

# Mobile Visual Analytics for Datacenter Power and Cooling Management

Ratnesh Sharma, Ming Hao, Ravigopal Vennelakanti, Manish Gupta,  
Umeshwar Dayal, Cullen Bash, Chandrakant Patel, Deepa Naik, A. Jayakumar,  
Sairabanu Z. Ganihar, Ramesh Munusamy, and Vani Mohan

1501 Page Mill Road

Hewlett-Packard Co. Laboratories, Palo Alto, CA

{ratnesh.sharma, ming.hao, ravigopal.vennelakanti, manish.gupta,  
umeshwar.dayal, cullen.bash, chandrakant.patel, deepa.naik,  
jayakumara, sairabanu-z.ganihar, ramesh.munuswamy, vanim}@hp.com

**Abstract.** The demand for data center solutions with lower total cost of ownership and lower complexity of management is driving the creation of next generation datacenters. The information technology industry is in the midst of a transformation to lower the cost of operation through consolidation and better utilization of critical data center resources. Successful consolidation necessitates increasing utilization of capital intensive “always-on” data center infrastructure, reduction in the recurring cost of power and management of physical resources. In this paper, we describe a tool that allows the data center facility managers and administrators to view and analyze the Key Performance Indicators (KPIs) associated with their data centers using pixel cell-based [10,11] visual analytics. The basic idea of our technique is to use the smallest element in the display to present the detailed information of the poser and thermal data records. Administrators can quickly recognize the patterns, trends, and anomalies. Furthermore, we discuss case studies of mobile visual analytics for energy and thermal state monitoring utilizing data from a rich sensor network.

**Keywords:** Visual Analytics, Sensors, Data center, Thermal Management, Data mining.

## 1 Introduction

The design and operation of the data center infrastructure is one of the primary challenges facing IT organizations. Unprecedented growth in demand for IT services has led to development of large, complex, resource-intensive IT infrastructure to support pervasive computing. Emerging high-density computer systems and consolidation of IT resources into fewer data centers are stretching the limits of data center capacity [1] in terms of power and resource utilization. The large number of components in a data center including cooling systems, power systems, and computer systems and the diversity of these components make data center design and operation complex.

To improve customers' RoIT (Return on Information Technology) [2], it is critical to maximize the resource utilization efficiency of the data center and simplify its management. To accomplish this goal, data center administrator need an agile sensor network coupled with powerful analytics to discover events and plan remedial action. Currently no tools are available to administrators to reveal interactions between data center elements like servers, network switches, and air conditioning units. The problem is compounded with the diversity of IT equipment and the scale of deployment. Since most thermal problems are sensitive to physical location, an onsite deployable tool is of great advantage to administrators. Onsite problem identification and escalation prevention is much easier with a mobile platform.

## 2 Background

### 2.1 Overview

Visual analytics combines various traditional data mining techniques with applications that help to manage datacenter power and cooling performance [3]. The combination of automatic analysis with human perception and understanding is key to this process. Large quantities of raw or processed data may be displayed in various dimensions using color shades or shapes. The change in color or shape enables a human to identify anomalies or events for possible corrective actions. The technique covers wide variety of data including time series. In this paper, we use combination of a pixel cell-based spatio-temporal overall view with high-resolution time series.

We employ visual correlation queries to be performed in an interesting area to meet the challenge of root cause and correlation analysis. The method of our visual correlation queries is to interactively select focus areas in the visualization. Then, according to the characteristics of the selected areas, such as temperature hot spots, we employ data mining techniques to compute relationships to other attributes of the data set. Then, the visual correlation query generates visual representations for users to view and refine the results [4].

For correlation analysis we include the Pearson Correlation  $r$  into our framework [5]:

$$r = \frac{\sum_{i=1}^n (A_{1i} - \bar{A}_1)(A_{2i} - \bar{A}_2)}{\sqrt{\sum_{i=1}^n (A_{1i} - \bar{A}_1)^2 \sum_{i=1}^n (A_{2i} - \bar{A}_2)^2}} \quad (1)$$

where  $n$  represents number of attributes in a sensor data set.

This function computes the pair-wise correlation between bi-variant data with attributes  $A_{1i}$  and  $A_{2i}$ . If two attributes are perfectly correlated, the correlation coefficient is 1; for an inverse correlation, -1.

### 2.2 Data Center Infrastructure

Figure 1 shows the basic data center building blocks from utility grid to the cooling tower [2] [5]. Switch gear comprising of transformers, static switches with associated panels distributes power to the cooling infrastructure and the IT infrastructure. Cooling infrastructure comprises of chillers, cooling towers, computer room air conditioning (CRAC) units and primary/secondary pumps. IT infrastructure includes servers, network devices and storage devices housed in standard racks. UPS maintains power quality during normal operation and provides energy storage to operate the IT infrastructure during brown outs or short power outages. Chillers provide chilled water to the data center room that houses the server racks and other IT equipment.

Figure 2 shows the detail of the datacenter room including CRAC units, server racks and air flow paths [5]. Data centers are typically air-cooled with a raised floor plenum to distribute cool air, power and networking. Figure 2 depicts a typical state-of-the-art data center air-conditioning environment with under-floor cool air distribution.

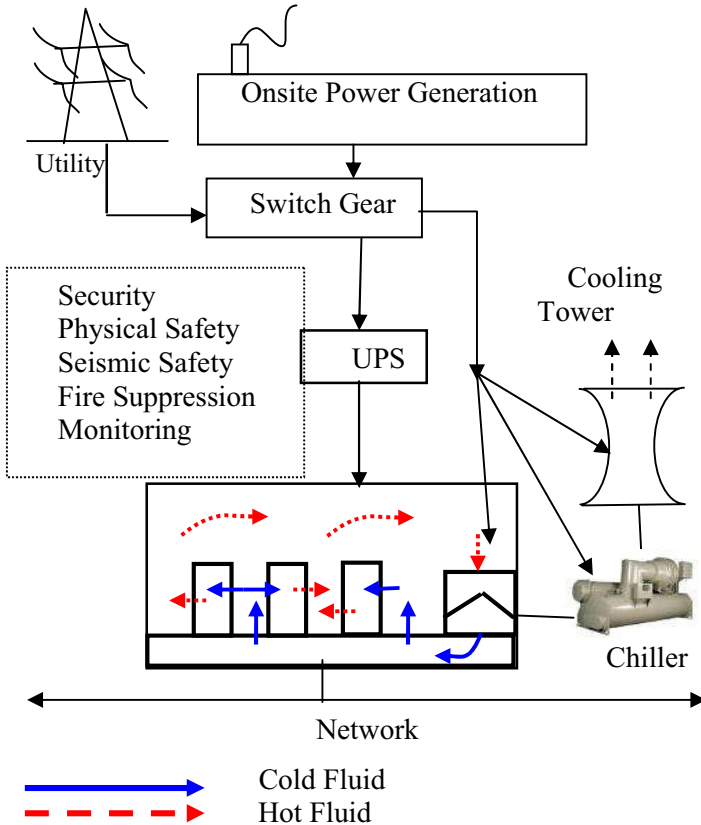


Fig. 1. Data Center Building Blocks




Fig. 2. Cross section of the datacenter

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 to the racks 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 air movers in the CRAC units pressurize the plenum with cool air which 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 by hot and cold aisles as shown in Figure 2. This separation is done for thermal efficiency considerations. Air inlets for all racks face cold aisles while hot air is expelled to hot aisles. A number of other equipment layout configurations and non-raised floor infrastructures also exist.

### 3 Architecture

The solution taps into a sensor grid distributed among racks and CRAC units [7] in the datacenter and implements data processing and analytics algorithms to extract information from it. It exposes a web-service interface to which mobile devices can connect to. The client application is implemented on HP iPAQ 210 mobile device, which runs Windows Mobile 6.0 and connects to the web-service over a Wi-Fi network (see fig.3). The mobile client application also provides an offline mode operation enabled with a rich set of visualization and analytics capabilities by enabling data mining on over 700,000 rows of data for every 24 hrs sensor data sampled on the mobile device. The choice of the mobile device was based on evaluating a number of physical and logistics constraints a datacenter operator would have to deal with while conducting either a routine or spontaneous operations workflow in a large data center environment including lack of desk space to operate more versatile large form-factor devices. This allows the data center administrators to use the solution in-the-field, i.e. inside the data center environment. The solution provides the administrator with up to the minute sensor information on any Data Center facilities infrastructure including information on a specific rack or set of racks that is being monitored by the sensors. The solution also enables a number of report generation capabilities that are available on demand backed by automated analytical engines. The data center administrator also has the ability to set specific filters based on thresholds for monitoring notification alerts. The client application automatically polls for the data and high-lights interest areas and creates canned reports for the operators use.




Fig. 3. Datacenter mobile studio – application architecture

Mobile devices have limited computing power, memory, storage space and connectivity options. Consequently, the design of the solution efficiently distributes the computation and data storage between the server and the mobile client device. Only the tasks that are dependent on the specific set of data under consideration and its spatio-temporal properties are executed on the device. Client component on the mobile device primarily consists of a correlation calculation component and visualization component. All the pre-processing requirements that are impervious to the dataset selected are performed on the server. Cleansing and transformations on the raw data being generated by sensor grid, which are compute intensive operations, are entirely performed on the server side web-service. The web-service also extends an interface to a repository of historical data staged and managed for its life cycle.

In the online mode of operation, which works over Wi-Fi, the client requests the web-service for specific data-sets based on user inputs, which get delivered in a ready-to-be-displayed form. Online mode allows for updated up-to-the-minute data to be available to the user. For the offline mode, the user can download a pre-packaged dataset containing complete but condensed 24-hour data for the entire data center on to the device. The pre-packaged datasets are periodically cached on the server side for client use. Implementation of visualization and correlation logic on the client, instead of the server allows the application to seamlessly and consistently work in connected and disconnected modes.

A reference implementation has been done and deployed in a 14,000 sq.ft. datacenter. This datacenter comprised sensor data monitored for over 500 racks, 11 CRAC

units consuming 600KW of power for the IT load. The data center was monitored by over 1350 sensors sampled at 1 minute interval.[8] The thermal and operational events, key performance indicators (KPIs) and key risk indicators (KRIs) are now available at the click of a button on the mobile device for further analysis and report generation.

The visualization in the iPAQ uses a spatio-temporal layout. Figure 4 illustrates a sample mapping of 4 sensors of multi-dimensional thermal data to a 2-D space. The layout provides both a “top” view for the data center thermal performance and interactive “data exploration” views for troubleshooting. Each rectangle contains a sensor temperature time series in a datacenter. Hot spots are marked for thermal correlation analysis and root-cause queries.

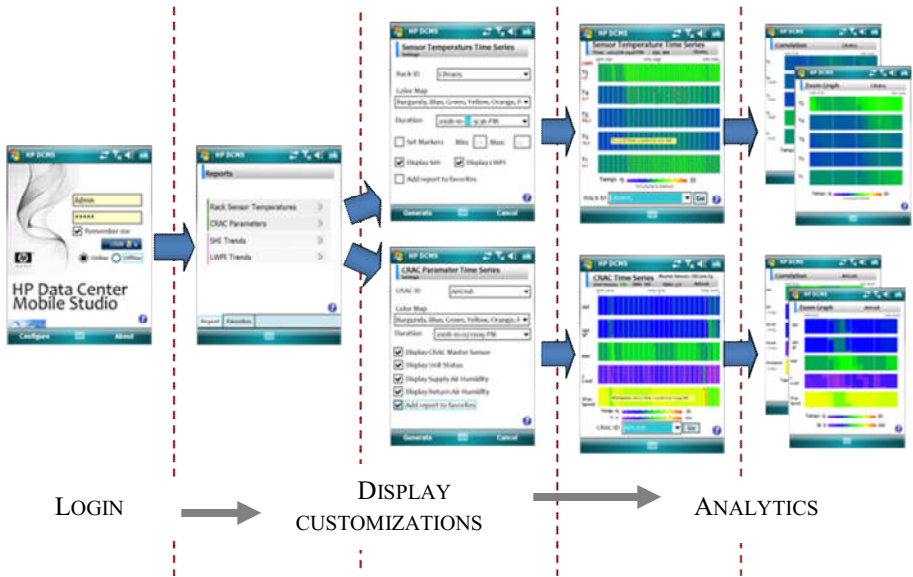


Fig. 4. Datacenter mobile studio – application pipeline

Figure 5 illustrate pixel cells of a time series which are arranged from bottom to top and left to right according to time. The color of a pixel is the temperature value of a sensor for each time interval. Some colors (i.e., red, orange, etc.) represent more severe conditions and alert the administrator to take immediate action. For example, T1 has a high temperature (over 35°C) at hour 11.




Fig. 5. A sensor temperature time series. The sensors are arranged in descending order with sensor T1 closest to the vent tile within a rack.

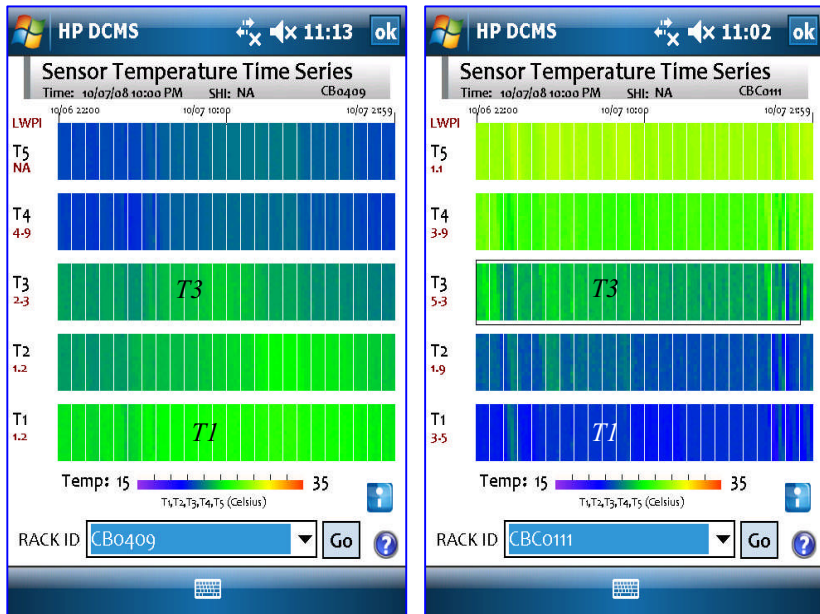
As described in the previous section internal datacenter infrastructure consists of racks and CRAC units. To visually analyze data center operations within the room, we use three different types of visual analytics queries: Finding correlation between inlet and outlet, at a change point, and among sensor attributes.

## 4 Case Studies

### 4.1 Thermal State Detection

Thermal state detection for distributed management and control of cooling resources is of great importance to datacenter operators, both on the facility and IT-side. High temperatures can adversely affect uptime and equipment performance and lead to premature end of life for equipment. Most data centers have four critical problem thermal states: (1) high sensor temperatures (e.g., temperatures exceeding a threshold), (2) high recirculation identified by temperature-based metrics such as Supply Heat Index (SHI) [2], (3) Large variation in temperature distribution identified by out-of-band sensor temperatures within a rack (e.g.,  $T1 > T2$ ), and (4) Oscillations between the previously described states and the normal conditions [5]. Detection of hot spots is necessary to prevent temperature related failures while detection of abnormal air flows and metrics helps to identify energy inefficiencies.

Figure 6 shows the screenshot of inlet air sensor temperature time series obtained from groups of 5 sensors from two alternate racks over a period of one day. Each screen displays a total of 7200 records. The color scale denotes temperature range



**Fig. 6.** Inlet air temperature display of alternate racks, the visualization on the left represents out of sequence temperatures (e.g.;  $T5 < T4 < T3 < T2 < T1$ ). The visualization on the right has out of sequence temperatures (e.g.,  $T5 > T4 > T3 > T2 > T1$ ) is a desirable time series.

between 15C and 35C. Observe the variation during the day and the reversal of temperature distribution between the two racks during the same time period. Such variations are prevalent in datacenters with heterogeneous equipment and knowledge of such states is important for onsite troubleshooting of thermal issues. Such a display can assist in normal operation by identifying mal configured racks or air distribution. A user can verify the rack assembly for potential air leakages as well. In many cases the server air flow path is not aligned with the air flow in the datacenter. Such mismatches can lead to temperature reversals as well.

### 4.2 Infrastructure Utilization

Utilization of infrastructure is often used as a yardstick to estimate return on investment [6]. The mobile datacenter studio provides methods to display utilization of power and cooling infrastructure, both historically and in real time. Users can apply analytics to obtain more information around inflexion points in the utilization display. Fan power utilization and chilled water utilization of CRAC units are a few of the several utilization metrics that can be displayed and analyzed. Among other things, infrastructure utilization can also assist in trouble shooting thermal state management issues. It can aid planning on addition and replacement of existing infrastructure and resource needs.

Figure 7 shows the CRAC unit time series display of data collected over a period of one day, from two units in the datacenter. On each screenshot, from top to bottom, air discharge temperature and set points are displayed followed by return air temperature, cooling load and blower fan power. While one of the units has 0% cooling load the




Fig. 7. CRAC unit displays - The pulldown window indicates the correlations between the supply air temperature and chilled water temperature. Figure on the left shows no chiller water utilization (i.e. zero). Figure on the right shows full utilization (100%).



other unit is at its maximum rated cooling load (100%) for a significant period of time. The demand variation between the units indicates the heterogeneity in heat load distribution in the datacenter. Users can investigate workload deployment in CRAC unit zones and balance out loads wherever possible. Such displays provide decision support related to cooling infrastructure maintenance/downtime schedules. Further analytics can help in linking CRAC unit events to IT power and cooling events.

Power and cooling accounts for significant portion of operating cost of datacenters. In addition to thermal management, energy consumption can be used to define operation state of the datacenter. Providing energy consumption enables a mobile user to factor it into management decisions. A high power usage level indicates that there is little margin for IT capacity expansion before power reduction can be achieved. Air flow studies can also be carried by data center operators to see real time impact of operation at different speeds.

### 4.3 Energy Consumption

Figure 8 shows the energy display on the CRAC unit time series screen. Tracking energy over time provides data administrator insights on time-of-use demand responses. Impact of rising power costs and routine downtimes can be reduced by appropriate demand response strategies. Variation in energy consumption can be correlated to thermal management states in the datacenter and power curtailment plans can be developed to manage demand during peak hours. Power and cooling delivery bottlenecks can also be identified using such displays.




Fig. 8. CRAC unit energy display – Average Fan Speed and power consumption for the rubber-banded region are shown in the pull-down menu

## 5 Discussions and Future work

### 5.1 Root Cause Analysis

Apart from access to time series and analytics thereof, datacenter mobile studio provides metrics that assist in assessment of datacenter operation. While thermal management is an ongoing challenge in datacenters, it can be a manifestation of several anomalies in the datacenter like mal configuration of racks, mal-operation of CRAC units, obstruction of vent tiles (see fig 2) or high workload.

Mixing of hot and cold air streams caused by mal-configuration of racks is captured by Supply Heat Index (SHI) [2]. Plots of historical and real-time values can be obtained by rubber banding the hot spot region in the sensor display screen for further diagnostics.


$$SHI = \frac{T_i - T_{ref}}{T_o - T_{ref}} \quad (2)$$

where  $T_i$ ,  $T_o$  and  $T_{ref}$  are the rack inlet, rack exhaust and CRAC unit supply temperatures (see fig. 2), respectively.

Another important parameter in this process is the Local Workload Placement Index (LWPI) [8]. It is defined as ratio of thermal management and air-conditioning margin at that location over the degree of recirculation at that location.

$$LWPI = AC \text{ Margin} + \text{Thermal Margin} - \text{Hot Air Recirculation} \quad (3)$$

Recirculation limits the ability to deploy workload while air conditioning margin improves the ability of the server (at the location of interest) to accept workload. It



**Fig. 9.** LWPI plot for a sensor location over a 24 hour period – Lower LWPI indicates lower cooling effectiveness and less favorable area for workload deployment

can be used to gauge the effectiveness of cooling for workload placement. A detailed definition of eqn.(3) can be found in [8]. Administrators are able to rubber-band an interesting area and query for the correlations or detailed information. Varying LWPI can initiate migration of workloads.

Figure 9 shows the plot of LWPI for a single sensor over a period of 24 hours. The region of low LWPI indicated in the graph shows a period of reduced effectiveness of cooling in the region around the sensor. Such downtrends can be an early trigger for workload migration process to begin. Workload can be relocated to regions of high LWPI during the course of such downshifts.

## 5.2 Data Collection

Data collection for onsite changes and measurement of select operational parameters are important for routine management of datacenters. Operators can download relevant data for analysis based on visual clues. Other data collection applications include logging of infrastructure changes, asset management, system operational states, generation of alarms and initiation of remediation procedures.

## 6 Conclusions

Advanced data analytics and knowledge models are needed to identify inefficiencies in the datacenter infrastructure. In addition to troubleshooting, information about infrastructure states can enable involvement of knowledge experts for assessments and analysis at the right time. Mobile studio combines this power of visual analytics with onsite datacenter management. It assists administrators in identifying key power and cooling events followed by root cause analysis during runtime. Future work will focus on combining historical data analysis with trends to predict anomalies and initiate preemptive measures.

## Acknowledgement

Authors wish to thank Vaibhav Bhatia, Rajkumar Velumani and Mohandas Mekanapurath for providing access to datacenter infrastructure for testing and evaluation.

## References

- [1] The Uptime Institute, Heat density trends in Data Processing, Computer systems and Telecommunications Equipment, White Paper issued by The Uptime Institute (2000)
- [2] Sharma, R.K., Shih, R., Bash, C.E., Patel, C.D., Varghese, P., Mekanapurath, M., Velayudhan, S., Kumar, M.V.: On building Next Generation Data Centers. In: Proceedings of Compute 2008, Bangalore, Karnataka, India, January 18-20 (2008)
- [3] Hao, M., Dayal, U., Keim, D.A., Schreck, T.: Multi-Resolution Techniques for Visual Exploration of Large Time-Series Data. In: Proceedings of IEEE VGTC Symposium on Visualization, EuroVis 2007 (2007)

- [4] Hao, M., Dayal, U., Keim, D.A., Morent, D., Schneidewind, J.: Intelligent Visual Analytics Queries. In: IEEE Symposium on Visual Analytics Science and Technology, CA (2007)
- [5] Sharma, S., Patel, S.: Application of exploratory data analysis techniques to temperature data in a conventional datacenter. In: Proc. ASME IPACK 2007, IPACK 2007-33170, BC, Canada (July 2007)
- [6] Patel, et al.: Energy flow in the information technology Stack: Introducing the coefficient of performance of the ensemble. In: Proc. ASME Intl. Mech. Eng. Cong. Exp. (IMECE), IMECE 2006-14830
- [7] Bash, C.E., Patel, S.: Dynamic Thermal Management of Air-Cooled Datacenters. In: ASME/IEEE ITherm 2006, San Diego, CA (2006)
- [8] Marwah, M., Sharma, R.K., Shih, R., Patel, C.D., Bhatia, V., Mekanapurath, M., Velumani, R., Velayudhan, S.: Data analysis, Visualization and Knowledge Discovery in Sustainable Data Centers. In: Proceedings of Compute 2009, Bangalore, Karnataka, India, January 9-10 (2009)
- [9] Bash, C.E., Forman, G.: Cool Job Allocation: Measuring the Power Savings of Placing Jobs at Cooling-Efficient Locations in the Data Center, USENIX, San Jose, CA (June 2007)
- [10] Hao, M., Sharmal, R.K., Dayal, U., Keim, D.A., Vennelakanti, R.: Application of Visual Analytics For Thermal State Management in Large Data Centers. CGF Journal (2009)
- [11] Hao, M., Keim, D.A., Dayal, U., Oelke, D., Tremblay, C.: Density Displays for Data Stream Monitoring. In: IEEE EuroVis 2008 (2008)