Anomaly-aware Resource Allocation: TARA

We propose a novel VM allocation solution (TARA), which increases the probability of detecting unexpected hotspots (thermal anomalies) when the number of computing resources that can be operational for a given load (set of workload requests) is increased. TARA harnesses information from our sensing infrastructure (thermal cameras, airflow meters, and internal sensors), the VM requests, and the prediction module (heat imbalance model) to allocate VMs such that the orthogonality between two consecutive thermal maps. Then, the original thermal map, captured online, is compared with the predicted one using our heat-imbalance model [66] to detect anomalous VM placement. TARA is comprised of thermal anomaly-aware VM placement and migration strategies.

Thermal Anomaly-aware VM Placement: Our novel anomaly-aware resource placement strategy enables self-protection by inducing as much difference between consecutive thermal maps (generated every δ seconds) as possible when placing VM in cloud datacenters. This strategy allows for early detection of thermal failures and easy detection of attacks such as illegitimate workloads because the solution tries to utilize as many servers as possible and changing their thermal map frequently. This frequent changes in thermal map gives more opportunity to detect long-lasting unexpected thermal hotspots, anomalies.

Energy-accuracy Tradeoff: Our VM placement method generate thermal maps every δ seconds given the budget (β), which restricts the number of servers that can be utilized in the datacenter. Parameter β represents the fraction of total datacenter resources that can be used and the corresponding power budget is given by, $P_\beta = P_{max} \cdot \beta$, where P_{max} is the power consumption when all the computing resources are utilized to the maximum in the datacenter. The higher the power budget, the greater the number of servers that can be utilized, and hence, the greater the difference between thermal maps (resulting

in a higher detection rate). TARA allows the datacenter managers to exploit the budget β to explore the tradeoff between energy expenditure and anomaly detection accuracy. Note that β is always greater than β_0, which is the minimum fraction of resources in the datacenter that is required for a given set of workloads to be completed without compromising their quality of service.

An uptime-conscious datacenter manager may be interested in eliminating hotspots so to increase service availability (and hence, use a high β). An energy-conscious datacenter manager, on the other hand, may be interested in eliminating coldspots so to improve energy efficiency (and hence, use a low β). TARA can autonomously allocate resources within the budget that the datacenter manager provides. Our VM allocation solution changes the thermal map when allocating resources to VMs every δ seconds. Frequent thermal map change increases the robustness of our anomaly detection module. We first formulate the VM allocation problem as an optimization problem, which employs our heat-imbalance model. The motivation for formulating the optimization problem is to gain insight and make key design decisions for our heuristic solution.

Optimization Problem: The goal is to find an *optimal* mapping of VMs to physical servers (represented by the binary associativity matrix $\mathbf{\Lambda}$) so to maximize the difference between the existing thermal map and the expected thermal map when the workloads in the VMs are all active. The known (given as well as measured) parameters and optimization variables of the optimization problem can be summarized as,

$$\begin{aligned}
&\textbf{Given (offline)} && : \mathcal{N}, T^{reco}, \delta, M_n, C_p; \\
&\textbf{Given (online)} && : \beta_0, \beta, \mathcal{M}, \mathbf{\Gamma}_m \ \forall m \in \mathcal{M}; \\
&\textbf{Measured (online)} && : T_n^{t_0}, m_n^{in}, T_n^{in}, \mathbf{\Lambda}_n \ \forall n \in \mathcal{N}; \\
&\textbf{Find} && : \mathbf{A} = \{a_{mn}\},\ m \in \mathcal{M}, n \in \mathcal{N}.
\end{aligned} \tag{5.3}$$

Here, $T_n^{t_0}$ and $\mathbf{\Lambda}_n = \{\lambda_n^s\}$ represent the current CPU temperature and the maximum residual capacity of each subsystem s at server n, respectively. The objective of the

optimization problem is,

$$\textbf{Maximize}: \sum_{n \in \mathcal{N}} |T_n^{t_0} - \tilde{T}_n^{t_0+\delta}|; \tag{5.4}$$

$$\textbf{Subject to}: \textbf{ C1, C2, C3, C4.}$$

Here, $\tilde{T}_n^{t_0+\delta} = T_n^{t_0} + \Delta T^n_{[t_0,t_0+\delta]}$ is the estimated temperature of an "active" server n at time $t_0 + \delta$ and $\Delta T^n_{[t_0,t_0+\delta]}$ is calculated using (3). If the server is unused, then we set $\Delta T^n_{[t_0,t_0+\delta]} = 0$. The first constraint,

$$\textbf{C1: } \sum_{n \in \mathcal{N}} a_{mn} = 1, \forall m \in \mathcal{M}, \tag{5.5}$$

ensures that a VM is allocated to *one* and *only one* server. The second constraint,

$$\textbf{C2: } \sum_{m \in \mathcal{M}} a_{mn} \cdot \gamma_m^s \leq \lambda_n^s, \forall n \in \mathcal{N}, \forall s \in \mathcal{S}, \tag{5.6}$$

ensures that the resource requirements of all VMs allocated to one server do not exceed the maximum capacity of a server subsystem. The third constraint,

$$\textbf{C3: } T_n^{t_0} + \Delta T^n_{[t_0,t_0+\delta]} \leq T^{reco}, \forall n \in \mathcal{N}, \tag{5.7}$$

ensures that the predicted CPU temperature – sum of the current CPU temperature $T_n^{t_0}$ and the predicted temperature increase $\Delta T^n_{[t_0,t_0+\delta]}$ calculated using the heat-imbalance model – is always below the recommended maximum operating temperature (T^{reco}), which is chosen by the datacenter manager. The fourth constraint,

$$\textbf{C4: } \beta_0 \leq \frac{\sum_{n \in \mathcal{N}} r_n}{|\mathcal{N}|} \leq \beta, \forall n \in \mathcal{N}, \tag{5.8}$$

ensures that the specified utilization factor of the datacenter (in terms of number of active servers) is not exceeded. Here, r_n is an indicator variable that conveys whether

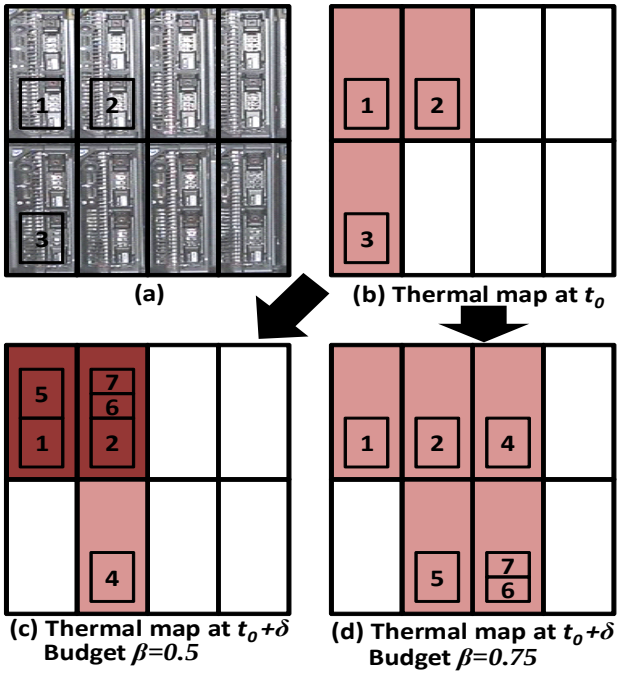

Figure 5.6: Toy example of VM allocation using TARA with different budgets ($\beta = 0.5, 0.75$): (a) An actual enclosure with eight blade-servers; (b) Thermal map of an enclosure at t_0 with three active VMs; A possible allocation of four newly arrived VMs (4 to 7) and the corresponding thermal map at $t_0 + \delta$ (c) when the budget $\beta = 0.5$ ($\beta_0 = 0.375$) and (d) when the budget $\beta = 0.75$ ($\beta_0 = 0.375$).

a server n is active or not,

$$r_n = \begin{cases} 1 & \text{if} \sum_{m \in \mathcal{M}} a_{mn} \geq 1 \\ 0 & \text{otherwise} \end{cases}, \forall n \in \mathcal{N}. \tag{5.9}$$

This optimization problem is NP-hard as it needs to find the maximum temperature difference when allocating $|\mathcal{M}|$ workloads to $|\mathcal{N}|$ servers (combinatorial problem). Hence, we present our heuristic solution in Algo. 2.

Algorithm 2: Thermal Anomaly-aware VM Placement Strategy

for $i = 1 \rightarrow |\mathcal{M}|$ **do**
1. **Find** server(s) $k \in \mathcal{K} \subset \mathcal{N}$, where $P_\beta/|\mathcal{N}| > P_s$;
2. **Find** a server $k \in \mathcal{K}$, where $|T_k^t - \tilde{T}_k^{t+\delta}|$ is maximal, $\tilde{T}_k^{t+\delta} < T^{reco}$, and $\gamma_i \leq \mathbf{\Lambda}_k$;
3. **Place** ith VM to the server k ;
4. **Update A** ;

end

The objective (maximize the temperature difference to maximize the detection accuracy) of TARA is in line with the objective of the optimization problem, i.e., the more the active physical servers, greater the detection accuracy. This is also made possible due to the logarithmic behavior (as shown in Fig. 4.3) of CPU temperature with respect to CPU utilization in multi-core multi-threaded systems which are measured using two different platforms, Dell cluster (PowerEdge M620 blade server) and IBM cluster (blade center) at Rutgers University (RU) and University of Florida (UFL) machine rooms (which are the most common computing equipment configuration in cloud datacenters). As the temperature difference is smaller when utilizing higher number of cores, the detection rate decreases under high utilization (at each server).

Toy Example: Figure 5.6(a) shows an actual enclosure of eight servers. Figure 5.6(b) shows an example temperature map at time t_0. Figure 5.6(c) and (d) show thermal map change from time t_0 to $t_0 + \delta$ under different power budgets of $\beta = 0.5$ and 0.75, respectively. For the scenario depicted in Fig. 5.6, β_0 is set to 0.375. The difference between the thermal maps at t_0 and $t_0 + \delta$ can be increased by increasing the budget β. Here, the difference between thermal maps at t_0 and $t_0 + \delta$ are $30°C$ and $48°C$ for $\beta = 0.5$ (Fig. 5.6(c)) and $\beta = 0.75$ (Fig. 5.6(d)), respectively. Note that VM 3 is not shown in (c) and (d) as it is assumed to have ended right after t_0. Thermal anomalies due to attacks can be easily detected and the ones due to failures can be detected early when the budget $\beta = 0.75$ as opposed to the case when $\beta = 0.5$ because the expected hotspots' intensities are kept low.

When placing VMs, the temperature of the server increases and the difference between the two maps is maximized given the budget as VMs. When some VMs are finished (e.g., VM number 3), temperature decreases and the difference is still high because they were placed where the heat-imbalance is maximum. However, the difference cannot be maximized as we do not know the length of the VMs and when the VMs finish. We will discuss more the length of the jobs and its impact in Sect. 5.5. The More servers are utilized in increase in budget, the higher intensity of thermal map change. The difference between thermal maps in temperature at t_0 and $t_{0+\delta}$ are $30, 48$, and $54°C$ for the budget $\beta = 0.5, 0.75$, and 1.0, respectively.

Thermal Anomaly-aware VM Migration: TARA has a VM migration capability to boost anomaly detection rate and energy savings within the power budget, while not exceeding the recommended operating thermal conditions. The goodness of detection rate may reduces over time because some VMs terminate their operation changing the observed thermal map. It may decrease the difference between thermal maps, and hence, decrease the sensitivity of anomaly detection rate. In other words, TARA may lose the opportunity to detect more anomalies over time, especially in the case when δ is large.

Migration strategy: TARA's migration strategy enables a strong defensive mechanism to deal with cyber-threats. It allows for creating periodically (δ) alternative random migrations (called *moving target*), while maintaining the power budget and reasonable service delay. When the average life time of the VMs is high compared to δ, then the average number of VMs at any point in time in the datacenter will be very high. Under such circumstances, $\beta_0 \simeq 1$ and β cannot be increased further. The only way the detection rate can be increased (especially for attacks such as spy VMs) is by migrating the high-intensity hotspots to various regions inside the datacenter. This "moving target" or "hopping" of high-intensity hotspots in space and time makes it a moving target problem for attackers.

Also, the attackers break the cloud computing system through specific IP that she/he has an access or a physical hardware depending on the the external or internal intrusion. However, the moving target makes invaders difficult to identify the physical or virtual resources. Even if the entire system is compromised and even the VM hypervisor, the anomalous behavior of the servers can be detected in thermal domain as the energy consumption based on the job allocation generates heats, and the thermal camera provides another layer of security. The frequent random migration of VM's physical location in the network will deceive the invader's discoveries.

TARA uses the information of VM terminations and the resource utilization (CPU, memory, disk I/O, and network subsystems) at the corresponding host servers to make migration decisions. TARA determines the subset of "highly-loaded" servers ($\mathcal{H} \subseteq \mathcal{N}$) and migrates VMs to already lighly-loaded servers. We know from our previous

work [66] that the temperature difference of running a VM in a highly-loaded server is lower than running it in a moderately-loaded or lightly-loaded server. As this different level of temperature difference determines the detection rate (high detection rate in a moderately-loaded or lightly-loaded server and low detection rate in a highly-loaded server), TARA strives to migrate VMs from highly-loaded server to moderately-loaded or lightly-loaded servers. We present our VM migration strategy in Algo. 3.

Algorithm 3: Thermal Anomaly-aware VM Migration Strategy

if *a VM in server n finished its operation at time t and the server n becomes lightly-loaded* **then**

1 **Find** a random server $h \in \mathcal{H}$;

2 **Find** a VM v running in h, where $\tilde{T}_n^{t+\delta} < T^{reco}$, and $\gamma_v \leq \mathbf{\Lambda}_h$;

3 **Migrate** VM (v) in server h to n

end

Migration cost: TARA also takes the migration cost (in terms of service delay, energy consumption, and heat generation) into consideration while choosing the candidate VMs for migration from the highly-loaded servers. The process of VM migration is in itself a network-intensive workload. Network-intensive workloads also result in an increase in the CPU utilization thus resulting in heat generation at the host and destination servers of the candidate VM under consideration [93, 94]. TARA takes into account the following key tradeoffs: detection rate vs. energy savings vs. delay and network overhead. As VM migration involves transferring the memory of the VM from the host to the destination server, the time taken and, hence, the energy consumption depends on the network bandwidth. The higher the network bandwidth the shorter the migration time and smaller the energy footprint and higher the detection rate. Excessive migrations can cause large delays and high energy consumption compared to the no-migration case. TARA migrates VMs only i) if migration does not result in excessive heat generation, ii) if the migration cost is within the power budget, and iii) if migration delay is within the SLA.

Toy Example: Let us consider a toy example in which we use only one of the dimensions (CPU) for ease of illustration. Let us define a highly-loaded server as one that uses >75% of CPU, a moderately-loaded server as one that uses 25-75% of CPU

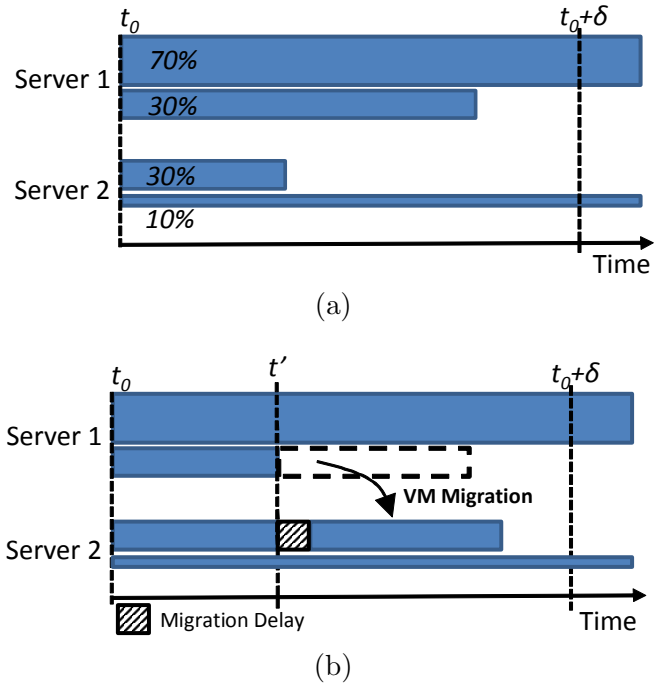

Figure 5.7: Placement and progress of VMs between two allocation instants (duration δ) (a) without VM migration; (b) with VM migration capability. Server 1 retains some margin to detect thermal anomalies by migrating the VM from the highly-loaded to lightly-loaded server.

and a lightly-loaded server as one that uses <25% of CPU. Once VMs are allocated using TARA at time t_0, the servers 1 and 2 are highly loaded as shown in Fig. 5.7(a). As the time elapses, one or more VMs terminate their operation and one or more the moderately-loaded ones may become lightly-loaded while some servers are heavily-loaded. If only TARA were to be employed at $t + \delta$, then the servers would continue operation under the current load conditions, not changing thermal map. In general, if the rate of VM arrival is greater than the rate of VM termination, there will be a high number of high and moderately-loaded servers.

However, TARA keeps track of highly-loaded and the lightly-loaded servers and relieves highly-loaded servers as shown at time t' in Fig. 5.7(b). Therefore, the thermal map of each server changes more frequently after migrating the VM to Server 2 making it moderately or lightly-loaded. However, the migrated VM experiences a temporary suspension in service referred to as migration delay (overhead). This delay is kept to pre-specified limits specified in the SLA. TARA strives to maximize the number of moderately-loaded and lightly-loaded servers and minimize the number of highly-loaded

servers within the power budget.

The datacenter manger has the capability to modify the definition of a highly-loaded server. If the definition is too strict, the migration delay is small as less amount of memory should be copied to the migration target (destination server). In general, servers and their VMs that consume small amount of resources can be easily migrated. However, if the definition of the high-loaded server is relaxed, the migration delay will increase as there will be more candidate VMs (also big in terms of resource requirements) for migration. TARA is suited for HPC applications in the cloud as these applications are not elastic in nature (i.e., fixed resource requirements and predictable performance). This non-elasticity allows TARA to estimate and control the overhead in terms of time (service delay due to VM suspension), energy consumption (as the migration process is modeled as a workload in itself), and heat generation.

5.4 Performance Evaluation

We evaluated the performance of TARA, in terms of thermal-anomaly-detection accuracy, via experiments on a small-scale testbed, and via trace-driven simulations. We performed small-scale experiments (16 servers) to design large-scale simulations (220 servers) under realistic assumptions. The system model used in our simulations has the same characteristics of our real testbed (e.g., temperature profile of CPU, room temperature, and workload profile). First, we provide details on our testbed and experiment methodology (workload traces, performance metrics, and competing approaches). Then, we elaborate on the simulation scenarios aimed at highlighting the benefits of anomaly-aware VM allocation for efficient anomaly detection.

5.4.1 Testbed and Simulation Methodology

Testbed: We have a fully equipped machine room in NSF CAC at RU with state-of-the-art computing equipment (modern blade servers in enclosures) and a fully controllable CRAC system. The blade servers are equipped with a host of internal sensors that provide information about server subsystem operating temperatures and utilization.

In addition, the machine room at RU is instrumented with an external heterogeneous sensing infrastructure [98] to capture the complex thermodynamic phenomena of heat generation and extraction at various regions inside the machine room. The sensing infrastructure comprises of scalar temperature and humidity sensors placed at the server inlet (cold aisle) and outlet (hot aisle), airflow meters at the server outlet, and thermal cameras in the hot aisle.

The computing equipment configuration is two Dell M1000E modular blade enclosures. Each enclosure is maximally configured with sixteen blades, each blade having two Intel Xeon E5504 Nehalem family quad-core processors at 2.0 GHz, forming an eight core node. Each blade has 6 GB RAM and 80 GB of local disk storage. The cluster system consists of 32 nodes, 256 cores, 80 GB memory and 2.5 TB disk capacity. The cooling equipment is a fully controllable Liebert 22-Ton Upflow CRAC system.

Workloads: We used real HPC production workload traces from the RIKEN Integrated Cluster of Clusters (RICC) [95]. The trace included data from a massively parallel cluster, which has 1024 nodes each with 12 GB of memory and two 4-core CPUs. As the RICC is a large-scale distributed system composed of a large number of nodes, we scaled and adapted the job requests to the characteristics of our system model. First, we converted the input traces to the Standard Workload Format (SWF) [96]. Then, we eliminated failed and canceled jobs as well as anomalies. As the traces did not provide all the information needed for our analysis, we needed to complete them using a model based on [97]. The entire trace consists of 400,000 requests spread over 6 months. The trace used in our simulation have 5,200 requests spread over 2 days.

Competing Strategies: We compared the performance of TARA against four strategies, namely, Random, Round Robin (RR), Best-Fit-Decreasing (BFD), and an energy-plus-thermal-aware consolidation technique VMAP [66]. In Random, the VMs are allocated in random sequence to any server. In RR, the VMs are allocated sequentially to servers. In BFD, the VMs are sorted according to volume. Then, each VM is allocated to the first physical server (w.r.t. server ID), which not only satisfies all the four subsystem utilization requirements but also has the least residual volume after packing that VM. In VMAP, VM is allocated to minimize the energy consumption

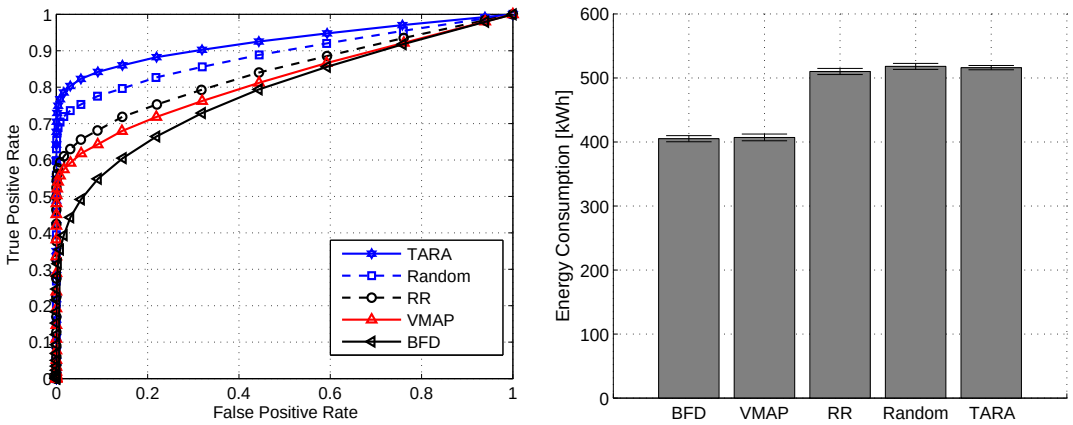

(a) (b)

Figure 5.8: (a) ROC of TARA with competing algorithms; (b) Energy consumption of TARA with competing algorithms.

while ensuring that the servers do not overheat.

Metrics: We evaluate the impact of our approach in terms of the following metrics: True Positive Rate (TPR) and False Positive Rate (FPR) of anomalies, and *energy consumption* (in kilowatt-hour [kWh]). TPR and FPR are depicted using *Receiver Operating Characteristic (ROC)* [109] curves. The TPR and FPR are calculated using different thresholds and the resulting points are plotted as a ROC curve. The larger the area under ROC curve, the better the performance of the detection in terms of accuracy.

Figure 5.8(a) shows ROC of TARA and competing algorithms. When the budget is high enough, TARA outperforms other competing algorithm. Random placement and RR show higher detection rate than VMAP and BFD because they inherently spread the VMs like TARA does resulting in a large number of lightly-loaded servers in which unexpected hotspots can easily be identified. However, VMAP and BFD consolidate VMs making the temperature map change due to anomalous events insignificant, resulting in low detection rate. Thus, energy consumption shown in Fig. 5.8(b) of VMAP and BFD is lower than TARA, VMAP, RR, and Random as they save a significant amount of energy by turning of the unused servers. TARA's anomaly-detection-rate is higher than those of other non-consolidation schemes even though its energy consumption is comparable to others'.

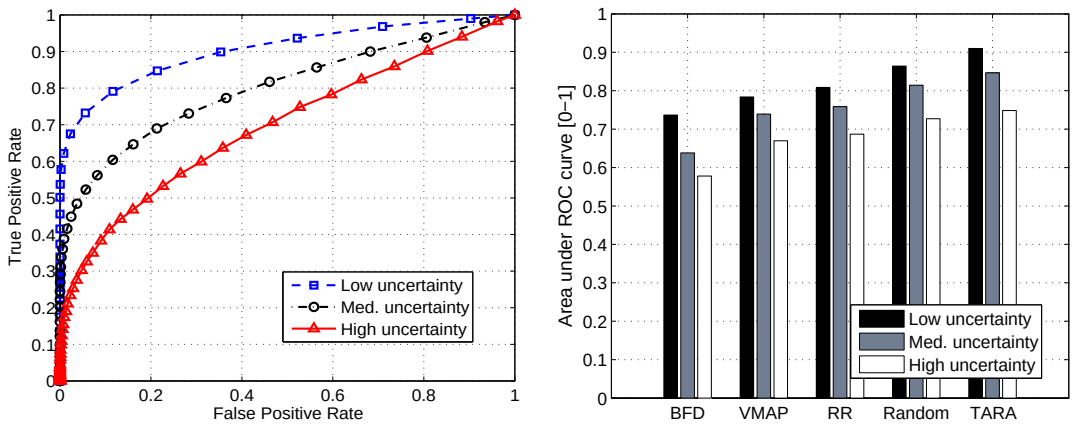

Figure 5.9: (a) ROC of TARA with different degree of uncertainty; (b) Area under ROC of TARA with different degree of uncertainty.

5.4.2 Model-based Anomaly Detection

We performed simulations to see the impact of uncertainty in data reported by the hardware sensors on the detection rate. Figure 5.9(a) shows ROC of TARA with different degrees of sensor-data uncertainty. As our heat imbalance model uses the temperature measurement we generate the Gaussian noise with the same mean ($\mu = 0°C$), but different standard deviations (σ = 0.5, 2.5, and 4.5$°C$, for low, medium, high environment uncertainties, respectively). The performance of model-based anomaly detection drops when uncertainty increases because our heat-imbalance model cannot perform well when the input is too noisy. Figure 5.9(b) shows the bar graph of area under ROC curve with different algorithms. The area under ROC curve of TARA's is higher than that of other schemes' representing higher detection accuracy for different degrees of uncertainty.

Figure 5.10(a) shows the ROC of TARA under different power budgets. Detection accuracy improves with increase in the power budget as more servers are available to distribute the workload and to maximize the difference between consecutive thermal maps (the objective of TARA). Figure 5.10(b) shows average power usage given the power budget. It shows that instead of exploiting the entire power budget TARA explores the solution space to find the most power efficient configuration that can provide the highest possible detection accuracy for a given power budget.

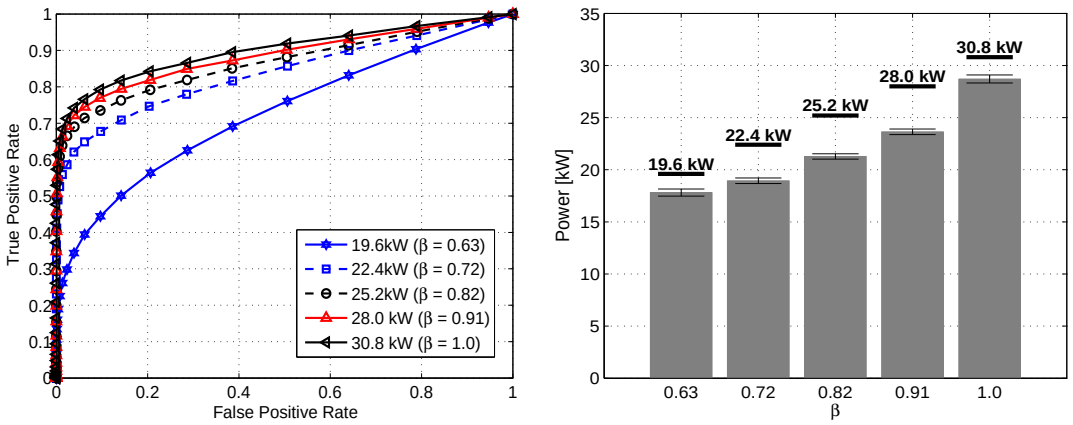

(a) (b)

Figure 5.10: (a) ROC of TARA with different power budget (b) Average power level of TARA with different power budget (here, $\beta_0 = 0.55$).

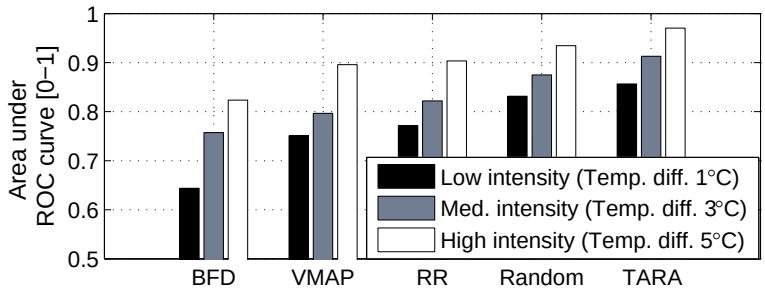

Figure 5.11: Area under ROC of TARA and competing algorithms in the presence of anomalies of different intensities.

Figure 5.11 shows the ROC of TARA given anomalies of different intensities. As the intensity and the distribution of unexpected hotspots caused by different types anomalies can be different as summarized in Table 5.1, we designed simulations with anomalies of different intensities (1, 3, and $5°C$, for low, medium, and high intensity, respectively). Our model-based anomaly detection mechanism in conjunction with TARA performs the best under all scenarios. It shows that the detection rate increase when the anomaly intensity increases even when a low threshold is used in our model-based anomaly detection mechanism.

5.4.3 Impact of VM Migration

We performed simulations to study the impact of the definition of "highly-loaded server" on the performance of TARA. Evaluation was carried out in a small-scale setting (180

Table 5.1: Intensity and distribution of thermal anomalies.

Causes of Thermal Anomaly		Intensity	Distribution
Attack	DOS attack	High	Group
	SPY VM	Low	Sporadic
Failure	CRAC fan	High	Group
	Server fan	High	Sporadic
Misconfig.	Misplacement	Normal	Diff. from Orig.
	Profiling Error	Diff. from Orig.	Normal

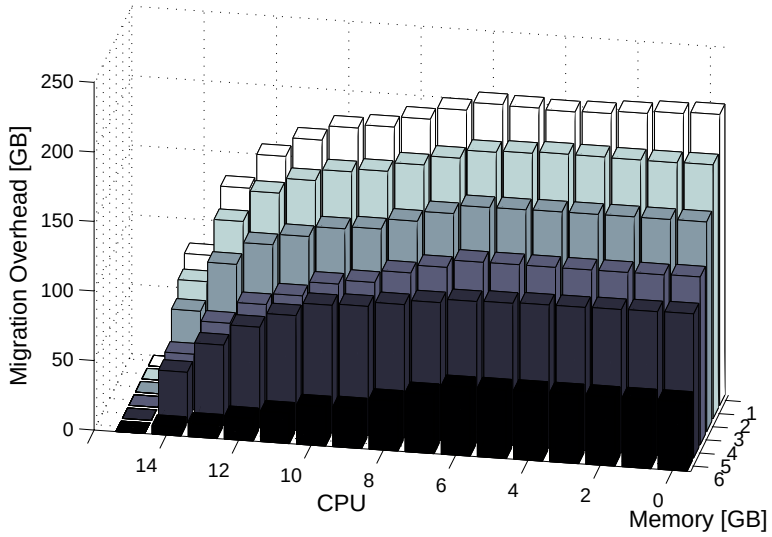

Figure 5.12: Definition of highly-loaded server (CPU-memory) and migration overhead [GB] (data), which should be migrated over the network.

servers and 5,200 VM requests spread over 2 days). It is clear that the migration overhead increases when there is an increase in the amount of memory (in bytes) to be migrated from source to destination servers. Figure 5.12 shows the migration overhead in GB along the z-axis and the parameters for the definition of a highly-loaded server along the x- and y-axes. If the definition of a highly-loaded server is strict (i.e., those servers in which a high number of CPUs and high amount of memory are being utilized), the migration overhead is small because there are small number of candidate servers for migration. The overhead increases when we relax the definition of highly-loaded server, and the overhead saturates when the definition is extreme (i.e., anything greater than 4 active CPUs and 1GB utilized memory) as the migration unnecessarily creates another highly-loaded server.

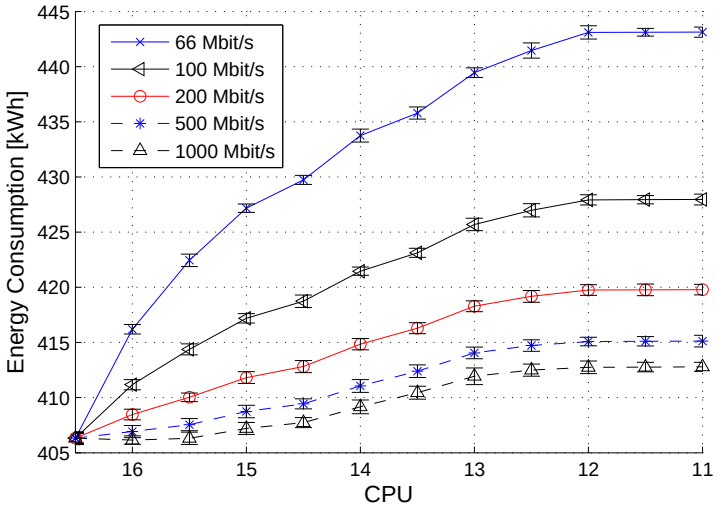

Figure 5.13: Energy consumption of TARA with various network speed based on the different definition of highly-loaded server (CPU).

Figure 5.13 shows the overall energy consumption under different average datacenter network bandwidths. It can be observed that the overall energy consumption for TARA decreases when the average datacenter network bandwidth is high as it decreases migration delay and hence, the energy consumption at the host and destination servers. Our inference from this study is that the migration strategy is beneficial only if the network speeds are high so that the benefits of further consolidation outweigh the cost for migration (in terms of service delays and energy consumption).

5.5 Discussion

Anomalies (i.e., attacks, misconfigurations, hardware failures) are becoming a significant concern for the datacenter managers as the failure of detecting them can cost a large business millions of dollars in loses. Our model-based thermal anomaly detection solution in conjunction with TARA can significantly improve the detection probability (7%, 15%, and 31% average improvement in detection with only 10% false positive rate) compared to model-based anomaly detection with traditional scheduling algorithms: random, round robin, and best-fit-decreasing.

Chapter 6

Conclusion

Due to the need of modern datacenters with the paradigm of cloud computing and big data analytics, the complexity of modern datacenters is growing at an alarming rate and their management is rapidly exceeding human ability, making autonomic approaches essential. In this thesis, we have proposed proactive thermal-aware datacenter management solutions, which include thermal- and energy-aware resource provisioning, cooling system optimization, and anomaly detection, to help minimize both the impact on the environment and the Total Cost of Ownership (TCO) of datacenters, making them energy efficient and green.

We have introduced four-layered architecture endowed with four different abstract components for the proactive thermal-aware solutions. Our proactive solutions autonomically manage datacenters using cross-layer information collected from the four-layered architecture and make decisions based on various application-specific optimization goals (e.g., performance, energy efficiency, anomaly detection rate). We have discussed the methods for acquiring thermal awareness using real-time measurements and heat and air circulation models. Also, we have proposed specific proactive thermal- and energy-aware solutions, which i) optimize cooling systems (i.e., air conditioner compressor duty cycle, and fan speed) to prevent heat imbalance and minimize the cost of cooling, and ii) maximize computing resource utilization to minimize datacenter energy consumption for computing and improve the efficiency of heat extraction.

References

[1] Growth in Data center electricity use 2005 to 2010.

[2] Uptime Institute Data Center Survey Results 2013.

[3] Justin Moore, Jeff Chase, Parthasarathy Ranganathan, and Ratnesh Sharma. Making Scheduling "cool": Temperature-aware Workload Placement in Data Centers. In *Proc. of USENIX Annual Technical Conference (ATEC)*, Anaheim, CA, 2005.

[4] Report to congress on server and data center energy efficiency. Technical report, U.S. Environmental Protection Agency, August 2007.

[5] Water Efficiency Management in Datacenters (Part I).

[6] Ishfaq Ahmad and Sanjay Ranka. *Handbook of Energy-Aware and Green Computing*. CRC Press, 2004.

[7] Wang, Jun and Feng, Ling and Xue, Wenwei and Song, Zhanjiang. A survey on energy-efficient data management. *SIGMOD Rec.*, 40(2):17–23, 2011.

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

[9] Cosmin Rusu, Alexandre Ferreira, Claudio Scordino, and Aaron Watson. Energy-efficient real-time heterogeneous server clusters. In *IEEE Real-Time and Embedded Technology and Applications Symp.*, pages 418–428, 2006.

[10] Yiyu Chen, Amitayu Das, Wubi Qin, Anand Sivasubramaniam, Qian Wang, and Natarajan Gautam. Managing server energy and operational costs in hosting centers. In *ACM SIGMETRICS Intl. Conf. on Measurement and modeling of computer systems*, pages 303–314, 2005.

[11] Jeffrey O. Kephart, Hoi Chan, Rajarshi Das, David W. Levine, Gerald Tesauro, Freeman Rawson, and Charles Lefurgy. Coordinating multiple autonomic managers to achieve specified power-performance tradeoffs. In *Intl. Conf. on Autonomic Computing*, page 24, 2007.

[12] Eun Kyung Lee, Indraneel Kulkarni, Dario Pompili, and Manish Parashar. Proactive Thermal Management in Green Datacenter. *The Journal of Supercomputing*, 60(2):165–195, May 2012.

[13] ChiehJan Mike Liang, Jie Liu, Liqian Luo, and Andreas Terzis. Racnet: Reliable acquisition network for high-fidelity data center sensing. http://research.microsoft.com/apps/pubs/default.aspx?id=70648.

[14] R. R. Schmidt. Thermal Profile of a High-Density Data Center-Methodology to Thermally Characterize a Data Center. *American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) Transactions*, 110(2):635–642, 2004.

[15] S. Kang, R. R. Schmidt, K. Kelkar, and S. Patankar. A Methodology for the Design of Perforated Tiles in Raised Floor Data Centers using Computational Flow Analysis. *IEEE Transactions on Components and Packaging Technologies*, 24(2):177–183, 2001.

[16] Jayanth Srinivasan, Sarita V. Adve, Pradip Bose, and Jude A. Rivers. The Impact of Technology Scaling on Lifetime Reliability. In *Proc. of the International Conference on Dependable Systems and Networks (DSN)*, June 2004.

[17] Jaeyeon Jung, Balachander Krishnamurthy, and Michael Rabinovich. Flash crowds and denial of service attacks: Characterization and implications for cdns and web sites. In *Proc. of the Intl. Conference on World Wide Web (WWW)*, Honolulu, HI, May 2002.

[18] Rongliang Zhou, Zhikui Wang, Cullen E. Bash, Tahir Cade, and Alan McReynolds. Failure resistant data center cooling control through model-based thermal zone mapping. In *Proc. of ASME Summer Heat Transfer Conference (HT)*, Puerto Rico, USA, July 2012.

[19] Krishna Kant. Data center evolution. *Comput. Netw.*, 53(17):2939–2965, 2009.

[20] Eun Kyung Lee, Hariharasudhan Viswanathan, and Dario Pompili. SILENCE: Distributed Adaptive Sampling for Sensor-based Autonomic Systems. In *Proc. of the ACM Intl. Conference on Autonomic computing (ICAC)*, Karlsruhe, Germany, June 2011.

[21] John Rath. Data Center St rategies. Technical report, Vantage Data Centers, 2011.

[22] U. Hoelzle and L. A. Barroso. *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*. Morgan and Claypool, 2009.

[23] Krishna Kant. Data center evolution: A tutorial on state of the art, issues, and challenges. *Computer Networks*, 53(17):2939–2965, 2009.

[24] C. Patel, C. Bash, L. Belady, L. Stahl, and D. Sullivan. Computational Fluid Dynamics Modeling of High Compute Density Data Centers to Assure System Inlet Air Specifications. In *Proc. of Pacific Rim/ASME International Electronic Packaging Technical Conferenceof (IPACK)*, Kauai, HI, 2001.

[25] A.H. Beitelmal and C.D. Patel. Computational Fluid Dynamics Modeling of High Compute Density Data Centers to Assure System Inlet Air Specifications. *Distributed and Parallel Databases*, 21(2-3):227–238, 2007.

[26] Eun Kyung Lee, Indraneel Kulkarni, Dario Pompili, and Manish Parashar. Proactive Thermal Management in Green Datacenter. *Journal of Supercomputing (Springer)*, 60(2):165–195, May 2012.

[27] 6sigmaDC. http://www.futurefacilities.com/software/6SigmaDCOverview.php.

[28] FloVENT. http://www.mentor.com/products/mechanical/products/flovent.

[29] TileFlow. http://inres.com/products/tileflow/overview.html.

[30] Paul Rad, Kailash Karki, and Tim Webb. Water Efficiency Management in Datacenters (Part I), 2008.

[31] R. R. Schmidt, E. E. Cruz, and M. K. Iyengar. Challenges of Data Center Thermal Management. *IBM Journal of Research and Development*, 49(4/5):709–723, 2005.

[32] Roger Schmidt and Ethan Cruz. Cluster of High-Powered Racks Within a Raised-Floor Computer Data Center: Effect of Perforated Tile Flow Distribution on Rack Inlet Air Temperatures. *Journal of Electronic Packaging*, 126(4):510–518, 2004.

[33] R. R. Schmidt. Thermal Management of Office Data Processing Centers. In *Proc. of Pacific Rim/ASME International Electronic Packaging Technical Conference (INTERpack)*, Kohala Coast, HI, 1997.

[34] T. Heath, A.P. Centeno, P. George, L. Ramos, Y. Jaluria, and R. Bianchini. Mercury and freon: temperature emulation and management for server systems. In *Proc. of the Architectural Support for Programming Languages and Operating Systems (ASPLOS-XII)*, San Jose, CA, 2006.

[35] ANSYS Inc. Ansys hpc for high-fidelity insight and design team productivity. http://www.ansys.com/products/hpc/default.asp.

[36] Qinghui Tang, T. Mukherjee, S.K.S. Gupta, and P. Cayton. In *Proc. of Intelligent Sensing and Information Processing (ICISIP), title=Sensor-Based Fast Thermal Evaluation Model For Energy Efficient High-Performance Datacenters*, pages 203 –208, 2006.

[37] J. Liu, B. Priyantha, F. Zhao, C.M. Liang, Q. Wang, and S. James. Towards Discovering Data Center Genome Using Sensor Networks. In *Proc. of the Embedded Networked Sensors (HotEmNets)*, Charlottesville, VA, 2008.

[38] Erin Mannas and Scott Jones. Add thermal monitoring to reduce data center energy consumption. http://pdfserv.maxim-ic.com/en/an/AN4334.pdf.

[39] R.K. Sharma, C.E. Bash, and R.D. Patel. Dimensionless Parameters For Evaluation of Thermal Design and Performance of Large-Scale Data Centers. In *Proc. of ASME/AIAA Joint Thermophysics and Heat Transfer Conference*, St. Louis, MO, 2002.

[40] R. R. Schmidt and E. Cruz. Raised Floor Computer Data Center: Effect on Rack Inlet Temperatures of Exiting both the Hot and Cold Aisle. In *Proc. of Itherm Conference (ITHERM)*, San Diego, CA, 2002.

[41] R. R. Schmidt, K. Karki, K. Kelkar, A. Radmehr, and S. Patankar. Measurements and Predictions of the Flow Distribution through Perforated Tiles in Raised Floor Data Centers. In *Proc. of Pacific Rim/ASME International Electronic Packaging Technical Conferenceof (IPACK)*, Kauai, HI, 2001.

[42] Jeffrey Rambo and Yogendra Joshi. Modeling of data center airflow and heat transfer: State of the art and future trends. *Distributed and Parallel Databases*, 21(2-3):193–225, 2007.

[43] S. Greenberg, E. Mills, and B. Tschudi. Best Practices for Data Centers: Lessons Learned from Benchmarking 22 Data Centers. In *Proc. of American Council for an Energy-Efficient Economy (ACEEE)*, Pacific Grove, CA, 2006.

[44] Hariharasudhan Viswanathan, Eun Kyung Lee, and Dario Pompili. Self-organizing Sensing Infrastructure for Autonomic Management of Green Datacenters. *IEEE Network*, 25(4):34–40, June 2011.

[45] Jim Fink. Impact of High Density Hot Aisles on IT Personnel Work Conditions. Technical report, Schneider Electric, 2011.

[46] Justin Moore, Jeff Chase, Parthasarathy Ranganathan, and Ratnesh Sharma. Making scheduling "cool": temperature-aware workload placement in data centers. In *Annual Conf. on USENIX Annual Technical Conf.*, pages 5–5, 2005.

[47] Justin D. Moore, Jeffrey S. Chase, and Parthasarathy Ranganathan. Weatherman: Automated, online and predictive thermal mapping and management for data centers. In *Intl. Conf. on Autonomic Computing*, pages 155–164, 2006.

[48] Taliver Heath, Ana Paula Centeno, Pradeep George, Luiz Ramos, Yogesh Jaluria, and Ricardo Bianchini. Mercury and freon: temperature emulation and management for server systems. In *Intl. Conf. on Architectural Support for Programming Languages and Operating Systems*, pages 106–116, 2006.

[49] Cullen Bash and George Forman. Cool job allocation: Measuring the power savings of placing jobs at cooling-efficient locations in the data center. In *USENIX Annual Technical Conf.*, pages 363–368, 2007.

[50] Qinghui Tang, Sandeep Kumar S. Gupta, and Georgios Varsamopoulos. Energy-efficient thermal-aware task scheduling for homogeneous high-performance computing data centers: A cyber-physical approach. *IEEE Trans. Parallel Distrib. Syst.*, 19(11):1458–1472, 2008.

[51] T. Mukherjee, A. Banerjee, G. Varsamopoulos, S. Gupta, and S. Rungta. Spatio-temporal thermal-aware job scheduling to minimize energy consumption in virtualized heterogeneous data centers. *Computer Networks, Elsevier*, 53(17):2888–2904, 2009.

[52] Ivan Rodero, Eun Kyung Lee, Dario Pompili, Manish Parashar, Marc Gamell, and Renato J. Figueiredo. Exploiting VM Technologies for Reactive Thermal Management in Instrumented Datacenters. In *Workshop on Energy Efficient Grids, Clouds, and Clusters in conjunction with IEEE Grid*, October 2010.

[53] Norman Bobroff, Andrzej Kochut, and Kirk Beaty. Dynamic Placement of Virtual Machines for Managing SLA Violations. In *Proc. of IFIP/IEEE Symp. on Integrated Network Management (IM)*, pages 119–128, Munich, Germany, May 2007.

[54] A. Kochut and K. Beaty. On Strategies for Dynamic Resource Management in Virtualized Server Environments. In *Proc. of the Intl. Symp. on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*, pages 193–200, Istanbul, Turkey, October 2007.

[55] Fabien Hermenier, Xavier Lorca, Jean-Marc Menaud, Gilles Muller, and Julia Lawall. Entropy: A Consolidation Manager for Clusters. In *Proc. of the ACM SIGPLAN/SIGOPS Conf. on Virtual Execution Environments (VEE)*, pages 41–50, Washington, DC, March 2009.

[56] Jan Stoess, Christian Lang, and Frank Bellosa. Energy Management for Hypervisor-based Virtual Machines. In *Proc. of USENIX Annual Technical Conf. (ATEC)*, pages 1–14, Santa Clara, CA, June 2007.

[57] William Voorsluys, James Broberg, Srikumar Venugopal, and Rajkumar Buyya. Cost of Virtual Machine Live Migration in Clouds: A Performance Evaluation. In *Proc. of Intl. Conf. on Cloud Computing (CloudCom)*, pages 254–265, Beijing, China, December 2009.

[58] Ying Song, Yuzhong Sun, Hui Wang, and Xining Song. An adaptive resource flowing scheme amongst vms in a vm-based utility computing. In *IEEE Intl. Conf. on Computer and Information Technology*, pages 1053–1058, 2007.

[59] Daniel A. Menasce and Mohamed N. Bennani. Autonomic virtualized environments. In *Intl. Conf. on Autonomic and Autonomous Systems*, page 28, 2006.

[60] Xiaoqiao Meng, Canturk Isci, Jeffrey Kephart, Li Zhang, Eric Bouillet, and Dimitrios Pendarakis. Efficient Resource Provisioning in Compute Clouds via VM Multiplexing. In *Proc. of Intl. Conf. on Autonomic Computing (ICAC)*, pages 11–20, Washington, DC, June 2010.

[61] Ripal Nathuji and Karsten Schwan. Virtualpower: coordinated power management in virtualized enterprise systems. In *ACM SIGOPS Symp. on Operating Systems Principles*, pages 265–278, 2007.

[62] Ripal Nathuji, Canturk Isci, and Eugene Gorbatov. Exploiting platform heterogeneity for power efficient data centers. In *Intl. Conf. on Autonomic Computing*, page 5, 2007.

[63] Gregor Laszewski, Lizhe Wang, Andrew J. Younge, and Xi He. Power-aware scheduling of virtual machines in dvfs-enabled clusters. In *IEEE Intl. Conf. on Cluster Computing*, pages 1–10, 2009.

[64] Sanjay Kumar, Vanish Talwar, Vibhore Kumar, Parthasarathy Ranganathan, and Karsten Schwan. vmanage: loosely coupled platform and virtualization management in data centers. In *Intl. Conf. on Autonomic Computing*, pages 127–136, 2009.

[65] Shen Li, Hieu Le, Nam Pham, Jin Heo, and Tarek Abdelzaher. Joint optimization of computing and cooling energy: Analytic model and a machine room case study. In *Proc. of IEEE International Conference on Distributed Computing Systems (ICDCS)*, pages 396–405, Macau, China, June 2012.

[66] Eun Kyung Lee, Hariharasudhan Viswanathan, and Dario Pompili. VMAP: Proactive Thermal-aware Vitual Machine Allocation in HPC Cloud Datacenters. In *Proc. of the IEEE Intl. Conference on High Performance Computing (HiPC)*, Pune, India, December 2012.

[67] Luiz Ramos and Ricardo Bianchini. C-oracle: Predictive thermal management for data centers. In *Intl. Symp. on High-Performance Computer Architecture*, pages 111–122, 2008.

[68] Timothy Wood, Gabriel Tarasuk-Levin, Prashant Shenoy, Peter Desnoyers, Emmanuel Cecchet, and Mark D. Corner. Memory Buddies: Exploiting Page Sharing for Smart Colocation in Virtualized Data Centers. In *Proc. of the ACM SIGPLAN/SIGOPS Conf. on Virtual Execution Environments (VEE)*, pages 31–40, Washington, DC, March 2009.

[69] Diwaker Gupta, Sangmin Lee, Michael Vrable, Stefan Savage, Alex C. Snoeren, George Varghese, Geoffrey M. Voelker, and Amin Vahdat. Difference Engine: Harnessing Memory Redundancy in Virtual Machines. *Communications of the ACM*, 53(10):85–93, October 2010.

[70] Sriram Govindan, Arjun R. Nath, Amitayu Das, Bhuvan Urgaonkar, and Anand Sivasubramaniam. Xen and co.: Communication-aware CPU Scheduling for Consolidated Xen-based Hosting Platforms. In *Proc. of the Intl. Conf. on Virtual Execution Environments (VEE)*, pages 126–136, San Diego, CA, June 2007.

[71] Q. Zhu, J. Zhu, and G. Agrawal. Power-aware consolidation of scientific workflows in virtualized environments. In *Proc. of High Performance Computing, Networking, Storage and Analysis (SC)*, New Orleans, LA, 2010.

[72] C. Hsu, Wu. Feng, and J.S. Archuleta. Towards Efficient Supercomputing: A Quest for the Right Metric. In *Proc. of IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Denver, Colorado, April 2005.

[73] Y. Liu and H. Zhu. A Survey of the Research on Power Management Techniques for High-Performance Systems. *Software: Practice and Experience*, November 2009.

[74] K.W. Cameron, R. Ge, and X. Feng. High-Performance, Power-Aware Distributed Computing for Scientific Applications. *Computer*, 38(11):40–47, 2005.

[75] EPA. EPA Report to Congress on Server and Data Center Energy Efficiency. Technical report, U.S. Environmental Protection Agency, 2007.

[76] ASHRAE Technical Committees. *Thermal Guidelines for Data Processing Environments*. American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE), 2004.

[77] M. Nakao, H. Hayama, and M. Nishioka. Which Cooling Air Supply System is Better for a High Heat Density Room: Underfloor or Overhead? In *Proc. of International Telecommunications Energy Conference (INTELEC)*, Kyoto, Japan, November 1991.

[78] H. Hayama and M. Nakao. Air Flow Systems for Telecommunications Equipment Rooms. In *International Telecommunications Energy Conference (INTELEC)*, Florence, Italy, October 1989.

[79] H. Hayama and M. Nakao. Airflow Distribution in Telecommunications Equipment Rooms. In *International Telecommunications Energy Conference (INTELEC)*, Orlando, FL, October 1990.

[80] PGnE. High Performance Data Centers: A Design Guidelines Sourcebook, 2006. http://hightech.lbl.gov/documents/data_centers/06_datacenters-pge.pdf.

[81] R.R. Schmidt and E. Cruz. Raised Floor Computer Data Center: Effect on Rack Inlet Temperatures of Exiting both the Hot and Cold Aisle. In *Proc. of Intersociety Conference on Thermal Phenomena in Electronic Systems (ITHERM)*, San Diego, CA, August 2002.

[82] R.R. Schmidt, K.C. Karki, and S.V. Patankar. Raised Floor Computer Data Center: Perforated Tile Flow Rates for Various Tile Layouts. In *Proc. of Intersociety Conference on Thermal Phenomena in Electronic Systems (ITHERM)*, Las Vegas, NV, June 2004.

[83] R. R. Schmidt. Thermal Profile of a High-Density Data Center-Methodology to Thermally Characterize a Data Center. *American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) Transactions*, 110(2):635–642, 2004.

[84] R. Subrata, A.Y. Zomaya, and B. Landfeldt. Cooperative power-aware scheduling in grid computing environments. *Journal of Parallel and Distributed Computing*, 70(2):84–91, 2010.

[85] S.U. Khan and I. Ahmad. A Cooperative Game Theoretical Technique for Joint Optimization of Energy Consumption and Response Time in Computational Grids. *Parallel and Distributed Systems, IEEE Transactions on*, 20(3):346–360, 2009.

[86] A.J. Shan and N. Krishnan. *Flow Resistance: A Design Guide for Engineers.* Hemisphere Publishing, 1989.

[87] G. Chandra, P. Kapur, and K.C. Saraswat. Scaling Trends For The On Chip Power Dissipation. In *Proc. of IEEE Interconnect Technology Conference (IITC)*, pages 170–172, Burlingame, CA, June 2002.

[88] Advanced configuration & power interface (acpi). Technical report, 2009.

[89] I. Rodero, J. Jaramillo, A. Quiroz, M. Parashar, F. Guim, and S. Poole. Energy-efficient Application-aware Online Provisioning for Virtualized Clouds and Data Centers. In *Proc. of Intl. Green Computing Conference (GREENCOMP)*, Chicago,IL, 2010.

[90] D. S. Johnson Andrew M. R. Garey, R. L. Graham. Resource Constrained Scheduling as Generalized Bin Packing. *Journal of Combinatorial Theory.*

[91] Richard M. Karp, Michael Luby, and A. Marchetti-Spaccamela. A Probabilistic Analysis of Multidimensional Bin Packing Problems. In *Proc. of the sixteenth annual ACM symposium on Theory of computing (STOC '84)*, pages 289–298, New York, NY, 1984. ACM.

[92] K. Maruyama, S. K. Chang, and D. T. Tang. A General Packing Algorithm for Multidimensional Resource Requirements. *International Journal of Parallel Programming*, 6:131–149, 1977.

[93] Ming Zhao and Renato J. Figueiredo. Experimental Study of Virtual Machine Migration in Support of Reservation of Cluster Resources. In *Proc. of the Virtualization Technology in Distributed Computing (VTDC)*, pages 5:1–5:8, Reno, Nevada, 2007.

[94] Senthil Nathan, Purushottam Kulkarni, and Umesh Bellur. Resource Availability Based Performance Benchmarking of Virtual Machine Migrations. In *Proc. of the ACM/SPEC International Conference on Performance Engineering (ICPE)*, pages 387–398, Prague, Czech Republic, 2013.

[95] The RIKEN Integrated Cluster of Clusters (RICC) Log.

[96] D. Feitelson. Parallel Workload Archive, 2010.

[97] Uri Lublin and Dror G. Feitelson. The Workload on Parallel Supercomputers: Modeling the Characteristics of Rigid Jobs. *Journal of Parallel Distributed Computing*, 63(11):1105–1122, November 2003.

[98] Hariharasudhan Viswanathan, Eun Kyung Lee, and Dario Pompili. Self-organizing Sensing Infrastructure for Autonomic Management of Green Datacenters. *IEEE Network*, 25(4):34–40, July 2011.

[99] Stockton Infrared Thermographic Services. Using thermal mapping at the data center. http://www.stocktoninfrared.com/using-thermal-mapping-at-the-data-center/.

[100] Energex Technologies. Datacenter thermal imaging and analysis. http://energextech.com/Data_Center_Thermal.pdf.

[101] Electronic Environments Infrastructure Solutions. Why thermal imaging of your data center infrastructure is important? http://www.eecnet.com/Resources/Articles/Why-Thermal-Imaging-for-Power—Cooling-Infrastruc/.

[102] C. Mazzariello, R. Bifulco, and R. Canonico. Integrating a network ids into an open source cloud computing environment. In *Proc. of Information Assurance and Security (IAS)*, Aug 2010.

[103] T. Alharkan and P. Martin. Idsaas: Intrusion detection system as a service in public clouds. In *Proc. of Cluster, Cloud and Grid Computing (CCGrid)*, May 2012.

[104] Rean Griffith, Gail Kaiser, and Javier Alonso López. Multi-perspective evaluation of self-healing systems using simple probabilistic models. In *Proc. of International Conference on Autonomic Computing (ICAC)*, June 2009.

[105] Javier Alonso, Rivalino Matias, Elder Vicente, Ana Maria, and Kishor S. Trivedi. A comparative experimental study of software rejuvenation overhead. *Performance Evaluation, Elsevier*, pages 231–250, 2013.

[106] M. Marwah, R. Sharma, and C. Bash. Thermal anomaly prediction in data centers. In *Proc. of IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm '10)*, Las Vegas, NV, June 2010.

[107] Robert M. Haralick, K. Shanmugam, and Its'hak Dinstein. Textural features for image classification. *Systems, Man and Cybernetics, IEEE Transactions on*, SMC-3(6):610 –621, Nov. 1973.

[108] Arati S. Kurani, Dong hui Xu, Jacob Furst, and Daniela Stan Raicu. Co-occurrence matrices for volumetric data. the. In *7th IASTED International Conference on Computer Graphics and Imaging, Kauai*, August 2004.

[109] John Kerekes. Receiver Operating Characteristic Curve Confidence Intervals and Regions. *Geoscience and Remote Sensing Letters*, 5(2):251–255, April 2008.