Self-Organizing Sensing Infrastructure for Autonomic Management of Green Datacenters

Hariharasudhan Viswanathan, Eun Kyung Lee, and Dario Pompili, Rutgers University

Abstract

The scale and complexity of modern datacenters are 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 this article, methods for acquiring thermal awareness using real-time measurements and heat and air circulation models as well as solutions for proactive autonomic datacenter management that exploit this awareness are discussed. Novel communication and coordination schemes that enable self-organization of a network of external heterogeneous sensors (e.g., thermal cameras, scalar temperature and humidity sensors, airflow meters) into a multi-tier sensing infrastructure capable of real-time datacenter monitoring are also presented.

Datacenters are a growing component of society's information technology (IT) infrastructure, enabling services related to health, banking, commerce, defense, education, and entertainment. Due to the rise in demand for computing and storage, *energy consumption*, *heat generation*, and *cooling requirements* of datacenters have become critical concerns in terms of both the growing operating costs as well as their environmental and societal impacts. It is predicted that datacenter energy consumption in the United States will reach 100 billion kWh/year by the end of 2011 with a corresponding energy bill of approximately $7.4 billion [1]. Impact on the environment and society may include increase in CO_2 emissions, overload of the electricity supply grid, and rising water usage for cooling leading to water scarcity. The scale and complexity of modern datacenters are growing at an alarming rate due to the rising popularity of the cloud computing paradigm as an effective means to cater to the aforementioned increase in demand [2]. Their management is rapidly exceeding human ability, making *autonomic* (*self-configuration*, *self-optimization*, *self-healing*, and *self-protection*) management approaches essential.

From our feasibility study and proof-of-concept experiments on our testbed [3] at the National Science Foundation (NSF) Center for Autonomic Computing (CAC) — a multi-institutional research center that comprises four university sites: the University of Florida, University of Arizona, Rutgers University, and Mississippi State University — we have inferred that one of the fundamental problems in existing datacenters is the local unevenness in *heat generation* and *heat extraction rates*. The former can be attributed to non-uniform distribution of workloads among servers and the heterogeneity of computing hardware, while the latter can be attributed to 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*. 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 in a particular region will lead to a significant temperature rise, which may increase the risk of equipment overheating and hence the chances of server system failures due to operation in the unsafe temperature range [4]. Thus, *thermal awareness*, which is the knowledge of heat imbalance in different regions inside a datacenter, is essential to maximize energy and cooling efficiency as well as to minimize server system failure rates.

Autonomic datacenter management, which includes *thermal- and energy-aware resource provisioning*, *cooling system optimization*, and *anomaly detection*, can help minimize both the impact on the environment and the total cost of ownership (TCO) of datacenters, making them energy-efficient and green. Autonomic datacenter management solutions require continuous processing andanalysis of real-time feedback. Modern blade servers are equipped with a number of internal sensors that provide information about server subsystem operating temperatures and utilization. Similarly, state-of-the-art programmable network elements such as OpenFlow switches [5] provide information regarding network traffic. However, the communication and computation overhead involved in the real-time collection and processing of all the aforementioned raw data would be huge. In addition, the extracted information cannot capture the complex thermodynamic phenomena of heat and air circulation inside a datacenter, and is unreliable in the event of security attacks or system failures.

This work was supported by the National Science Foundation under Grant No. CNS-0855091.

For these reasons, in this article 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. 1) — into a multitier 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*), thus enabling efficient real-time monitoring of datacenters. Incorporation of real-time measurements (temperature, airflow 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 [6, 7] in datacenters can leverage this thermal awareness. Our novel multitier approach eliminates the need for collecting and processing large volumes of data from the entire datacenter. Instead, it enables us to *zoom in* only on problem areas like undesired thermal hotspots (higher temperature regions), and to collect and analyze thoroughly the most relevant data from multiple sources (i.e., server and network element logs, internal and external sensors). This approach drastically improves the responsiveness of algorithms for time-critical management applications such as autonomic thermal-aware anomaly detection and classification.

In the next section, we present our solutions for autonomic adaptive sampling and thermal image compression, which enable efficient continuous real-time monitoring of datacenters. We then describe two datacenter management applications, thermal-aware virtual machine (VM) allocation and thermal-aware anomaly detection, and elaborate on how they exploit the information extracted from the sensing infrastructure before concluding the article in the final section.

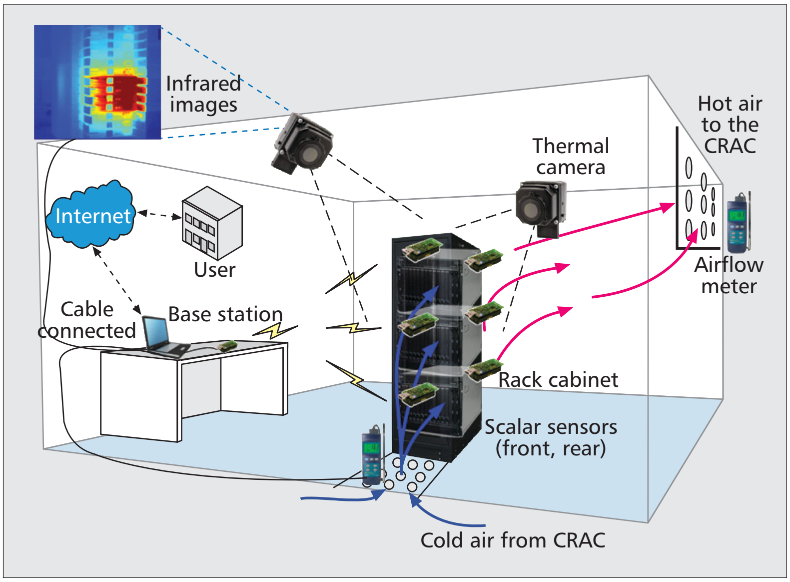

Figure 1. *Network of heterogeneous sensors (scalar temperature and humidity sensors, airflow meters, and thermal cameras) monitoring servers.*

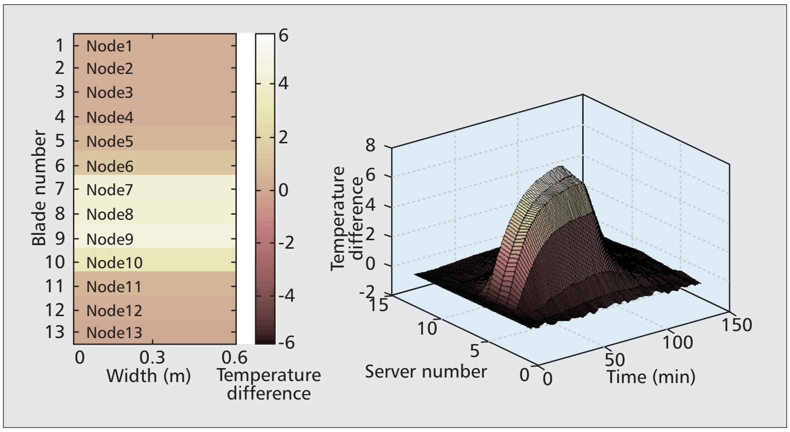

Figure 2. *Illustration of spatial (left) and temporal (right) correlation of temperature measurements, which are represented in relative terms as difference from a reference value of 23°C.*

Autonomic Sensing Infrastructure

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 (50,000 in total). The amount of sensed information collected and processed every second at a monitor node would be on the order of gigabits, thus increasing the strain on the communication and computation resources. We have designed two innovative solutions, *autonomic adaptive sampling* and *coordinated hotspot detection and localization*, for enabling self-organization of these heterogeneous sensors into an intelligent multitier 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.

Autonomic Adaptive Sampling

To understand the opportunities for adaptive sampling in a datacenter monitoring network, consider the following scenarios. If the sensed temperature values in a region inside the datacenter are highly correlated (e.g., in a rack in which workloads have been consolidated), the sensors can coordinate among themselves to eliminate redundancy in the data reported to a monitor node. Also, the temporal sampling and reporting rate of sensors can be lowered if the monitored phenomenon does not change rapidly over time (e.g., in idle servers). However, adaptive sampling should not affect the accuracy of reconstruction of the phenomenon at the monitor node as it will have an impact on the estimation of heat generation and extraction rates. Figure 2 (left) shows the temper-

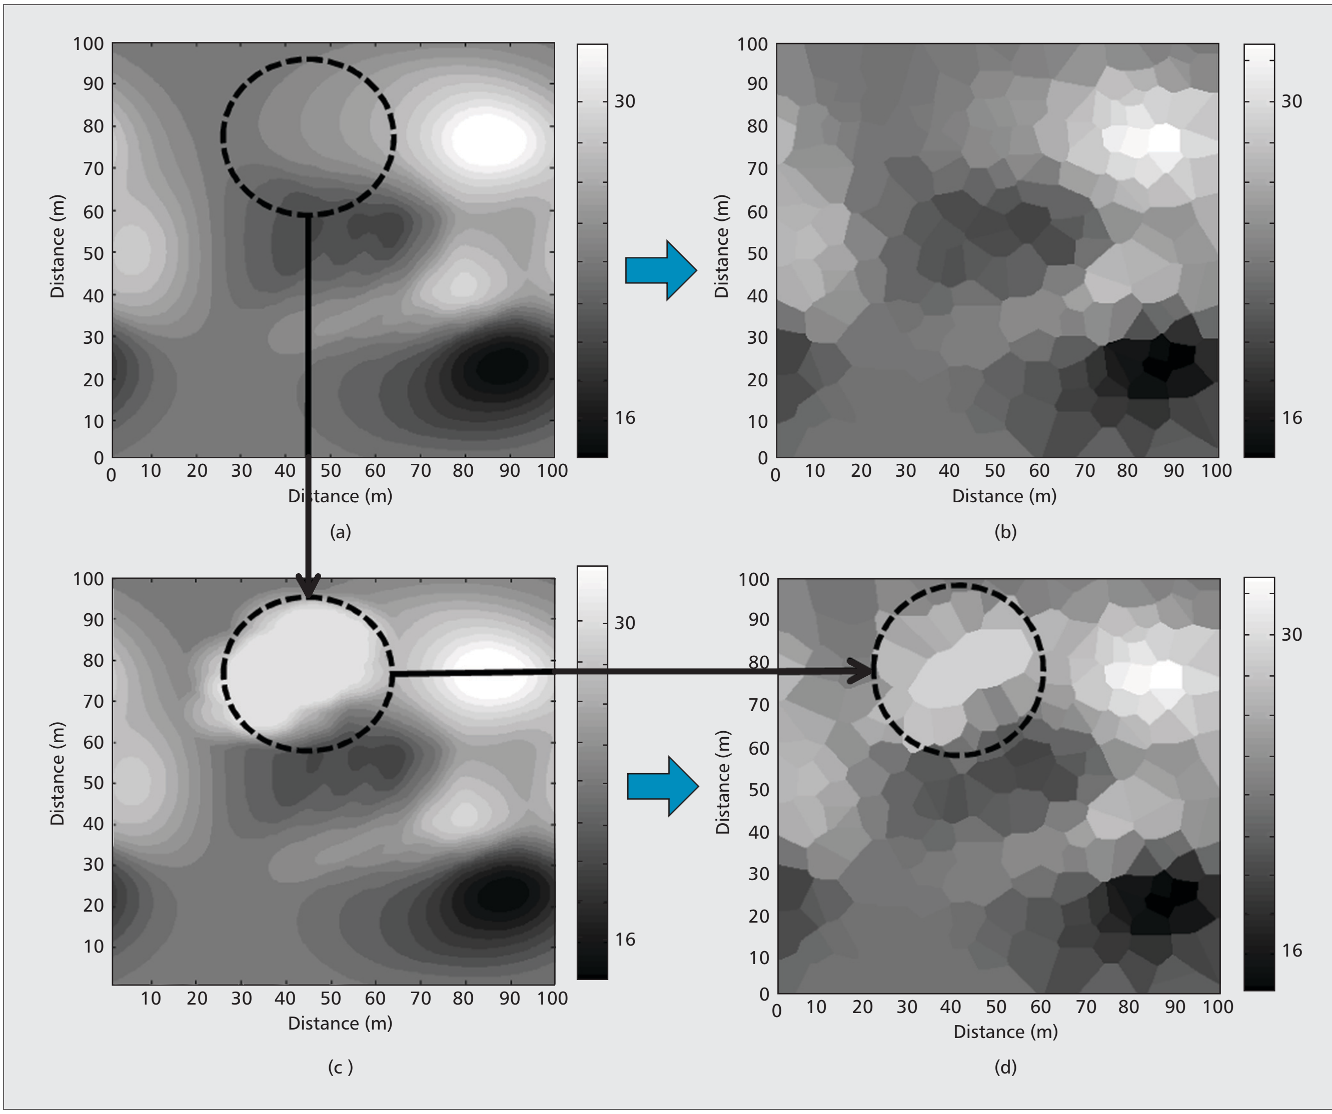

Figure 3. *a) Spatial distribution of a manifestation (temperature) in a 2D field; b) reconstructed field with data reported only by 210 REPs out of 400 nodes; c) variation in spatial distribution of the manifestation (area represented by dashed circle); d) reconstruction of the modified field with data from 218 REPs after some ASSOCs track the variation in the phenomenon.*

ature distribution in our laboratory testbed (a rack of 13 horizontally mounted servers) when only servers 7, 8, 9, and 10 are running workloads. The lighter the shade of the color, the higher the temperature. Temperature is represented in relative terms as difference from a reference value of 23°C. The spatial and temporal correlation among the measured temperature values at sensor nodes placed near the outlet fan of every server are also depicted in Fig. 2 (right).

Our novel adaptive sampling solution [8] enables scalar point-source sensor nodes in the vicinity to exploit the spatiotemporal correlation in measurements, and organize themselves into representative (REP) and associate (ASSOC) nodes. This self-organization is aimed at eliminating data redundancy and prioritizing data based on the value of information it possesses. A node's worthiness to be a REP for a set of ASSOC nodes is based on *similarity* and *correlation* of sensed data. Two nodes, i and j, are said to be sensing similar values if the difference between the means of their measured values, $e^k_{i,j}$, of a manifestation k of a phenomenon (e.g., temperature is a manifestation of the phenomenon of heat circulation) is less than a user-specified threshold, e^k_{th}. Measured values of manifestation k at i and j are said to be correlated if the correlation coefficient, $\rho^k_{i,j}$, which is calculated using a small finite number of samples from those nodes, is greater than a user-specified threshold, ρ^k_{th}.

If nodes i and j are sensing *similar and correlated* values, i is a potential ASSOC of j and vice versa. Actual ASSOCs and REPs are finalized after a localized election mechanism aimed at choosing the best set of REPs that have the maximum number of potential ASSOCs. The thresholds e^k_{th} and ρ^k_{th} may vary at different regions of the sensor network (i.e., the datacenter), and are provided by the datacenter management solutions like anomaly detection and classification mechanisms to either zoom in or zoom out for higher or lower resolution, respectively. The distribution of temperature in a simulated 100 × 100 m^2 field and its successful remote reconstruction (at a monitor node) with data obtained from the nodes of a sensor-based system that employs the proposed adaptive sampling scheme are shown in Fig. 3.

Detection and Localization of Hotspots

Our proposed heterogeneous sensing infrastructure also includes 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 through *remote sensing.* In addition to the significant reduction

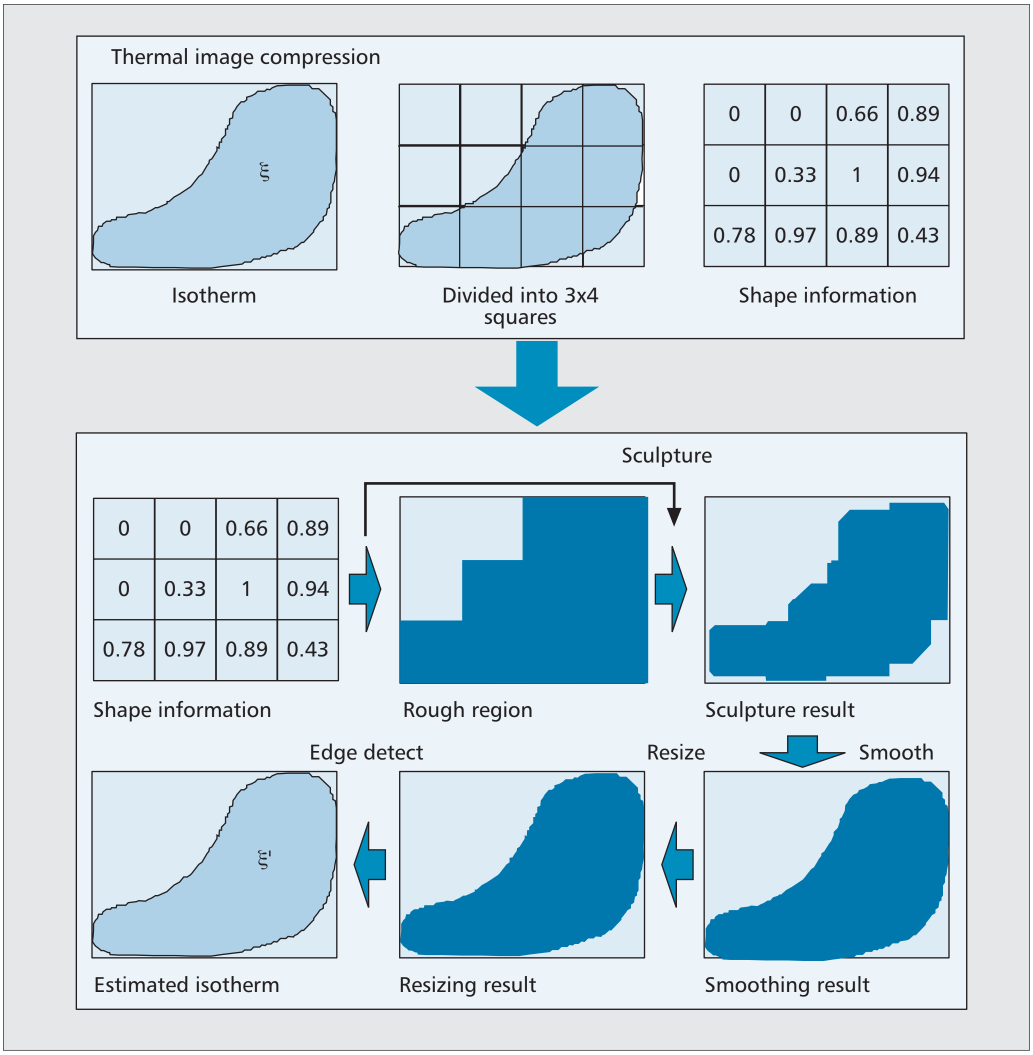

Figure 4. *Sculpturing, a novel technique for compression (top) of hotspot information in a 2D thermal image and reconstruction (bottom) of the same.*

in the number of data sources, using a few thermal cameras (approximately $2500 each) is cost effective in comparison to tens of thousands of scalar sensors (approximately $25 each). A thermal camera uses infrared radiation from objects in its field of view to create thermal images. The FLIR PathFinder thermal cameras (with multiple uncooled microbolometer detectors) we use in our testbed [3] can capture the surface temperature of objects in the range of –40 to 80°C. Apart from providing (remotely) point temperature values at a very high granularity, thermal cameras also help identify and characterize thermal hotspots more efficiently than a group of scalar sensors.

Multiple thermal cameras with pan and tilt capabilities are required to ensure non-overlapping coverage of a 3D region of interest (a number of racks in a datacenter) and to quickly detect thermal hotspots. However, for accurate 3D localization of thermal hotspots, coordination among these thermal cameras is essential as *ranging* (estimating the distance to a hotspot) cannot be done with a single 2D thermal image. Thermal cameras need to share the individual information about observed hotspot shape, position, and intensity for detection and localization of hotspots in 3D space.

Continuous exchange of actual 2D thermal images among multiple thermal cameras is prohibitive in terms of network bandwidth usage. We have designed an innovative data compression solution, *sculpturing*, which exploits specific features of thermal images of hotspots to enable efficient exchange of detected hotspot information among thermal cameras. Sculpturing uses *isotherm lines* to represent the border of a hotspot in a thermal image. Figure 4 illustrates the steps involved in sculpturing. A hotspot (ξ) is surrounded by an isotherm line l. First, a binary image I of the isotherm line from a thermal image is obtained and divided into $N \times M$ squares. In every square, the percentage of the region belonging to ξ is calculated; then this information is sent to the receiver along with the actual coordinates (in the thermal image) of the top left corner of the $N \times M$ matrix.

Our proposed sculpturing technique outperforms traditional image compression schemes like JPEG for thermal images, when the number of hotspots and hence the number of isotherms to be represented are low (which is usually the case as we are interested only in high temperature regions). In our simulations, sculpturing achieved a higher compression efficiency than JPEG (approximately 330 percent), when the number of isotherms to be represented by the sculpturing algorithm was set to three. When the number of isotherms is increased (for a higher level of detail), sculpturing loses its advantage. However, this problem can be overcome by using image segmentation so that each segment in turn has fewer hotspots and hence isotherm lines.

The receiver will use these percentage numbers to sculpt and reconstruct the 2D thermal image of the hotspot. Multi-

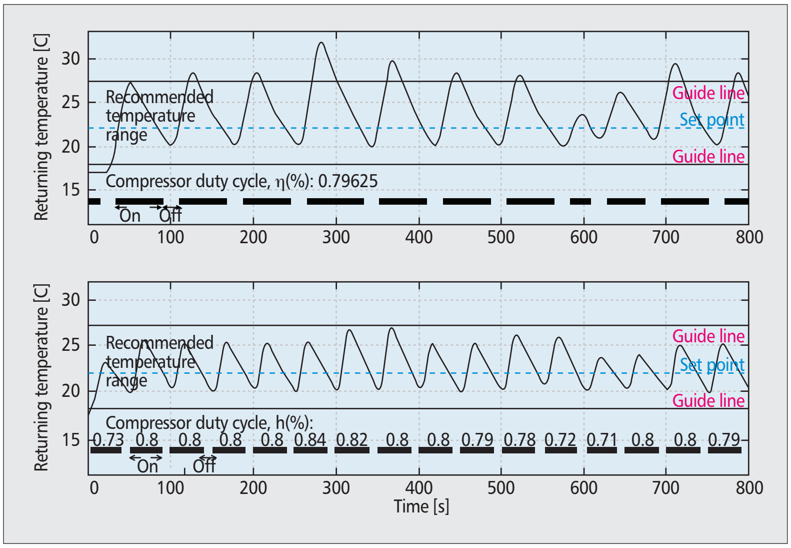

Figure 5. *Temperature variation in a datacenter whose servers are running HPL Linpack workload with reactive cooling (top) and proactive cooling (bottom).*

ple reconstructed 2D images of a hotspot are used to extract specific features such as its shape, area, centroid, and intensity. These features are then used to cooperatively localize hotspots. Thermal anomalies can be detected on the fly by comparing the thermal map recreated using these features against the predicted temperature distribution map (from heat imbalance estimates). If detected hotspots have not been predicted beforehand by the heat imbalance model, they are treated as thermal anomalies.

Thermal-Aware Resource Provisioning

In [7], we introduced the novel concept of heat imbalance estimation, which involves incorporation of measurements from the autonomic sensing infrastructure into heat generation and extraction rate models. This estimation allows prediction of future temperature trends at various regions inside a datacenter, which enables the design of proactive thermal-aware resource provisioning solutions. In virtualized HPC datacenters, one or more VMs are created for every workload or job request, and each VM is provisioned with resources that sufficiently satisfy the workload QoS requirements, which are based on service level agreements (SLAs). Once VMs are provisioned, they have to be allocated to physical servers.

VM Allocation Problem

We propose an energy-efficient strategy for VM allocation, which relies on *thermal-aware VM consolidation*. Minimizing the number of servers that are in operation (for a given set of VMs) will help reduce the energy overhead and hence the total energy consumption of a datacenter. Therefore, the goal is to pack VMs into the minimum possible number of servers, while simultaneously ensuring that the QoS requirements for each VM (or workload) in terms of processing power, memory, and communication bandwidth are met.

The total energy consumption in a datacenter can be split into energy consumption for computation (i.e., running the workloads [or VMs] on servers) and energy consumption for cooling. When the cooling system parameters (fan speed and compressor duty cycle) are fixed (i.e., the energy consumption for cooling is fixed), the only way to achieve energy efficiency is by minimizing the energy for computation through VM consolidation. VM consolidation also restricts heat generation to a smaller area inside the datacenter (instead of having the heat dispersed over a wide area), thus enabling efficient heat extraction and a higher coefficient of performance (COP) of cooling, which is the ratio of the amount of heat extracted to the work done to extract the heat. To achieve this, we formulate the VM allocation problem (VMAP) as a *variable size multidimensional bin packing problem*. This is a generalized version of the traditional fixed-size one-dimensional bin packing problem as the *bins (servers)* and *objects (VMs)* are represented by multiple dimensions (five 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 heat generation rate and the four subsystem (CPU, memory, disk storage, and network interface) utilization requirements, and they are computed by taking the corresponding workload characteristics into account. However, determination of the size (maximum capacity) of a bin or server along the five different dimensions is not straightforward. We define the maximum capacity along the first dimension as the maximum amount of heat that can be generated over a window of time while still ensuring that the server operating temperature is within the recommended range. Maximum capacity along the next three dimensions (CPU, memory, and storage) depends on the individual server configuration. Maximum capacity in terms of network bandwidth, however, is tricky to determine as it depends not only on the capacity of the individual server's network interfaces but also on the capacity of the top-of-the-rack switch. To determine accurately the capacity along this fifth dimension, we assume the use of fully-programmable OpenFlow switches [5].

Constraints for VMAP are chosen in order to ensure that:
- A VM is allocated to *one* and *only one* server.
- The QoS requirements of all VMs packed into one server are met; that is, the maximum utilization of a server subsystem is not exceeded.
- The heat imbalance is neither too high nor too low so that the server operating temperature remains within the recommended temperature range.

Multidimensional bin packing is an NP-hard problem that has been studied in detail, and many efficient approximation algorithms have been proposed for the same, as in [9]. VMAP has to be performed periodically on a set of VMs that arrive in a window of time. Extension of VMAP to include VM migrations involves usage of residual capacities for bin sizes (as some bins may be partly filled) and modification of the objective function to capture the additional cost of migration (extra energy consumption).

VMAP for Federated Datacenters

We extend VMAP to optimize resource provisioning across a network of heterogeneous datacenters, which is usually the case in a cloud infrastructure. Heterogeneity here refers to the difference in characteristics and capabilities of computing (e.g., heat generation rate of servers, processing power, network capacity) and cooling (e.g., efficiency of air-chilled vs. water-chilled) 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 water temperature increase caused by cooling systems). We propose two different approaches to solve VMAP for federated datacenters. The first is a two-step

Workload Type	Reactive (temperature-based)		Proactive (heat-imbalance-based)	
	Risk (%)	Energy (kWh)	Risk	Energy
FFTW	0.00	7.92	0.00	7.27
NAS	0.58	15.45	0.05	14.22
HPL	16.35	20.09	3.71	16.19

Table 1. *Risk of overheating and energy consumption for two different cooling system optimization techniques under different workload conditions.*

(generally suboptimal) approach in which the problems of deciding which datacenter should handle the VM and which physical server should host the VM (as in VMAP) are determined sequentially. The second is the more challenging approach in which the problem captures the characteristics and capabilities of all datacenters and reaches a globally optimal solution after considering all possible trade-offs. For example, if reducing the CO_2 footprint and the aggregate TCO are the goals, the solution should load datacenters that rely on renewable sources of energy as long as the following conditions are met: compliance with QoS requirements of workloads and with environmental regulations such as cap on water consumption and on water temperature increase caused by the cooling system.

Significant gains in terms of reduced energy cost and reduced hardware system failure rate can be obtained through proactive thermal-aware datacenter management as demonstrated in our prior work on cooling system optimization [7]. We have shown through simulations and experimentation that heat-imbalance-based proactive datacenter management is superior in terms of energy-efficiency and minimization of risk of equipment failures compared to its conventional temperature measurement-based reactive counterpart. Figure 5 shows the temperature variation inside a datacenter when reactive (top) and proactive (bottom) approaches to cooling system optimization are employed. It can be clearly seen that the proactive approach does not allow the servers to operate outside of the recommended temperature range, thus significantly reducing the risk of hardware failures due to overheating. Table 1 shows the percentage of servers under increased risk of failure and the energy consumption for the two different cooling system optimization techniques under different workload conditions.

Anomaly Detection and Classification

Unexpected changes in the local heat generation and extraction rates due to cooling equipment failures, misconfigurations, and attacks (e.g., illegitimate workloads) may over time cause large heat imbalances. If the heat imbalance is positive, it will create an unexpected thermal hotspot. Thermal hotspots may also result in a thermal fugue, which is characterized by a continuous increase in the rate of temperature rise. 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. Hence, novel autonomic approaches for online anomaly detection using thermal cameras and for identification of the causes of thermal anomalies are required so that their direct impact on cooling efficiency and server lifetime are minimized.

Detection of Anomalies

We propose an online autonomic thermal anomaly detection method (depicted in Fig. 6) that leverages the novel notion of the *thermal signature* of a datacenter. The heat imbalance model is used to estimate approximate intensities and distribution of expected hotspots for a specific workload (or VM) distribution (the datacenter's unique thermal signature). This estimation will incur only a marginal additional cost in terms of computation if the external sensor data is also used in the heat-imbalance-based proactive resource provisioning (say, VM provisioning). To detect autonomically the unexpected hotspots, the thermal signature is periodically compared against *features* extracted from thermal images (location, size, and intensity of hotspots) obtained using thermal cameras, which are part of the multitier sensing infrastructure.

The fundamental research challenge we had to address was the minimization of the amount of information needed from the multitier sensing infrastructure for detecting anomalies in a timely manner with a very low probability of error (or false alarms). The amount of information from the multitier sensing infrastructure and the accuracy of reconstruction of the underlying phenomena (temperature distribution and air circulation) are directly related. The amount of information obtained from the multitier sensing infrastructure for heat imbalance estimation at different points inside the datacenter (its thermal signature) is determined by the values of e^k_{th} and ρ^k_{th} (input parameters to our autonomic adaptive sampling solution). Also, accurate detection and "sculpturing" of hotspots from thermal images require specification of a temperature range of interest, which determines the number of different isotherm lines that will be extracted and hence the amount of information obtained from a thermal image (the raw data). Our autonomic thermal-aware anomaly detection mechanism chooses on the fly the best values for the aforementioned inputs to the adaptive sampling and thermal image

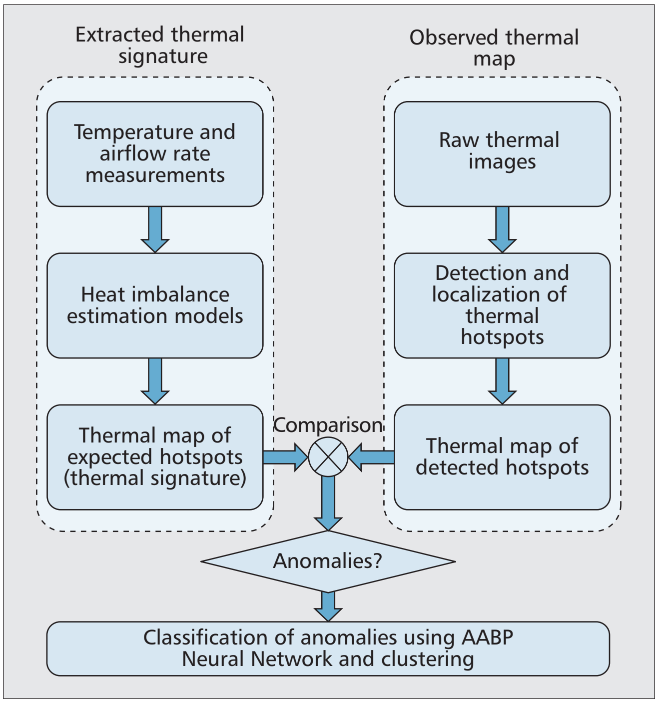

Figure 6. *Thermal-aware anomaly detection: extraction and comparison of the thermal signature of a datacenter with the observed thermal map.*

compression algorithms to strike a balance between the amount of raw data that needs to be processed (communication and computation overhead) and the probability of false alarms. Once an anomaly is identified, the adaptive sampling thresholds are manipulated to zoom in on the areas of interest to obtain external sensor data at a higher resolution, server subsystem utilization logs, and network traffic logs from the network elements for a thorough analysis of the cause of the anomaly.

Classification of Anomalies

It is necessary to pinpoint (or classify) the causes of the detected thermal anomaly to enable appropriate remedial action. Anomalies may span the entire datacenter (or multiple racks of servers), or be localized and span only a few racks (or servers). Problems with the cooling system such as inlet or exhaust fan failures, compressor failure, increase in cold water temperature, and/or broken cold water pipes usually lead to large-scale anomalies manifesting as a number of spatially distributed unexpected hotspots. On the contrary, misconfigurations, CPU fan failures, and illegitimate workloads running on servers (security attacks) result in fewer localized unexpected thermal hotspots.

Even though data from numerous sources is available for analysis (after zooming in), identifying the causes of the anomalies is not straightforward. *Multi-dimensional analysis* of the features extracted from the raw data is needed for the identification of the causes of anomalies. Supervised learning or semi-supervised learning techniques cannot be easily employed as the training data set needs to cover all possible normal events and rare anomalous events. These events are characterized by different combinations of feature values, which increase dramatically with increase in number of features. Hence, we advocate unsupervised learning strategies — specifically, the Auto-Associative Back Propagation (AABP) neural network [10], which involves training with unlabeled data — as it is more suited for classification of anomalies based on their causes. The AABP neural network operates as a smart compression operator and enables a compact representation of multidimensional data into smaller dimensions (2D or 3D) so that anomalies can be grouped into numerous clusters based on their causes.

Conclusion

We have presented innovative solutions for autonomic adaptive sampling and coordinated detection and localization of thermal hotspots. We have discussed in detail how these solutions can be effectively utilized in two important autonomic datacenter management problems, thermal- and energy-aware resource provisioning and thermal-aware anomaly detection. Our approach represents a transformative shift toward cross-layer autonomics for datacenter management problems, which have so far been considered mostly in terms of individual layers (application, virtualization, physical resource, and environment).

References

[1] "Report to Congress on Server and Data Center Energy Efficiency," U.S. Environmental Protection Agency, Tech. Rep., August 2007.

[2] A. Berl *et al.*, "Energy-Efficient Cloud Computing," *Comp. J.*, vol. 53, no. 7, 2010, pp. 1045-51.

[3] "Demo: Temperature and Heat Profiling in Green Datacenters," CPS Lab, Rutgers Univ., http://www.nsfcac.rutgers.edu/CPS/demo.html

[4] J. Srinivasan *et al.*, "The Impact of Technology Scaling on Lifetime Reliability," *Proc. Int'l. Conf. on Dependable Sys. and Networks*, Florence, Italy, June 2004.

[5] "The OpenFlow Switch Consortium", http://www.openflowswitch.org/

[6] "Thermo-Fluids Provisioning of a High Performance High Density Datacenter," http://www.hpl.hp.com/techreports/2004/HPL-2004-146R1.pdf

[7] E. K. Lee *et al.*, "Proactive Thermal Management in Green Datacenter," *J. Supercomputing*, June 2010, pp. 1-31.

[8] E. Lee, H. Viswanathan, and D. Pompili, "SILENCE: Distributed Adaptive Sampling for Autonomic Sensor-based Systems," *Proc. IEEE/ACM Conf. on Autonomic Computing*, Karlsruhe, Germany, June 2011.

[9] N. Bansal, A. Caprara, and M. Sviridenko, "Improved Approximation Algorithms for Multidimensional Bin Packing Problems," *Proc. IEEE Symp. on Foundations of Comp. Sci.*, Berkeley, CA, Oct. 2006.

[10] A. Herrero *et al.*, "Intrusion Detection at Packet Level by Unsupervised Architectures," *Proc. Int'l. Conf. on Intelligent Data Eng. and Automated Learning*, Birmingham, UK, Dec. 2007.

Biographies

HARIHARASUDHAN VISWANATHAN (hari_viswanathan@cac.rutgers.edu) started his Ph.D. program at the Department of Electrical and Computer Engineering (ECE), Rutgers University, in 2009. Currently, he is pursuing research in the fields of datacenter management, green computing, and cognitive radio emergency networking under the guidance of Dr. Dario Pompili at the NSF Center for Autonomic Computing (CAC). Previously, he received his B.S. in ECE from PSG College of Technology, India, and M.S. in ECE from Rutgers University in 2006 and 2009, respectively.

EUN KYUNG LEE (eunkyung_lee@cac.rutgers.edu) joined the Ph.D. program at the Department of Electrical and Computer Engineering (ECE), Rutgers University, in 2009. He is currently under the guidance of Dr. Dario Pompili as a member of the Center for Autonomic Computing (CAC), an NSF research center. His research interests include energy efficient datacenter management and wireless sensor networks. Previously, he received his B.S. in electrical and telecommunications engineering from Soongsil University, South Korea, and M.S. in ECE from Rutgers University in 2002 and 2009, respectively.

DARIO POMPILI (pompili@cac.rutgers.edu) joined the Department of Electrical and Computer Engineering (ECE) at Rutgers University as an assistant professor in Fall 2007. He received his Ph.D. in ECE from the Georgia Institute of Technology in June 2007. He had previously received his Laurea (integrated B.S. and M.S.) and doctorate degrees in telecommunications and system engineering from the University of Rome "La Sapienza," Italy, in 2001 and 2004, respectively. In 2011 he won the NSF CAREER award for his work on underwater multimedia acoustic communication.