Analysis of Environmental Data in Data Centers

Lola Bautista, Ratnesh Sharma
Enterprise Systems and Software Laboratory
HP Laboratories Palo Alto
HPL-2007-98
June 19, 2007*

smart-data center, dynamic smart cooling, thermal management, exploratory data analysis, principal components analysis, data mining

Significant work has been conducted in the past to understand the thermodynamic variables and parameters that influence the environmental conditions of the rooms vis-a-vis the computer room air conditioning (CRAC) units, racks and servers. A considerable amount of data had been collected from environmental sensors located at various locations within data centers, measuring the supply and return air temperatures in the CRAC units and both inlet and outlet temperatures in the racks.

This work describes the analysis done to discover and classify trends and patterns and relationships within temperatures and air flow data in a data center. Initially exploratory data analysis (EDA) techniques were used for reduction of data, visualization of deterministic behavior and identification of normal or abnormal environmental behavior in the control process. Principal Components Analysis (PCA) was used to capture the variables that contain the information about the influence of CRAC units over the racks.

* Internal Accession Date Only Approved for External Publication

Analysis of Environmental Data in Data Centers

Lola Bautista
Hewlett-Packard Laboratories
Palo Alto, CA 94304
lola.bautista@hp.com

Ratnesh Sharma
Hewlett-Packard Laboratories
Palo Alto, CA 94304
ratnesh.sharma@hp.com

Abstract

Optimization of data center performance for reduction in power and cooling costs is crucial to the growth of IT services both in enterprise and consumer sector. Significant work has been conducted in the past to understand the thermodynamic variables and parameters that influence the environmental conditions of the rooms vis-a-vis the computer room air conditioning (CRAC) units, racks and servers. A considerable amount of data had been collected from environmental sensors located at various locations within data centers, measuring the supply and return air temperatures in the CRAC units and both inlet and outlet temperatures in the racks.

This work describes the analysis of environmental data collected from such sensors to discover and classify trends and patterns and relationships within temperatures and air flow data in a data center. Data mining techniques have been used to extract knowledge from historical data to improve the dynamic monitoring and control of data center power and cooling resources. Initially exploratory data analysis (EDA) techniques were used for reduction of data, visualization of deterministic behavior and identification of normal or abnormal environmental behavior in the control process. Since the most important part of the analysis is to understand the influence of the CRAC units over the racks, Principal Components Analysis (PCA) was used to capture the variables that contain most of that information.

1. Introduction

Data center thermal management challenges have been steadily increasing over the past few years due to rack level power density increases resulting from system level compaction [1], [2] and energy demands. Cost effective management of such high power levels in the data center with reliable cooling solutions is essential to support the pervasiveness of computing needs. Nonetheless, energy consumption of data centers has also been severely increased by over-designed air handling systems and rack layouts that allow the hot and cold air streams to mix. Lack of local temperature sensors for monitoring and control has contributed to this gap in knowledge of air flow patterns and thermal management issues in conventional data centers.

With the advent of Smart Cooling [3], local temperature sensing is becoming a key part of datacenter thermal management. In the past, temperature sensing has been crucial for test cases and experimental needs [4]. However, significant research needs to be done to analyze and understand the immense amount of data collected from these sensors during runtime. Expedient evaluation [5] and inferencing is needed at runtime to reap the benefits of an adaptive control system [6], [7]. In this paper, we explore the application of data mining techniques, statistical methods [8], exploratory data analysis [9] techniques, and principal components analysis on temperature data sets collected from production data centers in HP Labs, Palo Alto for development of expert systems and advising solutions for runtime management.

2. Data Center Layout

Figure 1 displays a typical state-of-the-art data center air-conditioning environment with under-floor cool air distribution. Computer room air conditioning (CRAC) units cool the exhaust hot air from the computer racks. Energy consumption in data center cooling comprises work done to distribute the cool air and to extract heat from the hot exhaust air. A refrigerated or chilled water cooling coil in the CRAC unit extracts the heat from the air and cools it within a range of 10°C to 18°C. The flow of chilled water is controlled by an internal mixing valve that operates based on the air temperature measured at the return of the CRAC unit.

The air movers in the CRAC units pressurize the plenum with cool air. The cool air enters the data center through vented tiles located on the raised floor close to the inlet of the racks. Typically the racks are laid out in rows separated with hot and cold aisles as shown. The cold aisles supply cold air to the systems and the hot aisles receives hot air from the systems. A multitude of other equipment layout configurations and non-raised floor infrastructures exist and are applicable to the present study [10].

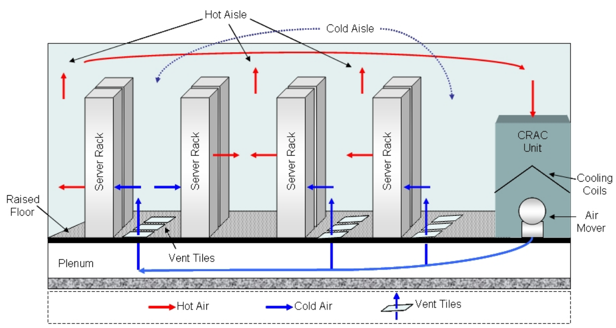

Figure 1. Typical Raised Floor Data Center Configuration

Each one of the racks in the datacenter has attached a set of five sensors at the inlet, and other five at the outlet. For this paper has been considered only the inlet sensors at the front denoted by T1, T2, T3, T4 and T5 from bottom to top (see Figure 2).

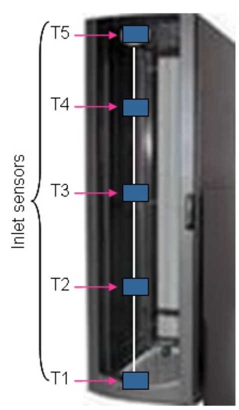

Figure 2. Rack Inlet Temperature Sensors

The rack units are grouped by rows (see Figure 3). Rows A, Aext, B, Bext, C and G were selected in order to cover different locations within the datacenter for the analysis.

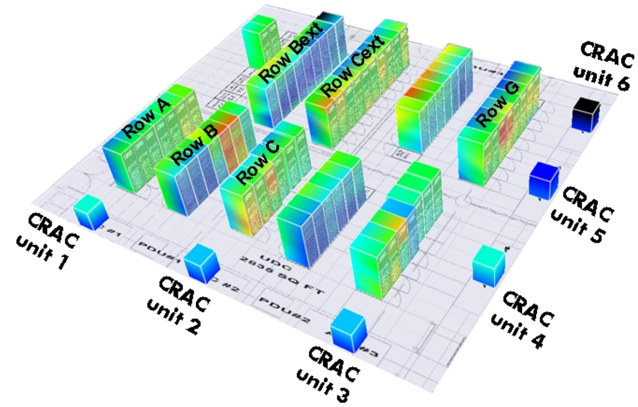

Figure 3. Data Center Layout

To understand the effect the CRAC units on the rack temperature sensors, perturbation experiments were conducted. CRAC unit temperature setpoint and air flow rates were changed to perturb the thermal equilibrium in the datacenter. The goal of the experiments were to better understand the transient fluid mechanics and heat transfer in a dynamic datacenter. Such knowledge can help in identifying potential anomalies and operational issues in datacenter operation [11].

3. SmartCooling Data Mining

A large amount of data had been collected over 2 years from the data centers located at Hewlett-Packard Laboratories running several experiments in order to understand the dynamic interaction existing between CRAC unit temperatures and the corresponding responses over the rack sensors. Automation of knowledge discovery and development of predictive control are some of the key goals of this research effort. Data mining techniques provide methods for reduction of raw data for analysis and tools for extraction of significant patterns for model development.

In this first approach to understand the influence of the CRAC units over the racks was used Exploratory Data Analysis and Principal Components Analysis in order to gain a deep insight into the raw data and to identify the significant variables that are acting over the whole system.

Here are presented the outputs from two of the conducted experiments, the first one with a perturbation length of 30 minutes (Figure 4a) called experiment A, and the second one of 60 minutes (Figure 4b) called experiment B. The plots show the progressive perturbation of CRAC Unit

1 (CU1) through CRAC Unit 6 (CU6). The perturbation of each CU starts when the previous CU is reaching its lowest temperature. For example, at index time 70 (Figure 4a), CU1 is reaching its lowest temperature (see first subplot), and at the same time is starting the perturbation over CU2 (see second subplot), and so on over the other CUs.

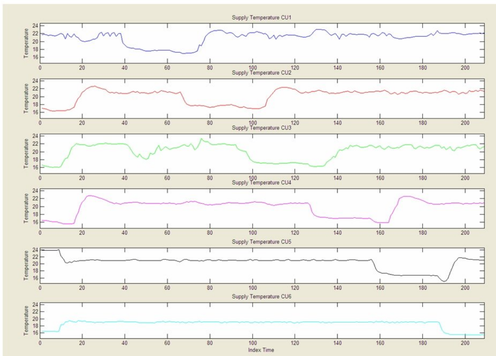

a) Experiment A

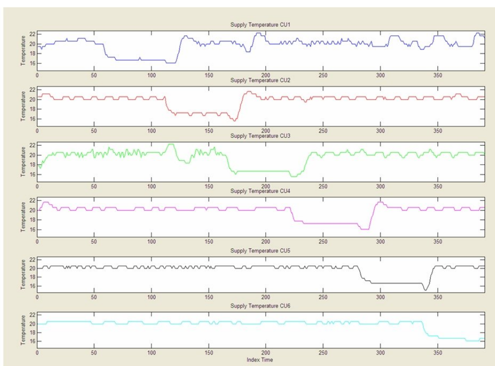

b) Experiment B

Figure 4. Perturbation over CRAC Unit Supply Temperatures

The response of the rack sensors to that perturbations are recorded simultaneously. Figure 5 shows the sensor response of the sensors located in different rack units for the two experiments. Initially, the response of the sensors appears to be identical for different periods of time.

The following subsection describes the results obtained from the Exploratory Data Analysis of this raw data.

3.1. Exploratory Data Analysis (EDA)

EDA enables greater insight into a raw data set. In this work it was used specially to test underlying assumptions such as randomness and distribution fitting. 13 racks are randomly selected from a datacenter for analysis. For the purposes of this paper, outputs from rack A7 sensors are

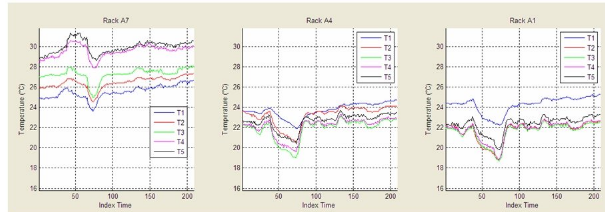

a) Experiment A

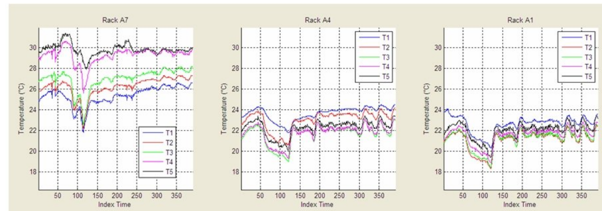

b) Experiment B

Figure 5. Rack Inlets Temperatures

obtained from the two experiments. Figure 6 shows the response of each one of the sensors from that rack for the two experiments.

The first test was conducted to determine the randomness of the data. Using lag-plots as shown in Figure 7, it is possible to visualize a strong relation between the current temperature, denoted Y and the previous, denoted Y_{i-1}. This relation suggests the presence of a deterministic nature of the data, indicating the choice of model that can be built from the data. In this case, an auto-regressive model could be the most suitable for this data.

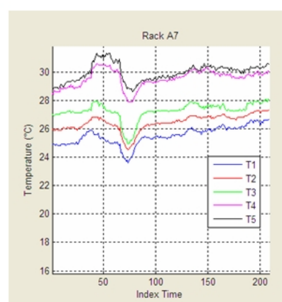

a) Experiment A

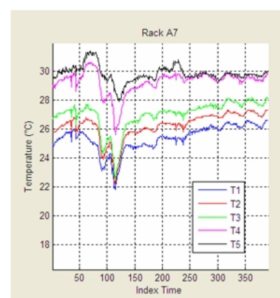

b) Experiment B

Figure 6. Sensor Response Temperatures at Rack A7

Understanding the shape of the distribution is necessary to create the appropriate predictive model. In that context, the second test was conducted to determine the distribution that best fit the data [12]. Figure 7 also shows the normal probability plot for both data sets. These plots in-

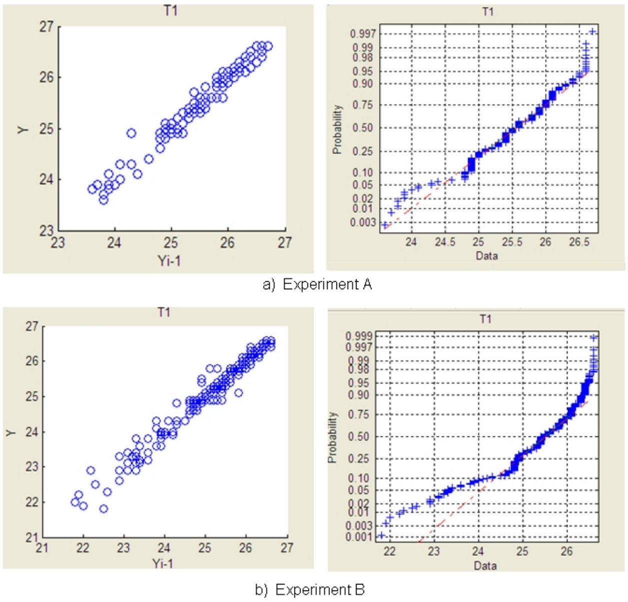

Figure 7. Lag Plot and Normal Probability Plot for Rack A7

dicate that normal distribution does not adequately fit the data. Probability Plot Correlation Coefficient Plot (PPCC) [13] is used to identify the kind of distribution appropriate for the data. The shape parameter λ is obtained using Tukey Lambda PPCC calculator [14]. This parameter indicates whether a distribution is short or long tailed, thus describing the best-fit family distribution for the data. In this case, Tukey Lambda PPCC plots for both data sets, conclude that U-shaped distribution families are appropriate for this purpose. This conclusion is confirmed by looking the histogram of the data set in Figure 8 where it is noticed the presence of low and high extreme values, which indicates the low and high limits of the temperatures within the datacenter.

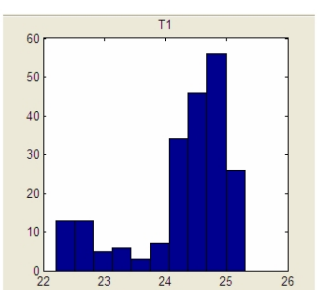

Figure 8. U-shaped Distribution

The relation between air supply temperatures of the CRAC units and the temperature response of the rack units for the two experiments are shown in Figures 9 and 10. The scatter plots show the variation of temperatures of rack A7 with air supply temperatures for all the CRAC units.

The most probable sensor response is shown for a range of CRAC unit air supply temperatures. The gaps corresponds to the abrupt transition between successive perturbations in each CRAC unit. Since the variation in the first data set (experiment A) is higher than that for the second one (experiment B), the data is much more dispersed. Unlike that in the former experiment, the data is more concentrated in each one of the stages of the performed experiment in the later case.

Following analysis of relationship between the CRAC units and the sensors, cross-correlation coefficients were calculated (see Figure 11) for rack A7. The highest correlation coefficient corresponds to CRAC Unit 1 and CRAC Unit 2. This output can be easily explained by the location of the rack as shown in Figure 3. Rack A7 is located in Row A which is closer to CRAC Unit 1 and CRAC Unit 2. CRAC Unit 1 correlations, however, decrease with height of the sensor. The lower sensors (T1 - T3) show a high correlation due to direct influence of air supply from the CRAC Unit 1 through the vent tiles. Sensors T4 and T5 are located close to the top of the rack and are affected by secondary air flows from other directions.

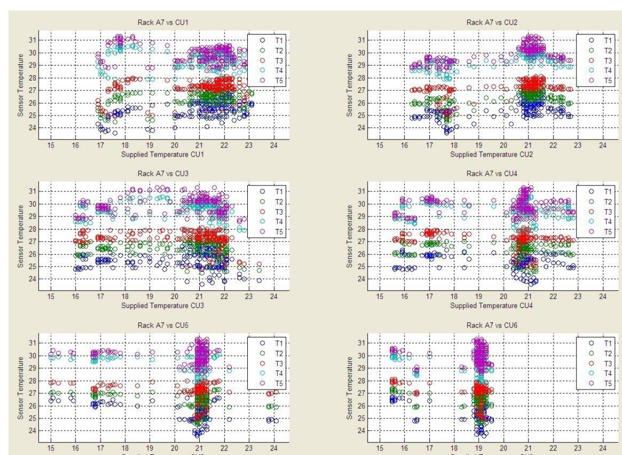

Figure 9. Crac Units Supply Temperatures Vs. Sensor Temperatures (Experiment A)

3.2. Principal Components Analysis (PCA)

Having gathered details about data set through EDA, Principal Component Analysis was conducted to identify

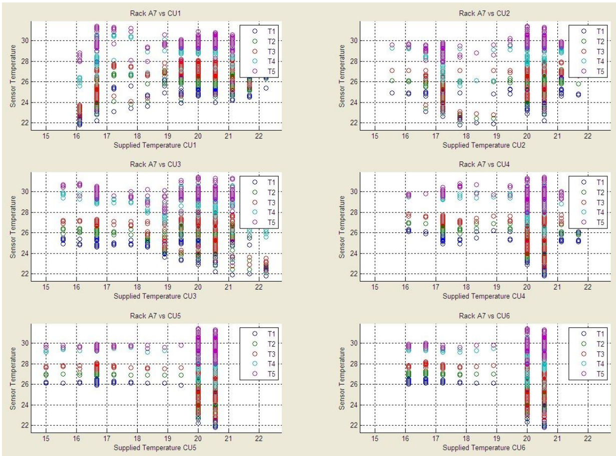

Figure 10. Crac Units Supply Temperatures Vs. Sensor Temperatures (Experiment B)

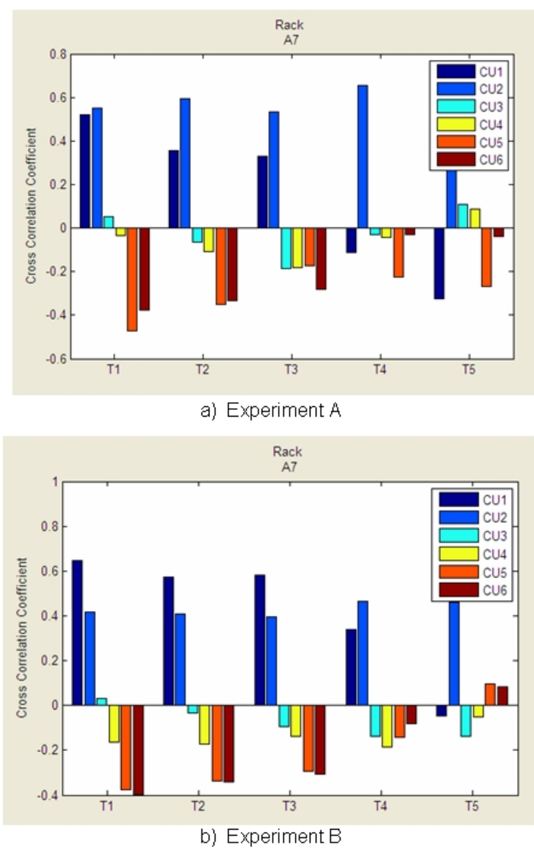

Figure 11. Correlation Coefficients respect to Crac Units Temperatures

the significant influences among the CRAC Units. Based on the inference from the plot of correlation coefficients (see Figure 11), the analysis can be extended to look for the correlated data between the recorded time series for each sensor in each rack unit. PCA is used to reduce dimensionality in the data set, while maintaining the variation present in the data set [15].

Let X be the raw data matrix of dimension $m \times n$, where m corresponds to the total of observations ($t_j, j = 0, ..., m-1$) and n corresponds to the total number of sensors (column vectors $x_l, l = 0, ..., n-1$). Figure 12 illustrates the configuration of the matrix.

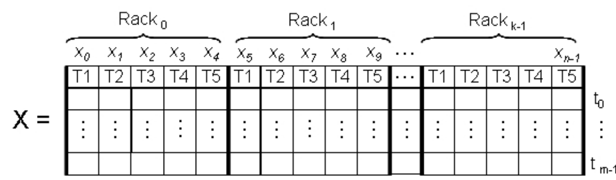

Figure 12. Input Data Matrix

The goal of PCA consists in finding an orthonormal matrix V such that the set of correlated variables X be transformed into a new data set Y of uncorrelated variables with maximum variance. This transformation is presented in Equation (1) as follows:

$$Y = VX \tag{1}$$

The set of vectors $v_i, i = 0, ..., k-1$ forms a new basis for the data set X and are called the *Principal Components* of X. These vectors have the following properties:

1. They are the eigenvectors of C_X, the covariance matrix of X

$$C_X = \frac{1}{m} X X^T \tag{2}$$

   where X^T is the transpose of X

2. The covariance matrix of Y, referred to as C_Y is a diagonalized matrix such as each element C_{Y_i} is the covariance of X along v_i. The elements of C_Y are known as the eigenvalues of C_X

3. The projection of Y into the original space is given by

$$\tilde{X} = YV \tag{3}$$

   It is obtained in such a way that the squared error $\sum_m \left\| X - \tilde{X} \right\|^2$ be minimum.

According to [16], the vector columns of Y are the hidden variables for X. Hidden variables are detected with the occurrence of anomalies in the process. In this case, the CRAC unit perturbations are equivalent to anomalies. The variation in the energy of the signal is captured by the number of hidden variables detected. The occurrence of a change in the energy of the signal increases the number of hidden variables. If the energy remains stable, the number of hidden variables is reduced. Detection of a single hidden variable indicates normal behavior.

A small subset of racks were selected to demonstrate the utility of this technique. The selection was based on knowledge the physical location of the racks to capture inter-row and intra-row (or inter-rack) variations. The first subset corresponds to rack units A1, A4 and A7, and the second subset corresponds to rack units B2 and B7, in order to compare the behavior of the two rows of racks in the data center. Figure 13 illustrates the specific location of each one of those racks within the datacenter.

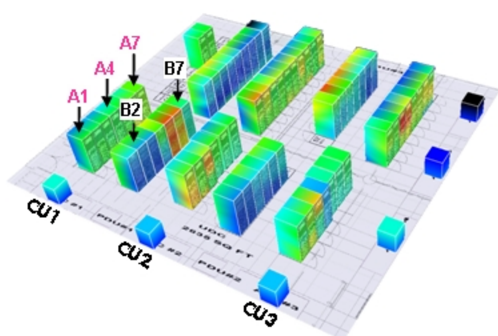

Figure 13. Location of racks within datacenter

The raw data was normalized using Equation (4) before analysis for detecting hidden variables.

$$\bar{X} = (X - \mu)/\sigma \qquad (4)$$

Figure 14 show the supply air temperatures of CRAC Units 1, 2 and 3 in experiment A and Figure 15 for experiment B. The vertical lines over the plots indicate the most significant index time where the number of hidden variables increased. Three runs for detecting hidden variables were executed. Run 1 included rack units A1, A4, and A7 (magenta line). Run 2 included rack units B2 and B7 (cyan line). Run 3 included all five rack units together (green line).

These plots indicate that the increase in the number of hidden variables coincides with the perturbations of

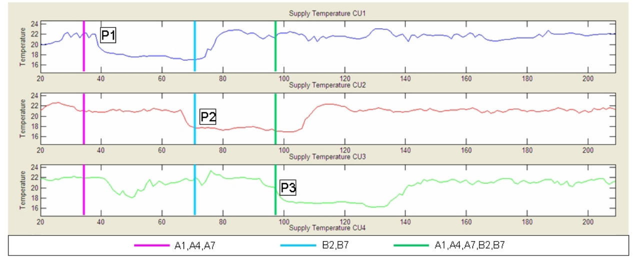

Figure 14. Hidden Variables (Experiment A)

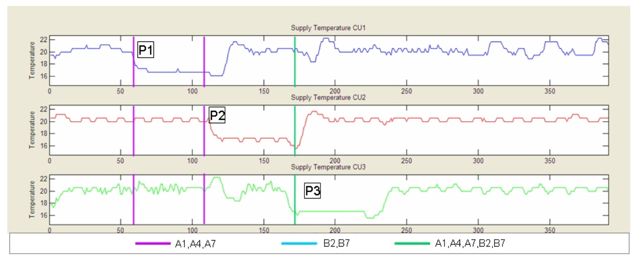

Figure 15. Hidden Variables (Experiment B)

the CRAC Units 1, 2 and 3 labeled as P1, P2 and P3 respectively. In Figure 14, perturbation P1 was detected by run 1, perturbation P2 by run 2, and perturbation P3 by run 3. In Figure 15, perturbations P1 and P2 were detected by run 1 and perturbation P3 by run 3. Looking at the Figure 13, these CRAC units are affecting the response of the rack sensors. The detection of hidden variables took less iterations when all the racks are analyzed together (A1, A4, A7, B2, and B7).

The estimation of time of ocurrence of a perturbation at rack level is very useful to further analysis of the correlations between the CRAC units and the racks. Similar detection can be used to perform change point detections and identify any periodicity in perturbations.

4. Conclusions

With the development of complex IT and facility infrastructure, rising energy costs and evolving data center service scenarios, it is important to understand the need for an expert management system that can manage the life cycle of the data center services. Although composite and advanced numerical models have been developed to design datacenter energy and thermal management infrastructure, none exist for real time management. Techniques like Exploratory Data Analysis and Principal Component Analysis are a necessity for the development of such models in real time. This paper is an attempt to lay the foundation for such an approach.

The current work was focused on reduction of raw data to gain a deep understanding of the influence of CRAC Units over rack sensors. EDA was useful in discovery of the deterministic nature of data and thus guide the building of an auto-regressive model. Creation of such models will help i providing future predictions about the energy and cooling management within the data center.

Identification of hidden variables by using PCA necessitates further analysis of multiple time series in order to determine accurately the location of change points across correlated data. This makes it possible to build a robust model for the thermal management in the data centers as well as to improve its physical deployment of resources. The comparison of two different experiments of different duration helps to estimate the characteristic time periods that can increase the statistical significance of the study.

Although much remains to be done in this area, such an analysis of environmental data is crucial to gather inferences from past performance, to control the present ensemble and to predict (or prevent) the occurrence of critical events within future datacenters.

References

[1] R. Schmidt. Computer and telecommunications equipment room cooling: A review of literature. *Eighth Intersociety Conference on Thermal and Thermomecanical Phenomena in Electronic Systems*, May 2002.

[2] R. Schmidt, E.E. Cruz, and M.K. Iyengar. Challenges of datacenter management. *IBM J Res. Dev.*, 49:709–723, 2005.

[3] Chandrakant Patel. A vision of energy aware computing - from chips to data centers. *The International Symposium on Micro-Mechanical Engineering*, December 2003.

[4] Chandrakant Patel, Cullen Bash, Ratnesth Sharma, Monem Beitelmal, and Rich Friedrich. Smart cooling of datacenters. *Proceedings of IPACK03 - The PacificRim/ASME International Electronis Packaging Technical Conference and Exhibition*, July 2003.

[5] Cullen Bash, Chandrakant Patel, and Ratnesh Sharma. Efficient thermal management of data centers - immediate and long-term research needs. *International Journal HVACR Res.*, 9, April 2003.

[6] J. Moore, J. Chase, P. Ranganathan, and R. Sharma. Making scheduling cool: Temperature-aware workload placement in data centers. *Usenix*, 2004.

[7] Cullen Bash and et. al. Dynamic thermal management of air-cooled datacenter. *ITherm 2006 - Intersociety Conference on Thermal and Thermo-mechanical Phenomena in Electronic Systems*, 2006.

[8] R. Ott and Longnecker M. *An Introduction to Statistical Methods and Data Analysis, 5th Ed.* Duxbury, 2001.

[9] J. Tukey. *Exploratory Data Analysis*. Addison-Wesley, 1977.

[10] Schmidt and et. al. Maintaining datacom rack inlet air temperatures with water cooled heater exchanger. *Proceedings ASME InterPACK 2005*, 2005.

[11] Ratnesh Sharma, Rocky Shih, Chandrakant Patel, and John Sontag. Application of exploratory data analysis (eda) techniques to temperature data in a conventional data center. *Proceedings of IPACK 2007, ASME InterPACK07*, July 2004.

[12] StatSoft. Statsoft electronic textbook. http://www.statsoft.com/textbook/stathome.html.

[13] National Institute of Standards and Technology. Nist/sematech e-handbook of statistical methods. http://www.itl.nist.gov/div898/handbook/.

[14] P. (2007) Free Statistics Software Office for Research Development Wessa and Education version 1.1.21. Tukey lambda ppcc plot calculator. http://www.wessa.net/.

[15] I.T. Jolliffe. *Principal Component Analysis*. Springer, 2002.

[16] Spiros Papadimitriou, Jimeng Sun, and Christos Faloutsos. Streaming pattern discovery in multiple time-series. *Proceedings of the 31st VLDB Conference*, pages 697–708, 2005.