Visual Analytics of Cyber Physical Data Streams Using Spatio-Temporal Radial Pixel Visualization

M. Hao[1], M. Marwah[1], S. Mittelstadt[2], H. Janetzko[2], D. Keim[2],
U. Dayal[1], C. Bash[1], C. Felix[1], C. Patel[1], M. Hsu[1], Y. Chen[1], M. Hund[2]
[1]Hewlett-Packard Laboratories, Palo Alto, CA [2]University of Konstanz, Germany

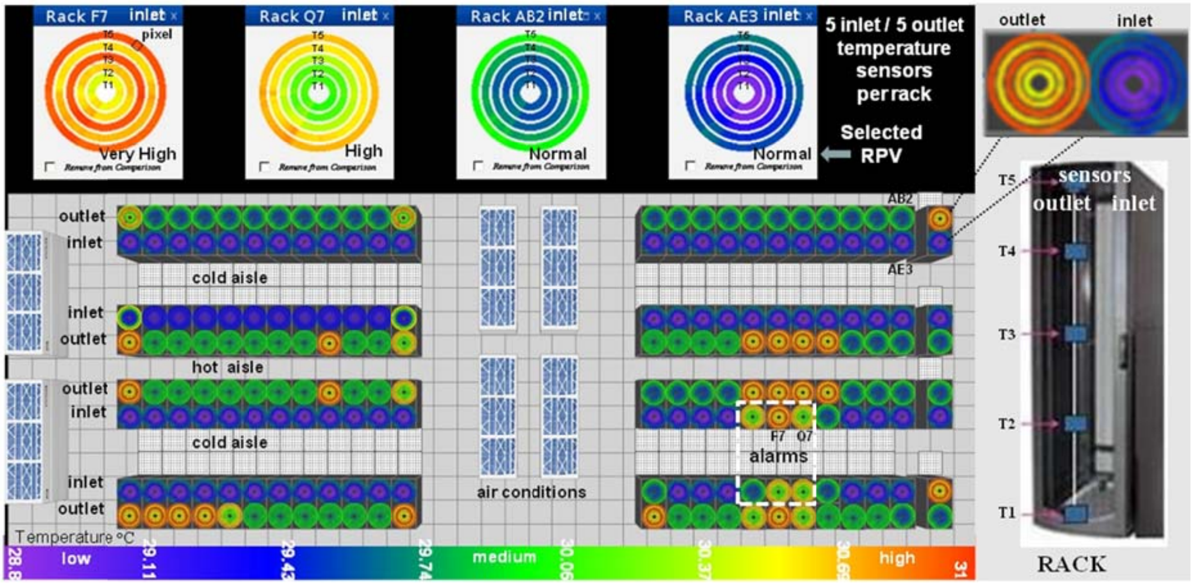

Figure 1: A 500 kilowatt data center showing 96 racks (8 rows, 12 racks per row) with temperature sensors to monitor the thermal state that allow the user to correlate and take action for disparate alarms. Each rack contains 5 inlet/outlet temperature sensors, thus monitoring large volumes of separate temperatures every minute over 24 hours in real-time. Each rack produces a pair of Radial Pixel Visualizations (RPVs) for their inlet and outlet sensors with color indicating temperature values. The temperatures can be checked against a user defined threshold and usually the temperatures are in an ascending sequence, that is, T1<T2<T3<T4<T5, where T1 is closest to the floor. Each minute of temperature measurements is represented by a colored pixel cell. Color depicts temperature using the standard data center color map ranging from low (purple, blue) to medium (green) to high (red). By correlating thermal alarms and their physical locations and by looking at temperature patterns in the recent past, administrators are able to quickly identify problems (e.g., Rack F7 (on top left) has an out of sequence sensor: T3>T4) and find the root causes of those alarms.

ABSTRACT

Cyber physical systems (CPS), such as smart buildings and data centers, are richly instrumented systems composed of tightly coupled computational and physical elements that generate large amounts of data. To explore CPS data and obtain actionable insights, we present a new approach called Radial Pixel Visualization (RPV); which uses multiple concentric rings to show the data in a compact circular layout of pixel cells, each ring containing the values for a specific variable over time and each pixel cell representing an individual data value at a specific time. RPV provides an effective visual representation of locality and periodicity of the high volume, multivariate data streams. RPVs may have an additional analysis ring for highlighting the results of correlation analysis or peak point detection. Our real-world applications demonstrate the effectiveness of this approach. The application examples show how RPV can help CPS administrators to identify periodic thermal hot spots, find root-causes of the cooling problems, understand building energy consumption, and optimize IT-services workloads.

Keywords: Radial pixel visualization, correlations, peaks, cyber physical system, sensor data streams

firstName.lastName@hp.com; firstName lastName@uni-konstanz.de

1.INTRODUCTION

1.1 Motivation

Cyber physical systems (CPS) [1, 2, 3] are systems that are characterized by a tight coupling between computational and physical infrastructures. Examples of such systems include smart buildings, data centers, smart electric grids, etc. Figure 1 shows a spatio-temporal data center layout of 96 racks with their inlet/outlet sensor temperatures measured every minute over 24 hours. Each rack has a pair of RPVs with their colors indicating the temperature values, thus allowing the user to correlate disparate alarms. Data center administrators are interested in exploiting this data in real-time to identify thermal problems, infer their root causes, determine under- or over-utilized resources, etc. Relevant questions include: Are temperature alarms that originate from different servers related? Are the alarms from servers in close physical proximity caused by the same underlying causes? What are the relationships between heterogeneous variables, such as the ambient temperature in a rack, CPU temperatures, and server utilizations? What periodic patterns occurred in the last few days? What are the causes of the alarms and how can the administrator take immediate action?

The above information discovered in the CPS data streams can be used to identify correlations, patterns, and failure conditions. The information can also be used to validate the current operational state and provide better overall management. In general, the basic requirements in visual analytics of CPS data streams are locality, compactness, and the ability to identify periodic patterns, correlations, and anomalies from multiple large time series.

1. 2 Related Work

To visualize CPS data streams (e.g., data center, IT-Services, traffic, etc.), there are several different advanced time series visualizations listed in recent survey papers [5, 6, 7]:

Cartesian Time Series Visualizations

Cartesian visualization techniques such as TimeSeries Bitmaps [8] or Recursive Patterns [9] have shown their ability to visualize large time series, but may fail to provide useful results for multiple periodic time series. The density display in [21, 22] places large volumes of data stream values in a spreadsheet-like row and column layout for discovering patterns and detecting alarms. Erbacher et al. [12] developed an environmental visualization for analyzing network traffic data to prevent the critical attacks. Andrienko et al. [25] support detection of events using statistical event detection methods and Schneiderman’s Time Searcher [13, 14, 15] suggests three interactive visual exploration techniques for analyzing large multivariate time-series (>10,000 data points) by showing up to ten simultaneous plots on the same screen. In [10], the authors discuss visualizing sets of non-equally spaced time series arising, e.g., from auction bids. Different from the above work, RPV displays the data streams with their physical locations and high resolution time series in a single view as shown in Figures 1 and 5.

Radial Time Series Visualizations

Radial visualizations play an important role in visualizing periodic data streams. From a survey paper [7], there are a number of radial visualizations used to visualize changes in large time-related multidimensional data sets, such as Circle View [16] and Circle Segments [17]. Further, radial hierarchical frequent pattern visualization has been used in market basket and web click stream analyses [11]. Stasko [18] enhanced the radical space hierarchy visualization with advanced focus-context techniques and a negation technique. An interesting technique in radial visualization is the spiral arrangement [19] to visualize time series data. Afterwards, Tominski et al. [20] enhanced the spirals with two-tone pseudo coloring. The spiral technique is an efficient and effective technique to visualize single variable data. However, it is difficult to visualize the relationships between multiple variables having different scales. A comparison of three radial visualization techniques for multivariate time series data is shown in Table 1.

All these prior techniques are able to visualize multiple time series and help to discover periodic patterns, but none of them tightly integrates correlation and peak point detection analyses, as well as semantic zoom and drilldown to detailed information.

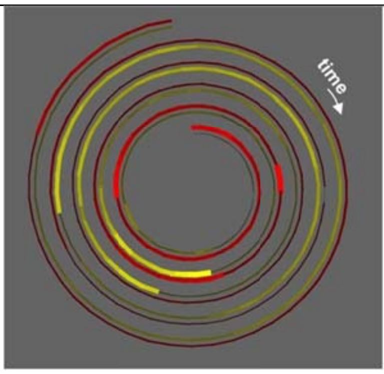	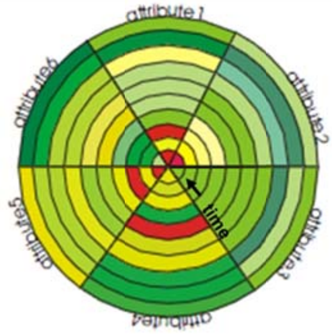	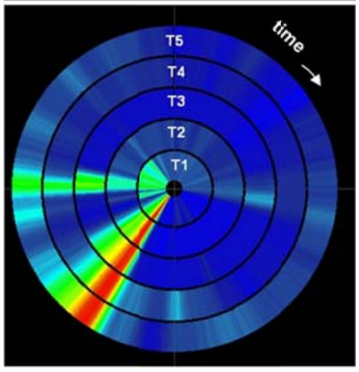
Multi Spirals - reference [19]: Shows two time series by rendering intertwined Spiral Graphs. Each spiral represents one time series. Color encodes different time series, e.g., Microsoft (yellow) and Sun Microsystems (red).	***Circle View - reference [16]:*** Shows a six attribute time series which is arranged in each segment. Color represents the aggregated value of each attributes. Each segment needs to be compared with the corresponding time slots of the neighboring segments to detect correlations.	***Radial Pixel Visualization - This Paper:*** Shows a five attribute time series. Each time series is arranged in one single ring using high resolution pixel cell-based time series. All time intervals (pixel cells) are aligned across different time series. Color represents the value of an attributes at a specific time interval.
- Each spiral represents one time series. - 4 to 8 spirals can be combined in a single graph using color coding. - To detect periodic patterns requires knowing the cycle length for proper alignment.	- Each segment represents one time series. - Can have many attributes per circle view. - To detect correlations requires placing related segments close to each other.	- Each ring represents one time series. - Can have many attributes per radial pixel visualization. - Able to detect periodic patterns (Figure 3) and correlations since the rings are time aligned. (Further contributions are described in Section 1.3)

Table 1: A comparison of three radial visualizations for multivariate time series data

1.3 Our Approach and Contributions

To leverage the prior work and to meet CPS challenges, we introduce a new approach of pixel cell-based Radial Pixel Visualization (RPV) with the following unique features:

1. The integration of pixel cells with radial visualization
   RPV uses multiple concentric rings to represent high resolution multivariate time series. Each measurement value is shown in a colored pixel cell. All pixel cells in the RPV are accessible to the users for drilldown.
2. Periodic pattern discovery
   The pixel cells are aligned by time in different rings which allows the easy detection of periodic pattern. In addition, the pixels are continuously placed in a ring over time (e.g., 24 hours). RPV is able to show data streams with cycles without a split as in Cartesian visualizations.
3. Multivariate correlation/peaks encoding
   To show the results of correlation analysis and peak point detection, we define an additional ring to show the analytic results. From the brightness of the analysis ring, analysts can detect the degree of correlation and dependency patterns. To prevent exceeding system capability, peaks can also be identified with the analysis ring, leading the analyst to focus on important and potentially dangerous events.
4. Physical infrastructure and location analysis
   To allow analysts to correlate the spatial relationships of different cyber data streams with their physical infrastructures, RPV uses x-y coordinates to relate the measurement data with their physical locations, e.g., longitude and latitude.

Using multiple RPVs, we can construct a real-time dashboard and combine it with advanced interaction techniques (e.g., semantic zoom and parameter control) to incorporate human domain knowledge into the process of solving problems.

2. RADIAL PIXEL VISUALIZATION (RPV)

Radial pixel visualization (RPV) is a temporally-aligned high resolution pixel time series graph for visualizing large volumes of data. RPV has two types of rings (data ring and analysis ring) and three different usages (single variable, multiple variables, and spatio-temporal layouts).

Data Rings

As illustrated in Figure 2, the ring construction has two parameters that define the size of the RPV: the minimum radius (minRadius), the size of the first inner data ring, and the maximum radius (maxRadius), the size of the outermost data ring. The construction then divides the 360° of a ring into a number of sectors according to the number of time intervals per period. Each sector is enclosed by two sector lines that start from a point with a distance of minRadius (Radius 1) from the center, and end at a distance of maxRadius (Radius 2) from the center. The pixel cell width (w) is equally distributed over each ring. Therefore, the sector lines are subdivided into nrOfRings (number of rings). The pixel cells are constructed by connecting the corresponding parts of each sector line to form a pixel cell.

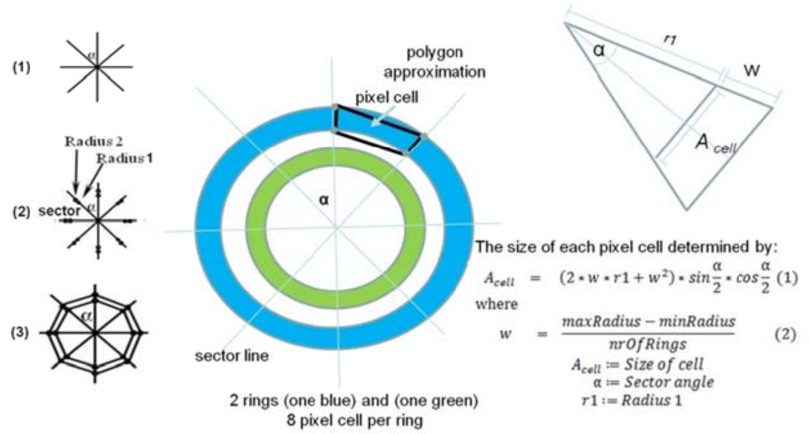

Figure 2: Radial Pixel Visualization (RPV) Construction

Analysis Ring (e.g., correlations, peaks)

In order to help users to quickly identify the important information in large data streams, we use an analysis ring on the outside of the data rings to highlight areas with a significant correlation between variables and/or interesting peak points. This unique feature did not exist in prior radial visualizations and is critical for CPS applications requiring anomaly detection and capacity planning. Details are described in sections 3.1 and 3.2.

2.1 Single Variable Rings

In a single variable RPV, each time period (one day, or week, or month, etc.) is assigned to a corresponding ring. This layout typically shows several time periods of a single variable. The ring alignment (from the inner to the outer) represents the time sequence. Each time period in a high resolution data stream is visualized to make periodic behavior perceivable.

Figure 3 illustrates a single variable (power consumption in a building) RPV over multiple time periods (from 6/6 to 6/8). As expected, each day shares a similar pattern: low consumption during the night and increasing consumption during the day with high values from noon to 6 pm. Note that 6/7 shows an area with high peak points (red, e.g., 960.32 KW) from 3 pm - 6 pm. The administrators need to investigate the source of this incident to determine the likely cause and whether it is related to a problem. In addition, knowing the periodic patterns of the daily power consumption enables administrators to predict future consumption.

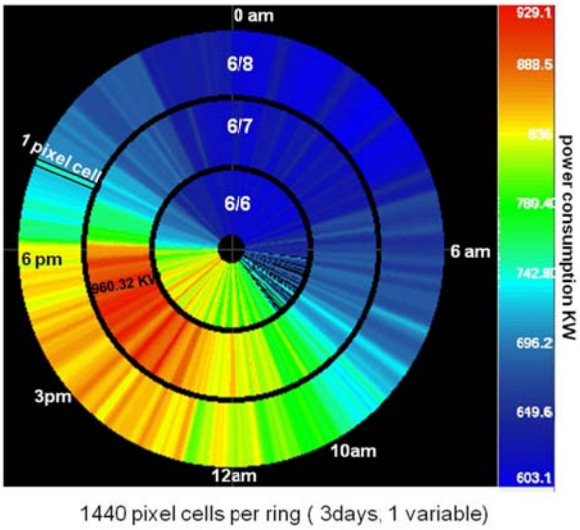

Figure 3: Single variable RPV with 3 one-day data rings (6/6-6/8)

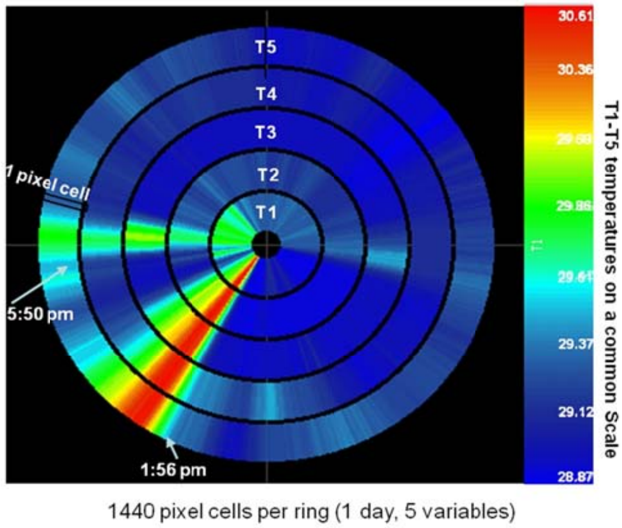

Figure 4: Multiple variable RPV with 5 one-day data rings (T1, T2, T3, T4, and T5)

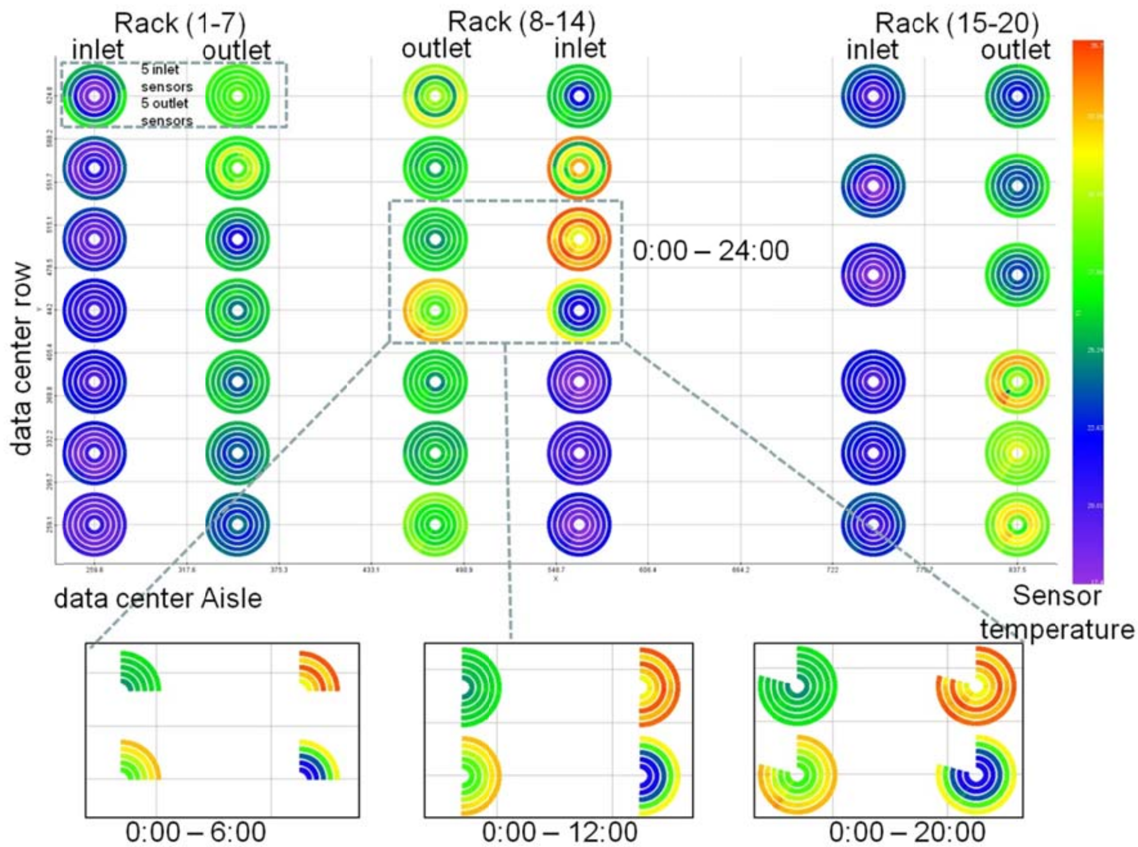

Figure 5: A 40 spatio-temporal RPVs in a data center in real-time
(x-axis: aisle, y-axis: row, color: temperature (each pair of rings represents a pair of in/out sensor (T1-T5) time series)

2.2 Multiple Variable Rings

The multi-variable RPV is used to analyze the relationships between multiple heterogeneous variables. Figure 4 is an example of temperatures from five data center sensors compared over the same day's time periods. Two patterns are visible. The first occurred at 1:56 pm and involved variables T1, T2, T3, T4, and T5 with correlated sensor high temperatures (red, yellow, and green); the second occurred at 5:50 pm and involved the same set of variables with correlated sensor medium temperatures (green and yellow). The first pattern lasts about twice long as the second. With this information, administrators are able to look into the root-causes of the high temperature periods.

2.3 Spatio-Temporal Layout

The spatio-temporal RPV layout allows an overview of a large number of data streams. Figure 5 shows 40 RPVs visualizing multiple rack temperatures in a data center. Each RPV uses the output of five different thermal sensors. In total, there are 200 data streams, containing interesting patterns. The analyst can easily identify problematic measurements in a system. For example, the hot temperature sensors are easily visible by their red colors.

By analyzing the spatial context in this overview, the analyst can build a hypothesis based on the spatial proximities and identify critical areas. Furthermore, spatio-temporal RPVs enable administrators to spot patterns across different locations and attributes and allow the identification of interesting incidents in real-time.

3. VISUAL EXPORATION

3.1 Correlation Detection

In this application, we compute the average of all pair-wise correlation-results shown by brightness in the analysis ring. One way to quickly identifying potential root causes is to find interdependencies. Correlation detection is a method to analyze the potential root-causes of a problem. In the early analysis, the sign of the correlation (positive/negative) is not as essential as the strength of the correlation. The positive or negative correlation will be analyzed when the user wants to do further analysis on the cause-effect between variables. The capability of highlighting the correlation strength guides the user to focus on interesting time intervals. We enable the user to find correlations between variables by showing them in the analysis ring (grey scale) in Figure 6.

First, we calculate the Pearson correlation coefficient to calculate the pair-wise correlations for a pair of the corresponding time slices. However, if the values of both variables remain constant over a time interval, the Pearson coefficient will be undefined (division by zero); in this case, we set their correlation coefficient to 0 in equation 3. The average of all pair-wise correlation-results is used to brighten the analysis ring. Negative and positive correlations do not compensate each other, because equation 3 only considers absolute values. Thus, the brightness of the analysis ring identifies the degree of the correlation between variables. Time intervals without correlations are faded out to reduce visual clutter (i.e., black: not correlated; white: highly correlated as in Figure 6).

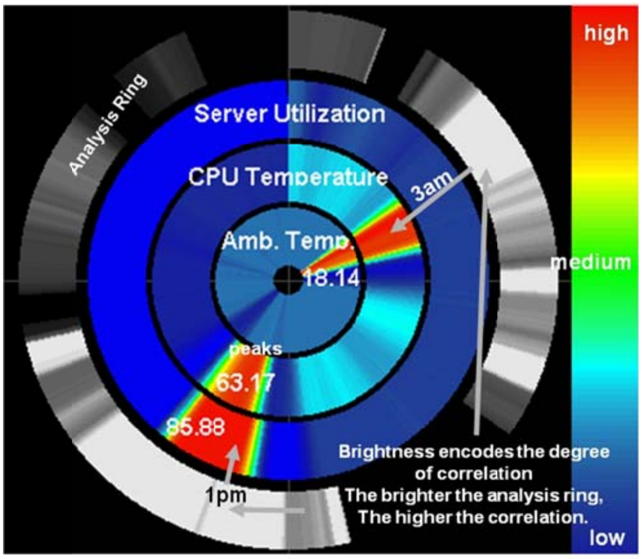

Figure 6: Daily Thermal State Analysis Using Analysis Ring (4 rings show: Server Utilization, CPU Temperature, Ambient Temperature, and analysis ring on correlation). Two extended high correlation areas are shown around 3 am and 1 pm.

$$C(A1,A2) = \begin{cases} 0, \mathit{if}\ \Delta A1, \Delta A2 \approx 0 \\ \left| \dfrac{Cov(A1,A2)}{\sqrt{Var(A1)} \cdot \sqrt{Var(A2)}} \right| \mathit{else} \end{cases} \qquad (3)$$

$C :=$ Correlation of A1 and A2
$A1, A2 :=$ Attributes
$\Delta A :=$ Gradient of Attribute
$Cov :=$ Covariance
$Var :=$ Variance

To illustrate the usage of the analysis ring on correlation, the data center administrator can analyze the relationships between heterogeneous variables, such as Server Utilization, CPU Temperature, and Ambient Temperature. As illustrated in Figure 6, CPU temperatures are high in two different time intervals (around 3 am and 1 pm). When observing only the CPU temperature ring, the root cause of the high temperature would remain unclear. By plotting both the ambient temperature and server utilization in Figure 6, it becomes evident that the first increase in CPU temperatures at 3 am is related to a period of high ambient temperature while the second increase at 1 pm is related to high server utilization. With this knowledge, the data center administrators are able to manage their resource consumption more effectively.

3.2 Automatic Peak Detection and Labeling

For thermal state analysis, peak points also have a high level of significance. Peaks in temperatures may, for example, reduce hardware reliability. We can use the analysis ring to show the significant peaks in the corresponding time intervals instead of showing correlations.

The automatic peak area detection algorithm is based on a variant of the Douglas Peucker algorithm equation that only considers absolute values [23]. The data stream is first reduced to a number of peak points as shown in Figure 7 (A) and (B). Each of the peak points higher or lower than a given threshold in Figure 7 (C) is encoded in the analysis ring. To enhance the visibility of the discrete peak points, we use a Gaussian kernel to brighten the analysis ring at the corresponding position. The Gaussian curves will overlap and increase if two or more peaks from the same or different rings are close to each other, which enables users to easily detect the high density peak areas.

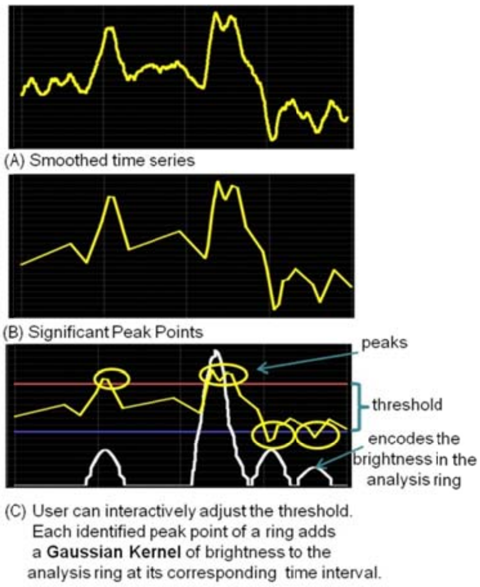

Figure 7: Automated peak points detection

Figure 8 illustrates the three most significant peak points (i.e., 85.88% in server utilization, 63.17^{O} C in CPU temperature, 18.14^{O} C in Ambient temperature in the data rings). These peaks were automatically labeled at the time that the peaks (exceed maximum values, red) occurred. Discovering these peaks during CPS operation is essential for administrators to manage their resource consumption and for capacity planning; e.g., before new equipment can be added at a particular location in a data center, the administrator needs to make sure that the peak and cooling requirements can be met. Using the highlighting from the peak point detection analysis ring enables administrators to quickly focus on the peak resource. From Figure 8, administrators are able to compare the peaks with the normal temperature readings (35^{O} C) and server utilization (80%) to make a decision about whether or not to add additional computer power and/or cooling resources.

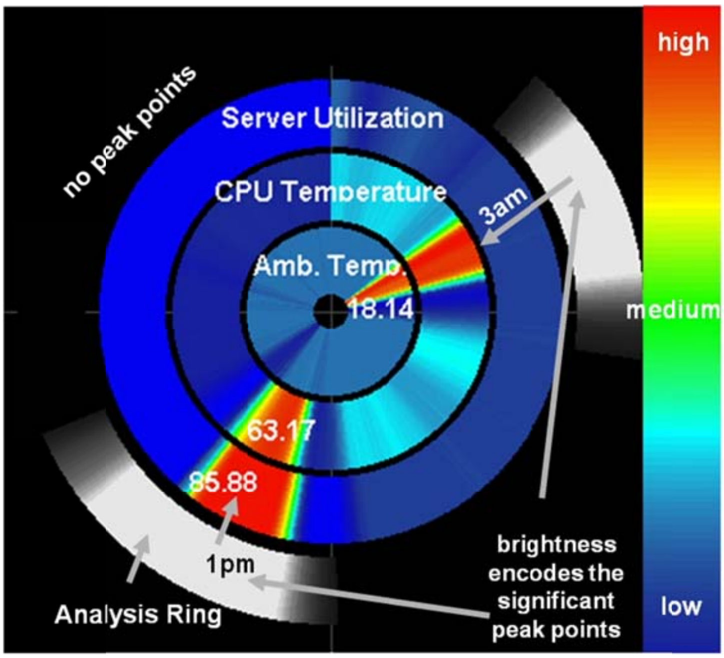

Figure 8: Encode automatic peak detection in the analysis ring. The brightness indicates significant peak point areas

3.3 Semantic Zoom and Drilldown

Besides changing the size of radial pixel visualization and controlling the parameters for calculating the correlations and peak points, the user is able to perform a semantic zoom of a problem sensor (red) as shown in Figure 9 (A). The zoomed RPV can be displayed at a higher resolution (the color map is rescaled) to identify high temperatures (red) and peaks as shown in the bright area of the analysis ring in Figure 9 (B). A tooltip displays detailed information (location, attribute name, value, and time stamp) of the selected data point. In addition, the corresponding RPV can be displayed as line charts in Figure 9 (C) with more accurate details before smoothing.

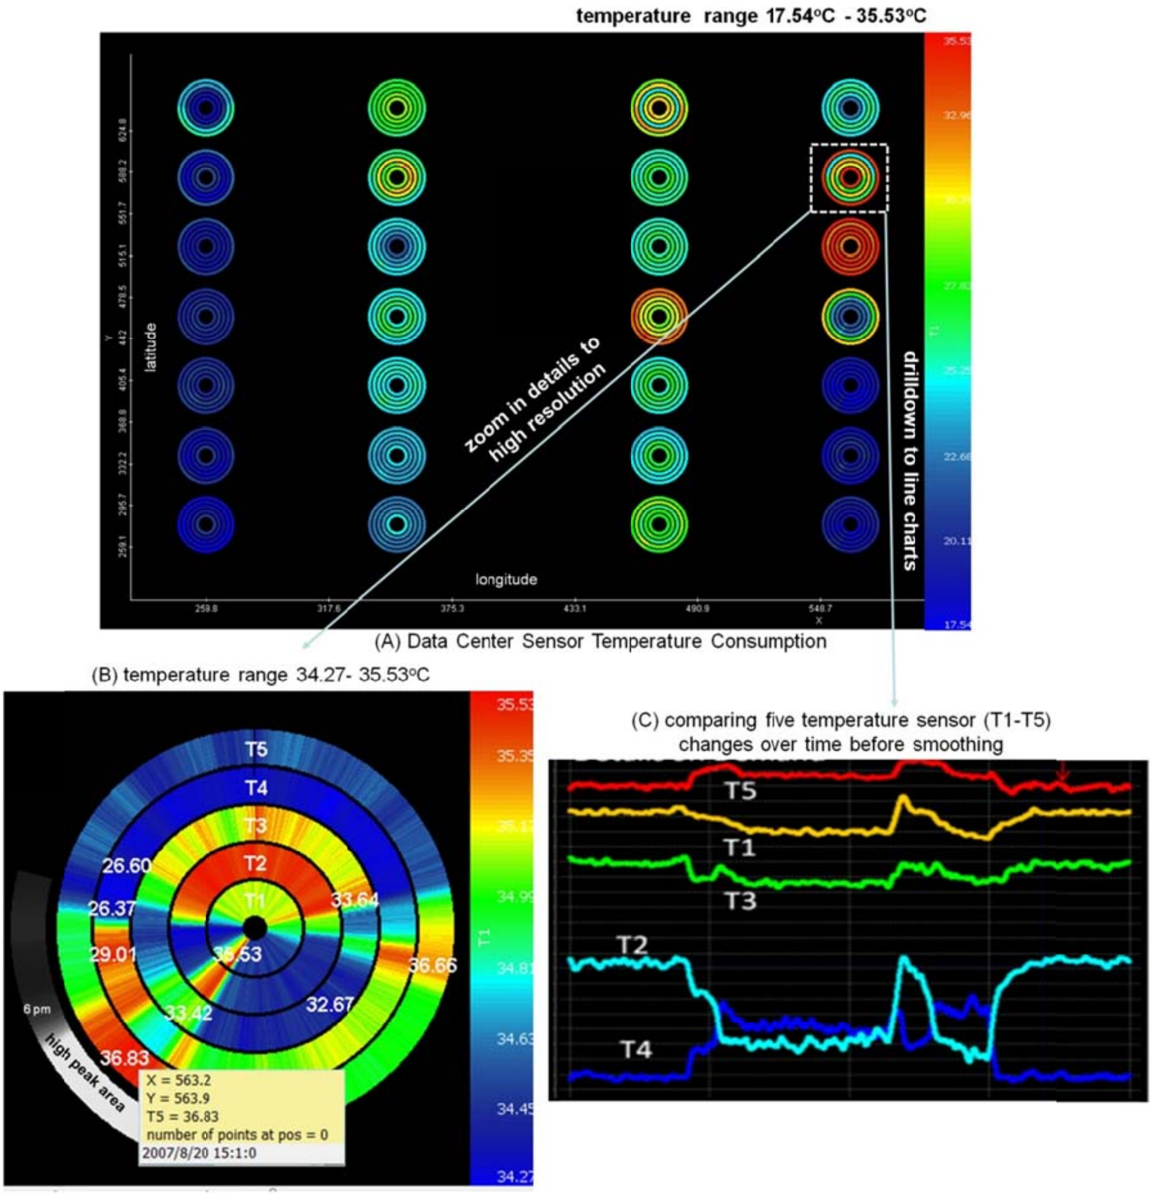

Figure 9: Interactive semantic zooming and drilldown to detailed information

4. USE CASES

4.1 IT Workload Optimization

Figure 10 shows attributes pertaining to IT workload and cooling power in a data center over a 24-hour period. Administrators are able to use RPV to visualize the power consumption differences and reschedule the non-critical IT workload to use the photovoltaic (solar) output in the daytime. Before optimization, the non-critical workload was spread throughout the day in Figure 10 (A); after optimization, it is concentrated during the period of PV power generation (solar supply) as shown in Figure 10 (B). The visualization confirms the benefits of rescheduling the workload and reducing the overall daily cooling power consumption [24].

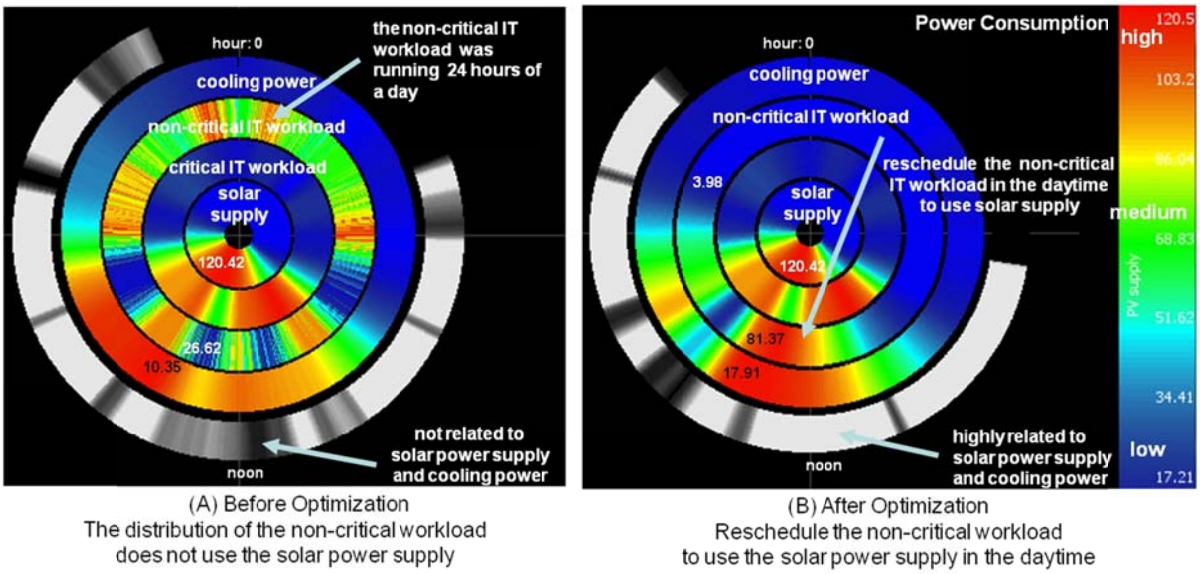

Figure 10: Use RPV to visually validate the benefits of reschedule the not-critical IT workload
Use normalized scale for each variable (low: blue, medium: green, high: red)

4.2 Building Energy Consumption Characterization

Figure 11 shows the power consumption of three buildings at an urban campus over three days. The daily usage patterns are easy to see for Buildings 1 and 2. As expected, consumption is high during the working hours (9 am to 6 pm). Building 3's consumption is flat, due to the presence of solar panels that offset part of the demand during the day. The visualizations allow administrators to compare usage patterns between buildings and validate the impact of solar panels under different weather conditions (e.g., sunny, cloudy, etc).

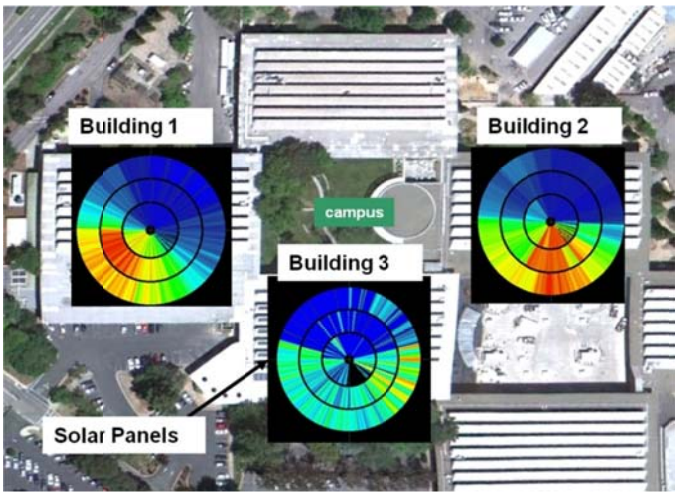

Figure 11: Building 3 has flat energy consumption (Solar)
(Each ring shows one day's consumption, 6/6-6/8)

5. EVALUATION

To evaluate the effectiveness of Radial Pixel Visualization (RPV) and compare it with three well-known visualizations (Multiple Line Charts, HorizonGraphs [4], and RecursivePattern [9]) as illustrated in Figure 12, we conducted an informal user study with 11 domain experts from four different cyber physical system areas: IT Services, Data Center, Campus Building, and Research Labs. Both HorizonGraphs and RecursivePattern use Cartesian coordinates.

Figure 12 illustrates an IT service performance analysis. The data set consists of five different variables (Memory, I/O Operation, %Utilization, #Users, and #Transactions) recorded over eight days in 5-minute intervals. In order to compare four different techniques, each one is given exactly the same physical space.

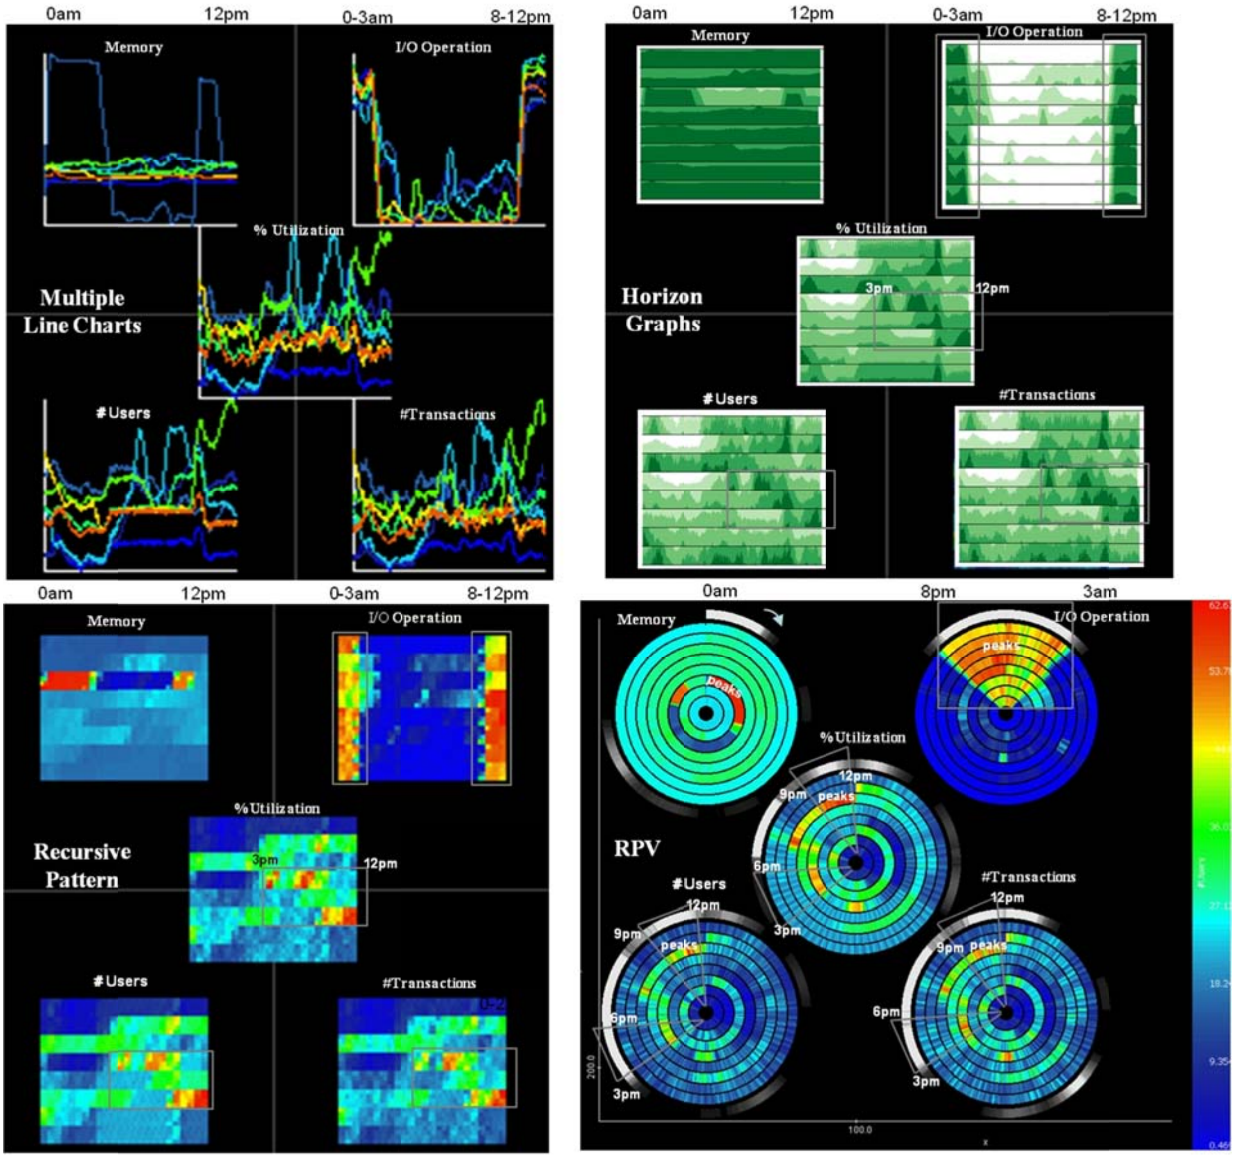

Figure 12: Evaluation of Multiple Line Charts, HorizonGraphs, RecursivePattern (Cartesian Layout) and Radial Pixel Visualizations (RPV) with an analysis ring highlighting peaks.

The user study led to the following observations:

The strengths of Radial Pixel Visualization (RPV):

1. ***The integration of pixel cells and periodic pattern discovery***: Using pixel cell-based RPV, users are able to visualize large data streams with drilldown capability. RPV is able to show a continuous data stream (from comparing evening to next morning) without any split as illustrated in Figure 12's I/O operation (8 pm to next day 3 am) showing a clear advantage compared to HorizonGraphs and RecursivePattern.
2. ***Multivariate information encoding***: This only exists in RPV. The results of an automatic correlation or peak point detection analysis of multiple time series are encoded in the analysis ring for helping users to quickly identify problems in real-time. Two significant peak point areas are highlighted by the brightness of the analysis ring around 3 pm to 6 pm and 9 pm to 12 pm as shown in Figure 12, on the %Utilization, #Users, and #Transactions RPVs. These three RPVs are highly correlated such as when the number of users and transactions increase, the corresponding %utilization increases accordingly.
3. ***Locality***: The following are three use cases to validate the importance of visualizing locality:
   a) In Figure 1, the spatial layout of a RPV can be used to detect the root cause of alarms in a data center. For example, if the sensor temperature alarms (e.g., near RPV F7) are all from the same aisle of the data center, the floor vent tiles in that aisle may require adjustment.
   b) In Figure 11, three building RPVs are placed on a geographical map based on their longitude and latitude to compare energy consumption between buildings as illustrated in Section 4.2.
   c) In Figure 12, three highly correlated RPVs with attributes %Utilization, #Users, and #Transaction are placed together to form a cluster for quick identification.

The Weaknesses of Radial Pixel Visualization (RPV):

1. ***Variable sized pixel cells***: The inner rings have smaller pixel cells than the pixel cells in the outer rings. To overcome this difficulty, we need to place the least important variables in the inner rings, e.g., the temperature sensor least likely to violate the threshold (usually T1 in Figure 1).
2. ***Variable sized data rings:*** RPV has variable size data rings because of its variable size pixel cells. In Cartesian coordinates, all the data segments have an equal size which could make the pattern comparisons easier than RPV. However, all three radial visualizations (HorizonGraphs, RecursivePattern, and RPV) show similar patterns and anomalies in attributes (%Utilization, #Users, and #Transaction) in Figure 12. Using different colors, RecursivePattern, and RPV visualizations are much easier to find patterns than HorizonGraph.

Of 11 domain experts, 6 preferred RPVs, 3 preferred RecursivePattern and 2 preferred HorizonGraphs. To our surprise, none of the experts preferred the multiple line charts which were most likely due to the high degree of overplotting. From the user study results, we have learned that a user's preference is highly dependent on the application requirements and the user's personal experience. Both Cartesian coordinate visualizations (HorizonGraphs and RecursionPatterns) and RPV are able to provide an overview of the relationships among multiple data streams. Furthermore, RPV is able to visualize the entire CPS system with locality and the results of an automated analysis.

6. CONCLUSION

In this paper, we presented RPV (Radial Pixel Visualization), a new approach for CPS (Cyber Physical System) visualization with comprehensive features for visualizing large amounts of multi-attribute data. RPV combines pixel cell-based radial visualization with efficient and effective knowledge discovery techniques. We have applied the RPV idea to real data sets from data centers and smart buildings. The resulting Radial Pixel Visualizations provide significantly more information than radial visualizations without using pixel cells. In the future, we will explore the potential of embedding RPV into building or site-level energy management dashboards. Furthermore, we will apply RPV to audience sense applications which are important in analyzing customer shopping behavior.

ACKNOWLEDGMENTS

The authors wish to thank Walter Hill for providing the IT service data and his comments and suggestions.

REFERENCES

[1] R. H. Campbell, et.al. Cyber-Physical Systems: position paper: CPS Environments. University of Illinois, Urbana-Champaign, 2012.

[2] NSF: Workshop on Cyber-Physical Systems, Austin, TX., 2006.

[3] A. D. Wood and J. A. Stankovic: Human in the Loop: Distributed Data streams for Immersive Cyber-Physical Systems. ACM SIGBED Review – Special Issue on the RSS forum on deeply embedded real-time computing. Volume 5 Issue 1, January, 2008.

[4] H. Reijner: The Development of the Horizon Graph. Panoption Software.

[5] W. Aigner, S. Miksch, H. Schumann, C. Tominski: Visualization of Time-Oriented Data, 2011, XVI.

[6] S. Diehl, F. Beck, and M. Burch: Uncovering Strengths and Weaknesses of Radial Visualizations – an Empirical Approach. IEEE Transactions on Visualization and Computer Graphs. VOL. 16, NO. 6, November/December 2010.

[7] G. M. Draper, Y. Livnat, R. F. Riesenfeld: A Survey of Radial Methods for Information Visualization. IEEE Transactions on Visualization and Computer Graphs. VOL. 15, NO. 5, September/October 2009.

[8] N. Kumar, N. Lolla, E. Keogh, S. Lonardi, C. Ratanamahatana: Time-Series Bitmaps: A practical visualization tool for working with large time series databases. In SIAM 2005 Data Mining Conference.

[9] D. A. Keim, M. Ankest M., H. Kriegel. Recursive pattern: A technique for visualizing very large amounts of data. In Proceedings of the 6th conference on Visualization'95 (1995), IEEE Computer Society, p. 279.

[10] A. Aris, B. Shneiderman, C. Plaisant, G. Shmueli, W. Jank: Representing Unevenly-Spaced Time Series Data for Visualization and Interactive Exploration. Proc. of International Conference on Human-Computer Interaction, pp. 835-846, 2005.

[11] D. A. Keim, J. Schneidewind and M. Sips. FP-Viz: Visual Frequent Pattern Mining. Proceedings of IEEE Symposium on Information Visualization (InfoVis '05), Poster Paper, 2005.

[12] R.F. Eracher et al., Designing Visualization Capabilities for IDS Challenges. VISCE '05 Proceedings of the IEEE Workshops on Visualization for Computer Security. IEEE Computer Society Washington, DC, 2005.

[13] H. Hocheiser, B. Shneiderman: Dynamic Query Tools for Time Series Data Sets, Timebox Widgets for Interactive Exploration, Information Visualization, 2004.

[14] P. Buono, A. Aleks, B. Shneiderman, C. Plaisant, A. Khella: Interactive Pattern Search in Time Series. Proc. of Visual Data Analysis Conference VDA05. San Jose, CA, 2005.

[15] P. Buono, C. Plaisant, A. Simeone, B. Shneiderman, G. Samuel, W. Jank: Similarity-Based Forecasting with Simultaneous Previews: A River Plot Interface for Time Series Forecasting. Proc. 11th International Conference on Information Visualization. Zurich, Switzerland, 2007.

[16] D. Keim, J. Schneidewind, M. Sips: Circle View – A new Approach for Visualizing Time-related Multidimensional Data Sets. ACM Advanced Visual Interfaces, International Working Conference, AVI 2004, May 25-28, Gallipoli (Lecce), Italy, 2004.

[17] M. Ankerst, D. A. Keim, H. Kriegel: "Circle Segments": A Technique for Visually Exploring Large Multidimensional Data Sets. Visualization '96, Hot Topic Session, San Francisco, CA, 1996.

[18] J. Stasko, E. Zhang: "Focus+Context Display and Navigation Techniques for Enhancing Radial, Space-Filling Hierarchy Visualizations, InfoVis2000.

[19] M. Weber, M. Alexa, W. Muller: Visualizing time-series on spirals. In proceedings of the IEEE Symposium on Information Visualization, 2001.

[20] C. Tominski, H. Schumann: Enhanced interactive spiral display. In Proceedings of the Annual SIGRAD Conference, Special Theme: Interactivity (2008), pp. 53–56.

[21] M. C. Hao, U. Dayal, D. A. Keim, R. K. Sharma, A. Mehta: Visual Analytics of Anomaly Detection in Large Data Streams. Visual Data Analysis Conference, CA, 1/20/09.

[22] M. C. Hao, D. A. Keim, U. Dayal, D. Oelke, C. Tremblay: Density Display for Real-Time Data Monitoring, IEEE EuroVis, 2008.

[23] D. Douglas & T. Peucker: "Algorithms for the reduction of the number of points required to represent a digitized line or its caricature", The Canadian Cartographer 10(2), 112–122 (1973).

[24] Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwah, C. Hyser, "Renewable and Cooling Aware Workload Management for Sustainable Data Centers", ACM SIGMETRICS/Performance, June 11-15 2012, London, UK.

[25] G. Andrienko et al. Extracting Events from Spatial Time Series. IV '10 Proceedings of the 2010 14th International Conference Information Visualization. Pages 48-53, IEEE Computer Society Washington, DC, 2010.