Visual Exploration of Air Quality Data with A Time-Correlation Partitioning Tree Based on Information Theory

WEI CHEN, Zhejiang University, P.R.China
FANGZHOU GUO, Zhejiang University, P.R.China
FEIRAN WU, Huawei Technologies Co., Ltd., P.R.China
QI WANG, Zhejiang University, P.R.China
LEI SHI, Institute of Software Chinese Academy of Sciences, P.R.China
HUAMIN QU, Hong Kong University of Science and Technology, P.R.China

Discovering the correlations among variables of air quality data is challenging because the correlation time-series are long-lasting, multi-faceted, and information-sparse. In this paper, we propose a novel visual representation, called Time-Correlation Partitioning (TCP) tree that compactly characterizes correlations of multiple air quality variables and their evolutions. A TCP tree is generated by partitioning the information-theoretic correlation time-series into pieces with respect to the variable hierarchy and temporal variations, and reorganizing these pieces into a hierarchically nested structure. The visual exploration of a TCP tree provides a sparse data traversal of the correlation variations, and a situation-aware analysis of correlations among variables. This can help meteorologists understand the correlations among air quality variables better. We demonstrate the efficiency of our approach in a real-world air quality investigation scenario.

CCS Concepts: • **Information systems** → **Data analytics**; • **Human-centered computing** → **Visual analytics**;

Additional Key Words and Phrases: Sensor; Multivariate time-series; Information Theory; Transfer Entropy;

ACM Reference Format:
Wei Chen, Fangzhou Guo, Feiran Wu, Qi Wang, Lei Shi, and Huamin Qu. 2017. Visual Exploration of Air Quality Data with A Time-Correlation Partitioning Tree Based on Information Theory. 1, 1, Article 1 (August 2017), 20 pages.
https://doi.org/10.1145/nnnnnnn.nnnnnnn

This work is supported by National 973 Program of China (2015CB352503), Major Program of National Natural Science Foundation of China (61232012), National Natural Science Foundation of China (61422211).
Authors' addresses: W. Chen, State Key Lab of CAD&CG, Zhejiang University, 866 Yuhangtang Rd, Hangzhou, Zhejiang, P.R.China; F. Guo, State Key Lab of CAD&CG, Zhejiang University, 866 Yuhangtang Rd, Hangzhou, Zhejiang, P.R.China; F. Wu, Huawei Technologies Co., Ltd., Hangzhou, Zhejiang, P.R. China; Q. Wang, State Key Lab of CAD&CG, Zhejiang University, 866 Yuhangtang Rd, Hangzhou, Zhejiang, P.R.China; L. Shi, Institute of Software Chinese Academy of Sciences, Beijing, P.R.China; H. Qu, Hong Kong University of Science and Technology, Hong Kong, P.R.China.

XXXX-XXXX/2017/8-ART1 $15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

The rapid growth of industrial economy and oil-fueled vehicles has dramatically increased the global air pollution all over the world. According to WHO, ambient air pollution contributes to 6.7% of all deaths [1]. Due to this strong tie between air quality and health [26], the air quality problem has attracted growing attentions. For many years, meteorologists have been analyzing the air pollutants (such as oxynitride and particulate) together with weather variables (such as temperature and relative humidity) , which are monitored in modern cities, in order to understand the dynamics of air pollutants.

At the heart of fighting global air pollutions, analyzing air quality data requires interdisciplinary knowledge and techniques to exploit of the time-oriented, multivariate nature of this data, and to enhance situation awareness for domain users. In this practice, data visualization techniques incorporated with clustering, dimension reduction and data simplification analysis can be important to provide a clear view of multiple air quality variables and their evolutions [7, 16]. Previous studies have made significant progress on monitoring, analyzing, and forecasting the air quality and weather conditions. In visualization, the weather data and air pollutants are often displayed on a map monitoring the air quality in certain area [33, 45]. Regression analysis [8], statistical analysis [26] and correlation analysis [21] are often used to analyze the patterns in air quality data. As for forecasting, most studies concentrated on the visualization of predictive models and the ensembling of data for more precise forecast result [22]. The correlation analysis between air pollutant and weather has been focusing on static times, e.g., year by year, without considering the evolution of the correlations along time [21]. In particular, few attentions have been paid on the mutual and dynamic influences of multi-faceted variables.

This paper aims at the temporal correlation analysis of air pollutants and weather variables collected from multiple sensors. Importantly, we introduce both symmetric and asymmetric information-theoretic measures to capture the correlation portrait among variables. While most correlation visualization and analysis techniques [4] [40] can be applied to our scenario to interpret multiple sensor data, the resulting correlations displayed are temporally long-lasting, dynamically changing, multi-faceted, and information-sparse, making the task of interactive exploration time-consuming. We note that the temporal coherence is frequently used in the correlation analysis of time series. Using this coherence appropriately, we can effectively abstract the time series to support efficient sparse data traversal. Meanwhile, physically meaningful correlations only exist in a limited set of variable tuples. Exploiting this sparsity can greatly reduce the analysis overhead. Conventional time-partition [11] or variable-based graph structures [1, 21] have been proven to be effective in characterizing the coherence, trend and similarity in terms of time or variable. However, treating the temporal and multi-faceted variations equally may prevent the possibility of detecting interesting correlations of a specific variable pair in a small time interval.

Based on the above observations, we proposed a novel hierarchical data structure named time-correlation partition tree (TCP tree) and embedded this novel structure in a visualization system. The TCP tree is capable of capturing the sparsity in both the temporal domain and the variable domain. Integrating two domains into a single tree structure enables users to explore the correlations in different level of details and analyze the temporal patterns of the correlations within a consistent visualization. By allowing dynamic tree navigation and on-demand visual exploration of local correlation time-series, users are empowered with a capability of locating and identifying interesting correlation patterns in a context-aware fashion.

In summary, the contribution of this paper can be summarized as:

[1] http://www.who.int/gho/phe/outdoor_air_pollution/burden/en/

- A novel hierarchical data structure, TCP tree, which organizes the correlations among air quality variables in both temporal and variable domain. The TCP tree intuitively enables users to explore and analyze the evolution of correlations among sets of air quality variables.
- An interactive visualization system illustrating the TCP tree and a set of novel visual designs enable users to interactively construct TCP tree and explore the correlations.

The rest of this paper is structured as follows. Section 2 summarizes the related work. We describe the analytical tasks and design goals in Section 3. Section 4 explains how the information-theoretic correlations are calculated. Section 5 introduces the structure and construction of the TCP tree. The visual design is elaborated in Section 6. Section 7 summarizes the interactions in the system, and the case study is introduced in Section 8. Finally, we conclude this paper in Section 9.

2 RELATED WORK

2.1 Air Quality Analysis

Analysis of multivariate air quality data turns out to be a prolonged scientific battle involving analysts from diverse academic domains. Qu et al. [21] integrated a suite of novel visualizations into their comprehensive system, including circular bar charts and weighted complete graphs, in support of the analysis of air pollution problem in Hong Kong. Although they take into account the key role played by wind direction and speed in weather data visualization, the lack of corresponding geographical information maintains a fatal shortcoming of their research work. Zheng et al. [46], on the other hand, employed a co-training-based semi-supervised learning approach to improve the air quality inference accuracy. Both spatially and temporally related features are identified in their approach.

Air quality monitoring is of great assistance to analysts during the air quality analysis process. V*ö*lgyesi et al. [33] proposed SensorMap, an overall air quality monitoring system based on car-mounted sensor data, to gain a detailed picture of the air quality in a large area at a low cost. Unfortunately, their work lean more toward hardware platform development rather than visualization and analysis. Another essential task of air quality analysis lies in prediction. Different from works that only focus on measuring temporal correlations of weather variables and air pollution [2, 14], Demuzere et al. [8] took it a step further and extended their method to an alternative air quality prediction tool on the basis of similar correlation investigation. WeaVER [22] presents a series of practical encoding choices to interpret multiple weather features as well as their interactions, which benefits weather forecasts in an intuitive way. However, no formal evaluations are provided to prove the validity of their visualization designs. In this paper, we construct a TCP Tree structure tailored for air quality analysis which characterizes correlations of multiple variables and their evolutions based on information entropy measurements.

2.2 Information Theory in Visualization

The information theory has recently attracted much attention in the visualization field [38] [5]. Using information theory in data analysis and visualization can help build connections between data communications and data analysis and visualization [38]. Theoretically, the stages of a visualization process can be interpreted using the taxonomy of information theory [5].

Generally, the information entropy is employed to measure information quantitatively. It is quite useful for locating important regions and improving the analysis and visualization efficiency, e.g., placing seeds for streamline generation [43]. One example is the view selection that can be optimized by measuring the information entropy associated with different views [3] [29] [15]. The quality of the LOD view [37] can also be evaluated by using the information entropy. Similarly, the importance-driven focus of attention [36] can be captured by building an information channel between objects and viewpoints.

An important usage of the information theory is to measure the correlation between two variables. The symmetrical mutual information can be used to evaluate the similarity between isosurfaces [4]. Likewise, the relative information between multi-modal [12] or time-varying datasets [40] is essential to achieve importance-driven visualization. The recently developed transfer entropy [25] can be used to characterize the asymmetrical correlations between two time-series, and has proven to be effective for volume visualization [39], neuroscience [35] and social media analysis [34]. More recently, an information-aware framework was introduced to explore multivariate datasets [1]. Our approach advances the scheme by exploiting the temporal variations of a large-scale multi-variate time-series.

2.3 Time-series Data Structuring and Visualization

Visualizing and structuring time series has been a classical research topic. The standard visualization for linear time dimension would be a two dimensional plot: one axis for time, the other for data value. Weber et al. [42] designed a spiral-shaped time axis where careful selection of cycle length could reveal the cyclic pattern of the data. If the time dimension refers to date, a calendar view [32] can be adopted to visualize the value changes in different days. Tominski et al. [30] employed parallel coordinates to represent time series.

Structuring time is an important scheme to capture the semantic evolution along the timeline. For instance, the time line structure [20] is commonly used to represent events, activities or even status. Storyline [19] and ThemeRiver [13] can be used to represent the evolution of multiple-variables. We employ a ThemeRiver-like structure to display the computed correlation time-series. Other data structures like trees and graphs can also be used to depict the time-oriented evolution structure. For instance, a tree structure is automatically generated to incorporate animations into time-varying data for illustrative narration. The event graph [23] is widely applied to capture the connections among different time pieces. Likewise, a TransGraph [10] was designed to organize a time-varying volume data set into a hierarchy of states and visualize the resulting transition relationships. A pioneering work similar to ours is the time-space partitioning (TSP) tree that reformulates a time-varying volume dataset to a nested tree structure for the purpose of sparse data traversal and rendering acceleration. To the best of our knowledge, the proposed data structure is the first to characterize the space of time, variable and correlation in an information-theoretic way.

3 ANALYTICAL TASKS & DESIGN GOALS

In this section, we first describe the features of the air quality data, thereafter, we introduce the analytical tasks of analyzing the evolution of correlations among air quality variables, then we summarize the design goals of the system to fulfill these tasks.

3.1 The Time-Correlation-Variable Space

Generally, air quality data is a set of time-series of weather variables and air pollutants monitored by multiple sensors. The complete set of symmetrical and asymmetrical time-varying correlations among these time-series spans a triple space, called the *time-correlation-variable* (TCV) space, whose three dimensions are time, variable and correlation, respectively (see Figure 1 (c)).

In particular, variables are the basic units, and can be grouped from multiple perspectives, namely, locations, sensor types, and variable types. We denote the hierarchical organization of variables in terms of locations, sensor types, and categories as the *variable space*. This actually organizes the TCV space along the variable dimension, as illustrated in Figure 1 (c). Meanwhile, along the time dimension a time-series can be recursively subdivided into a hierarchical *time space* by either exploiting their temporal coherence or using a uniform subdivision scheme.

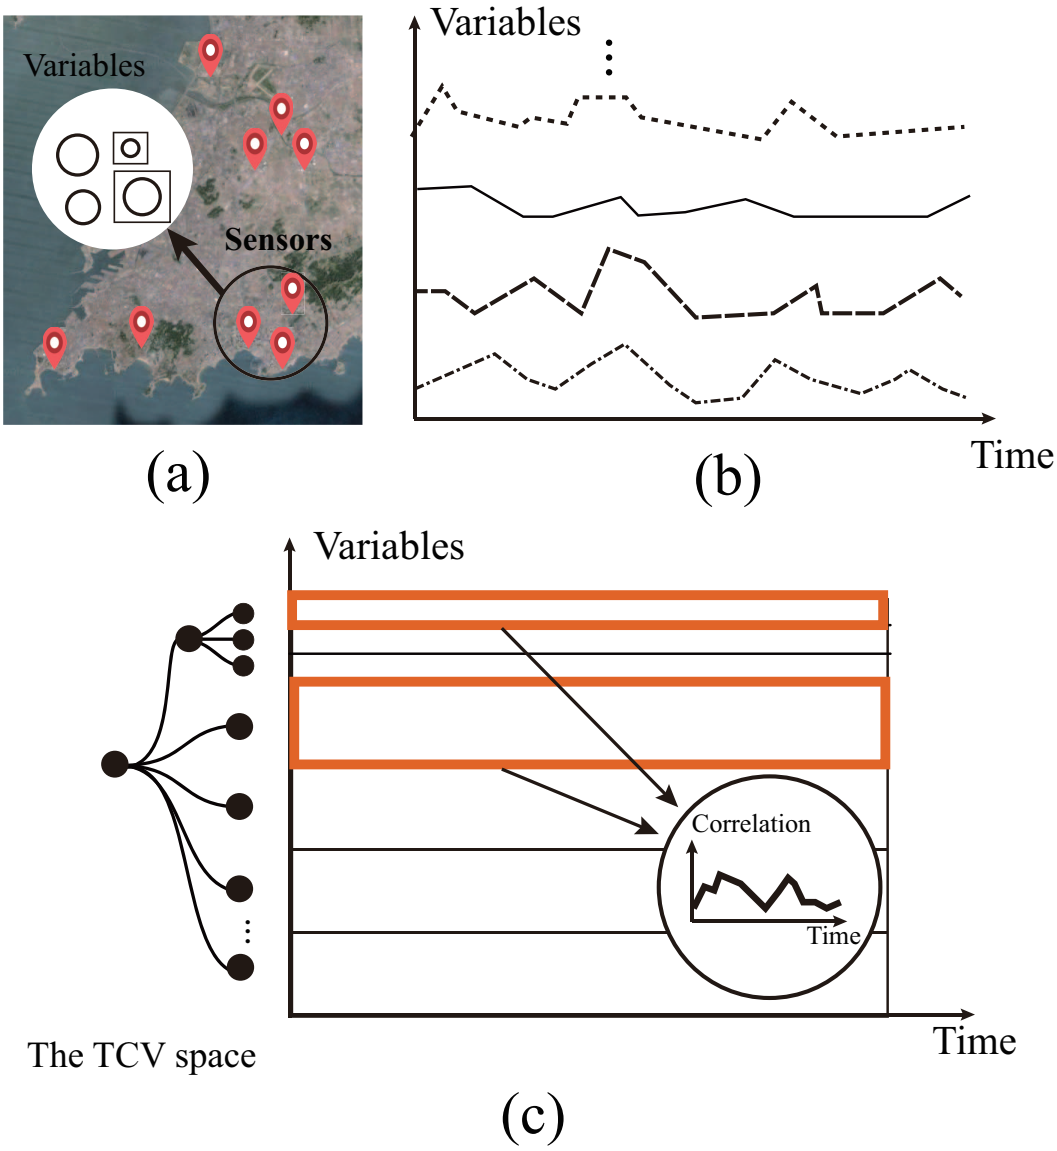

Fig. 1. (a) Sensors and variables; (b) The input dataset, e.g., time-series of senor readings; (c) The time-correlation-variable space converted from the input dataset.

3.2 Analytical tasks

We first summarize the analytical tasks of analyzing the correlations among different air quality variables and time steps.

- **Identify the time slices when the correlations among air quality variables are significant.** The correlations of air quality variables are time-varying, therefore, there are time slices when the correlations are weak and strong. Filtering the time slices according to the correlation strength and

identifying when variables have significant correlation is the basic task when analyzing temporal correlations among variables.

• **Discover the periodic patterns of the correlations of air quality variables.** The correlations of air quality variables may change periodically along time. Discovering these periodic patterns helps meteorologists understand the evolution pattern of the air quality.

• **Discover the transition patterns of the asymmetrical correlations among air quality variables.** For asymmetrical correlations, transition patterns (i.e. X affects Y and Y affects Z) may exists among the variables. Discovering these patterns helps meteorologists explore the order of importance of air variables.

3.3 Design Goals

Due to the complexity of the data space of the correlation among air quality variables, directly visualize the correlations is not capable for fulfilling the analytical tasks. Therefore, we designed a novel data structure, named time-correlation partitioning (TCP) tree, to organize the data. In order to assist users to accomplish all the analytical tasks, we summarized the design goals of the visualization of TCP tree, which includes three aspects:

• **Hierarchy** While it is easy to present the TCV space as a large pixel-based map or streamgraph, depicting the entire dataset with a hierarchically abstracted bundle of informative pieces makes the understanding and exploration of the dataset more efficient.

• **Informativeness** The key of the visualization is to provide a clear view of the evolution and variations between correlations of multiple variables. Two representative entropy measures (symmetrical or asymmetrical) are encoded.

• **Completeness** The visualization should be self-contained, i.e., present all relevant information, including the overall structure of the TCV space, the detailed correlation time-series of all variable pairs and their aggregations, as well as the temporal coherence and variations.

4 INFORMATION-THEORETIC CORRELATIONS

We employ the concept of information entropy to represent the correlations between two time-varying sequences within a specific time interval. Below we first briefly describe two types of information entropy measures.

For two time-series $X = (x_1, x_2, ..., x_n)$ and $Y = (y_1, y_2, ..., y_n)$, $1 \leq n \leq m, n, m \in \mathbb{N}$, the *mutual information* I(X;Y) is defined as:

$$I(X;Y) = I(Y;X) = H(X) + H(Y) - H(X,Y) \tag{1}$$

Here, $H(\cdot)$ denotes the entropy of a time-series. Equation 1 explains the reduction in the uncertainty of X due to the knowledge of Y [6] and vice versa. Note that the mutual information between two time series is measured along an identical timeline.

If the time delay of the information transfer is taken into account, or say, consider the information transfer from the time series X to the time series Y in a time interval, the *transfer entropy* [25] from X to Y can be defined as:

$$
\begin{aligned}
T_{X\to Y} &= \sum_{1\le n\le m} p(y_{n+1}, y_n^{(l)}, x_n^{(k)}) \log \frac{p(y_{n+1}|y_n^{(l)}, x_n^{(k)})}{p(y_{n+1}|y_n^{(l)})} \\
p(y_{n+1}|y_n^{(l)}, x_n^{(k)}) &= \frac{p(y_{n+1}, y_n^{(l)}, x_n^{(k)})}{p(y_n^{(l)}, x_n^{(k)})} \\
p(y_{n+1}|y_n^{(l)}) &= \frac{p(y_{n+1}, y_n^{(l)})}{p(y_n^{(l)})}
\end{aligned}
\tag{2}
$$

where $x_n^{(k)} = (x_n, ...x_{n-k+1})$ and $y_n^{(l)} = (y_n, ...y_{n-l+1})$ denote the past states of X and Y with two Markov processes of order k and order l, and $p(\cdot)$ denotes the proportion of a specific sequence in X and Y. Note that k and l are two adjustable constants. Please refer to [25] for more details.

In principle, the transfer entropy explains the reduction of uncertainty in Y due to the past states of X. $(T_{X\to Y} - T_{Y\to X})$ indicates the dominant strength of influence from X to Y. Thus, we can judge that X influences Y if it is larger than zero and vice versa. We denote $TD(X,Y) = (T_{X\to Y} - T_{Y\to X})$ as the *transfer entropy difference* and $TS(X,Y) = (T_{X\to Y} + T_{Y\to X})$ as the *transfer entropy summation*. When $TD(X,Y) > 0$, we say X affects Y, and when $TD(X,Y) < 0$, we say Y affects X.

Essentially, $I(X;Y)$ represents a symmetrical correlation, while $TD(X,Y)$ encodes an asymmetrical correlation. Each measure computes a numerical value for two time-series whose time ranges are supposed to be limited. For long-term time-series, we sample the entire time range with a sequence of shifted windows where the shift and width are denoted by Δt and w respectively. We compute the *correlation time-series* with respect to the window sequence.

Temporal aggregation of correlations calculated by aggregation operations, such as sum, average, median, peak, and valley, has three forms, including **one-to-one**, **one-to-many**, and **many-to-many** correlations. An one-to-one correlation is the correlation time-series between a pair of variables, which can be summarized into a single value by applying an aggregation operation. An one-to-many correlation is a vector formed by the one-to-one correlations between a variable and a set of variables, and shows the summarized correlation of one specific variable. A many-to-many correlation is a matrix formed by the one-to-one correlations between two sets of variables, and overviews the correlations among all variables.

5 TIME-CORRELATION PARTITIONING TREE

In this section, we firstly introduce the structure of TCP tree and how it organizes the correlation data, and then introduce the two partition operations for constructing the TCP tree.

5.1 The Tree Structure

The TCP tree is designed to characterize the time-varying correlations hierarchically in both the variable domain and the temporal domain. In each tree node, its associated correlation time-series and temporally aggregated correlations are recorded. In particular, there are three types of tree nodes in the tree corresponding to the three forms of the temporally aggregated correlations. An one-to-one node shows the aggregated correlation between two variables. An one-to-many node shows the aggregated correlations among one variable and a set of variables. A many-to-many node shows the aggregated correlations between two sets of variables. The state diagram of the node type and partitioning operations is shown in Figure 2. Time partitioning divides the temporal domain of one-to-many nodes and one-to-one nodes into multiple segments. Each segment is a child node of the divided node.

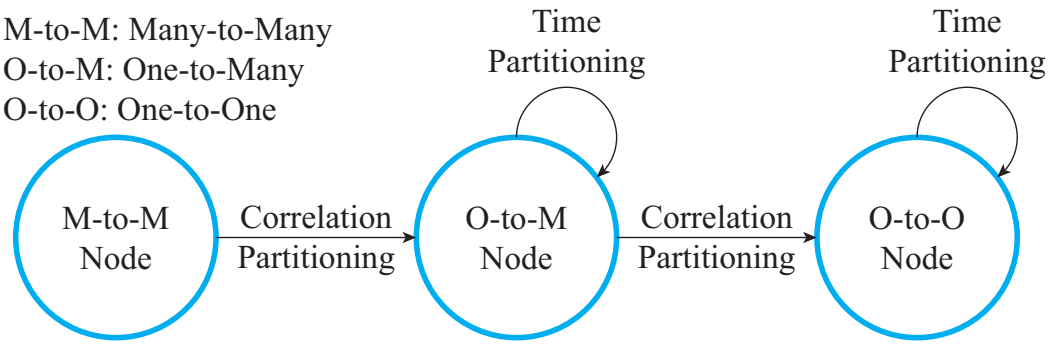

Fig. 2. The state diagram of node types and partitioning operations.

The initial state of a TCP tree, which is also the root of a TCP tree, represents the aggregated correlations among all variables along the whole time axis and thus it is a many-to-many node, showing an overview of the correlations. The root node is partitioned into multiple sub-nodes iteratively by a sequence of time partitioning and correlation partitioning operations. After each partitioning operation is applied, new nodes are generated and appended to the partitioned node as its children.

The order of the correlation partitioning follows the variable hierarchy. The variable hierarchy can be built upon natural hierarchy of the underlying dataset, including:

- The variable types that are relevant to the monitored objects (e.g., PM10 and PM2.5 belong to Particle Matter series) can be used to group variables or sensors;
- The relations among sensors (e.g., the sensor network) and the dependency of variables to sensors, can be used to group variables;
- Environment-related factors and the spatial locations of sensors can be used to categorize variables associated with sensors;
- A group of sensors, or a group of variables can be hierarchically organized based on the domain experience or analysis tasks.

In the following sections we elaborate the correlation partitioning and the time partitioning operations.

5.2 Correlation Partition & Time Partition

The construction of the TCP tree is a dynamic procedure. Two different partition operations are applied to tree nodes iteratively, including correlation partition and time partition. The correlation partition is capable on many-to-many nodes and one-to-many nodes, and the time partition is capable on one-to-many nodes and one-to-one nodes.

Correlation partition is based on the hierarchical structure of the data. Without time partition operations, a many-to-many node is partitioned to a set of one-to-many nodes and then a one-to-many nodes is partitioned to a set of one-to-one nodes. In this way, users are enabled to concentrate on different correlation sets with correlation partitions. The reason why we supports users to partition and explore the TCP tree hierarchically is that the results of time partitions on one-to-many nodes and one-to-one nodes are different.

Time partition is applied on one-to-many nodes and one-to-one nodes in variable tree. In a TCP tree, each tree node is attached with a correlation time-series. By time partition, time pieces with significant relevance are extracted, each of which is appended to a new tree node. A threshold-based

method is designed to obtain the user-interested partitions, e.g. the partitions have high $TD(X,Y)$ and so on, as shown in Figure 3. For one-to-one nodes, the method is directed applied. For one-to-many nodes, one critical problem is the conflict process of different partitions of all involved time-series, because each one-to-one correlation time-series can have independent partition scheme. We solve this problem with a two-stage process. In the first stage, the threshold-based method is applied to all correlation time-series in the node. In the second stage, we construct a joint partition by leveraging the obtained partitions in the first stage, and then build a time tree based on the partition.

In this way, the structure of TCPTree is dynamically built according to the partition operations.

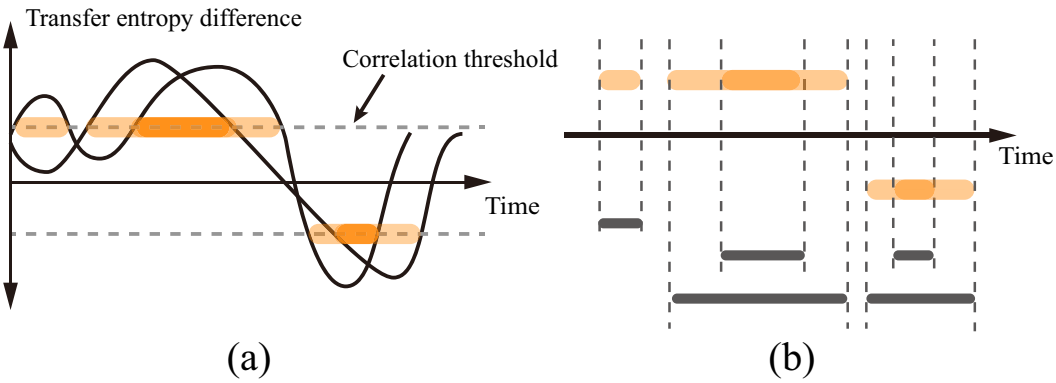

Fig. 3. Adaptive partition demonstrated with two asymmetrical correlation time-series. (a) Filtering the time-series with a high correlation threshold. (b) Constructing a joint partition by different operations.

6 VISUALIZATION OF A TCP TREE

A TCP tree is an information-theoretic, compact, and hierarchical characterization of the correlations among air quality variables.

6.1 The tree structure

The tree structure is dynamically built and modified by a series of user-defined partitioning operations. The initial design of the visualization of the tree structure is a single node link diagram with the layout shown in Figure 4. However, there are two major problems of this design. First, it lacks space efficiency. Second, it cannot offer flexible navigation after multiple partition operations.

Thus, we use a more compact design which combines a sunburst diagram and a node-link diagram to solve the two problems (see in Fig. 5). Initially, a single node is used to represent the entire set of the correlations, which contains many-to-many correlations. After partition is applied, the partition result is represented by a group of nodes surrounding a sunburst diagram (see in Fig. 5 A), which represents the tree structure before partitioning. Nodes generated by the partition surround the sunburst diagram and segments on the helix in a node are preview of the time partition result on the node, as shown in Figure 5 B.

Because each tree node corresponds to a set of or a single correlation time-series, it is necessary to label the time range. A variable tree node extends the same time range as its parent node, however, a time tree node only contains one of the segments partitioned from the its parent node. A helix outside each node is used to show the time range of the time-series. Segments are added to the helix to give a preview of potential result of time partition on the node, as shown in Figure 5.

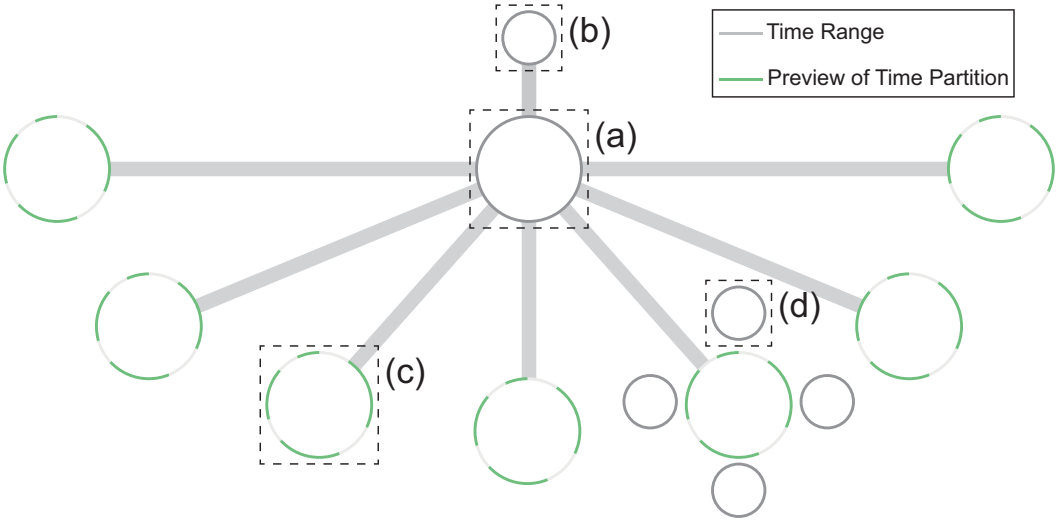

Fig. 4. The first visual design of the tree structure. (a) The latest partitioned node; (b) the ancestor of the latest partitioned node; (c) newly generated node by partitioning with a preview of time partitioning operation; and (d) the nodes generated by time partitioning operation

In each tree node, no matter it is a variable tree node or time tree node, it supports users to freely modify the variables in it. For example, in a many-to-many variable tree node, users can remove PRESS and WS to explore the correlations among the remaining variables. The modification of variables will effect the result of following partition operations as the set of variables is consistent between parent node and leaf nodes.

6.2 Aggregated correlations

To increase the readability of the tree nodes of TCPTree, aggregated correlations are visualized inside each node. The first design of the aggregated correlations inside the tree node is shown in Figure 6. However, we identify two major limitations in this design. First, the length of the links interfere with users' cognition of the strength of the relations, because of the varying lengths of lines without any information encoded. Second, users' have to switch between the asymmetrical correlation and symmetrical correlation repeatedly for comparison.

Therefore, we improve the design by using different layouts with the same glyph design and color encodings to visualize the three types of nodes, as shown in Fig. 7. The time partition results of one-to-many node and one-to-one node are different, it is the reason why we still remain the one-to-one nodes in the TCP tree node.

- **Many-to-many correlations.** The basic visual scheme of the temporally aggregated many-to-many correlations is a radial node-link graph. It is an overview of correlations among variables in the TCP tree. Each attribute is a node and all the nodes are uniformly distributed on a circle. To avoid the visual clutter, the correlations among attributes are represented by small rectangles distributed around the node, as shown in Figure 7. For example, the aggregated correlations among CO and other attributes are encoded by a group of small glyphs (see in Figure 7). Each glyph is formed by a circle and a rectangle. The color of the circle encodes $I(X;Y)$; and the length of the rectangle encodes $TS(X,Y)$. The rectangle is divided into two sub-rectangles, the length of the one close to the circle encodes $T_{X\to Y}$ and the length of the other encodes $T_{Y\to X}$. The brightness of two sub-rectangles

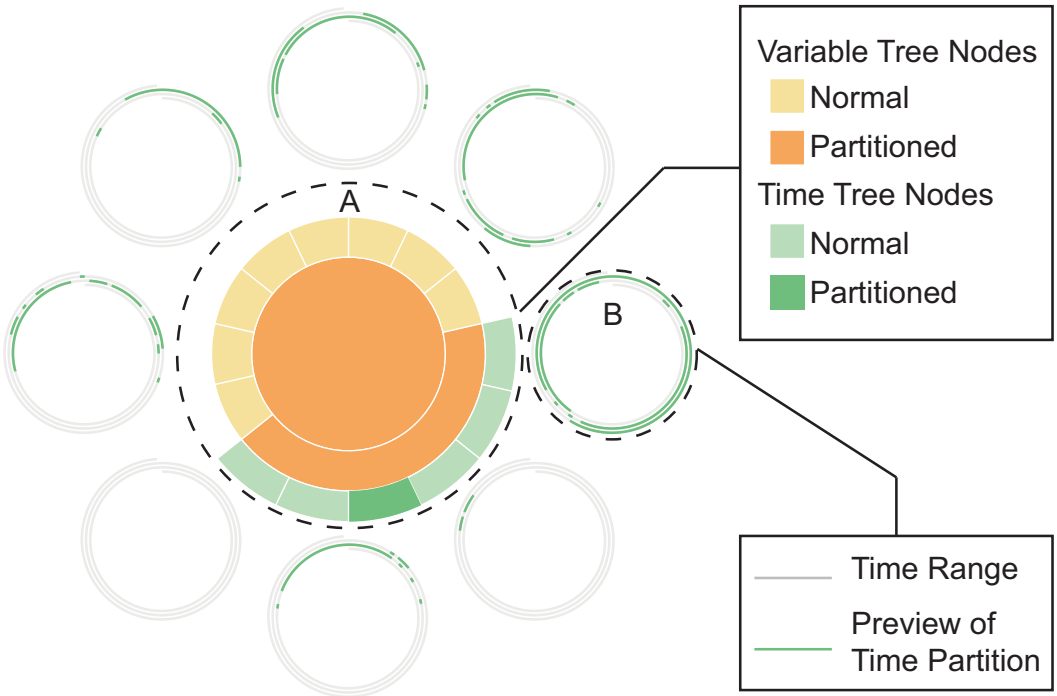

Fig. 5. A TCP tree consists of two parts. Part A is the structure of the tree which are formed by the tree nodes generated by partitioning before the latest partitioning (green nodes are time tree nodes and yellow nodes are variable tree nodes). Part B is the tree nodes generated by the latest partitioning, the green segments on the helix are the preview of the time partitioning.

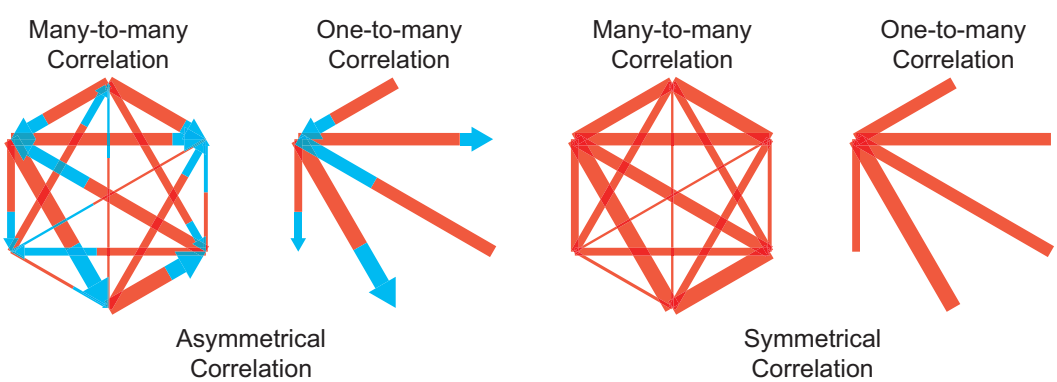

Fig. 6. The original design of aggregated correlations. Asymmetrical correlations are represented by directed links and symmetrical correlations are represented by undirected links.

encodes $TD(X,Y)$. In this way, a pair of variables which have large $TS(X,Y)$ and large $TD(X,Y)$ is actually highlighted from others.

- **One-to-many correlations.** For one-to-many correlations, say the correlations among variable X and a set of variables S, a radial layout is used: X is placed on the center and variables in S

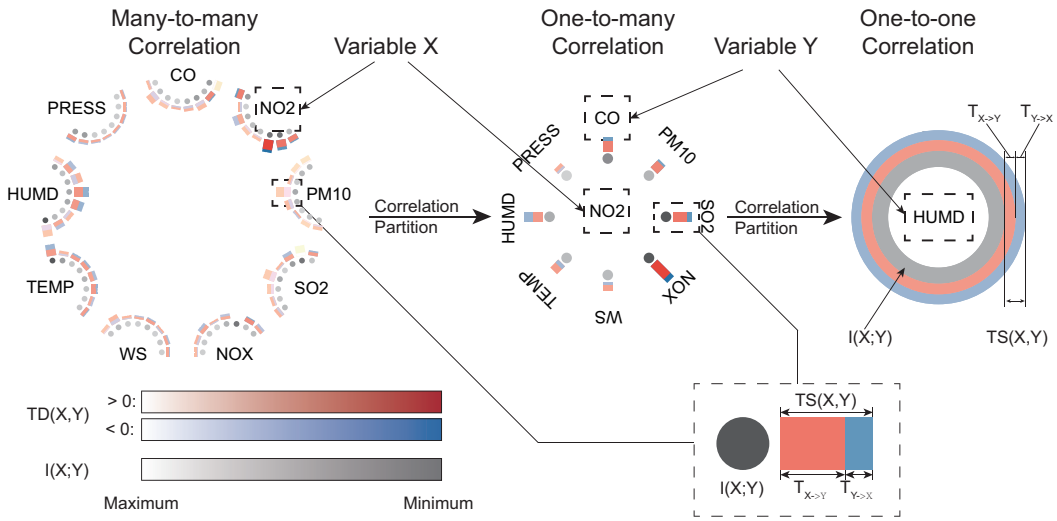

Fig. 7. The visual design of three types of correlations. For many-to-many nodes and one-to-many nodes, we use a glyph design to encode $TD(X,Y)$, $TS(X,Y)$, $T_{X\to Y}$, and $T_{Y\to X}$. For one-to-one nodes, we use three circles to encode these values.

are placed around X. The aggregated correlations are represented by rectangles the same as the many-to-many correlations, as shown in Fig. 7.

- **One-to-one correlations.** The visual encoding of one-to-one correlations is very simple, as shown in Fig. 7. For a pair of variables (X,Y), the name of variable Y is placed in the center. Three concentric circles from the inside out represent $I(X;Y)$, $T_{X\to Y}$, $T_{Y\to X}$ respectively, with the color mapping shown in Fig. 7.

6.3 Details of correlation time-series

In tree view, the hierarchy and aggregated correlations are visualized, however, the details of the correlation time-series are still missing. For completeness, the details are visualized by colored 2D *pixel maps* and modified line charts, as shown in Fig. 8 (b). When using 2D pixel maps, each asymmetrical correlation time-series ($T_{X\to Y}$ and $T_{Y\to X}$) is represented by two rows of pixelbars and each symmetrical correlation time-series is represented by a single line of grey pixelbars. The modified line chart is used to show the $TD(X,Y)$. After the pixelmap of the asymmetrical correlations between X and Y is expanded, the modified line chart is shown.

When using modified line charts, $T_{X\to Y}$ and $T_{Y\to X}$ are represented by two polylines and $TD(X,Y)$ is emphasized by the colored region (the region between the two polylines). Users can freely switch between the two visualization forms by clicking.

Initially, the order of the pixel maps and line charts is decided by the attribute order in the data. While users are exploring the correlation in the tree view, associated pixel maps and line charts will be highlighted and reordered to the top of the list. When time partition operations are performed in the tree view, associated pixel maps are firstly expanded to line charts and all associated line charts are reordered to the top. Grey rectangles are added on line charts to label the partitioned time regions.

7 VISUAL EXPLORATION WITH THE TCP TREE

The TCP Tree is dynamically built by the combination of correlation partitioning and time partitioning. Users can freely filter the variables and change the time partitioning parameters before applying partitioning operations.

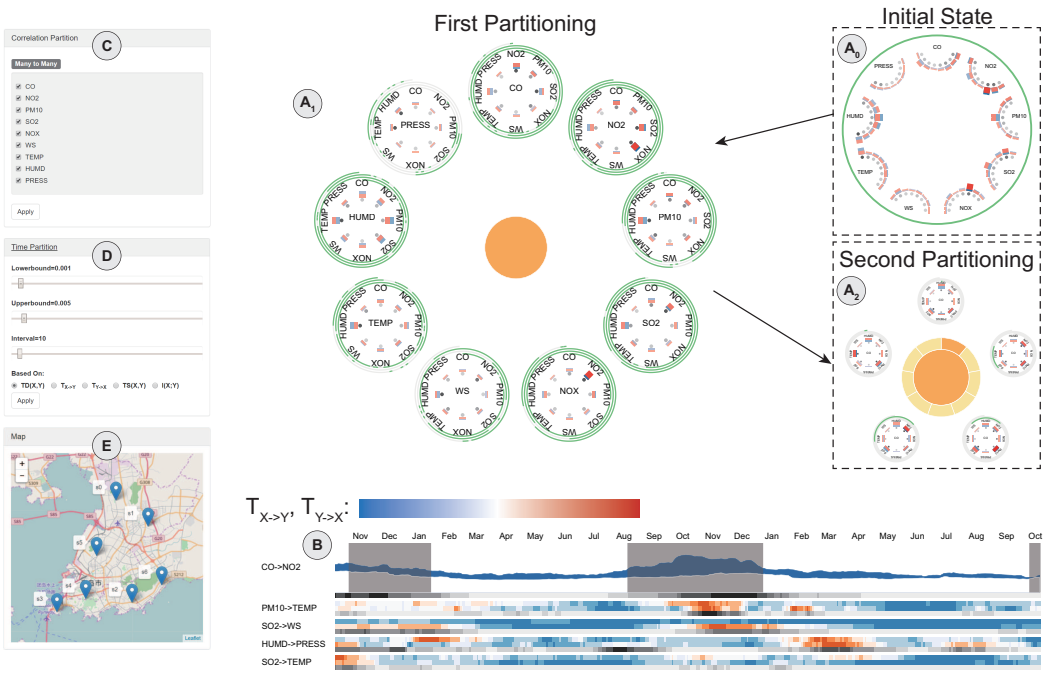

Fig. 8. Information-theoretic visualization of the air quality data with our TCP tree structure. A0, A1 and A2 are three sequential states of a time correlation partition (TCP) tree view. A0 is the initial state of the TCP tree, and represents the aggregated correlations among all variables along the entire time axis. A1 is the state after applying a correlation partition on A0, and A2 is the state after applying a time partition on the node CO in A1. B denotes a hybrid visualization of pixel map and line chart and shows the details of variable correlations. C and D are the parameter panels for variable-oriented and temporal partitions, respectively. E is a map that shows the spatial distribution of the sensors.

7.1 Integrated Visual Interface

We implemented the visual interface by Javascript with D3 and angularjs. The integrated visual interface (Figure 8) consists of three views: a *TCP tree view*, a *correlation time-series view*, and a control panel. All views are coordinated with dynamic query interactions.

The TCP tree view depicts a TCP tree, by which users can freely construct the tree structure and explore the correlations among arbitrary combination of the air quality variables (top of Figure 8).

The correlation time-series view shows all the one-to-one correlation time-series with both a pixel based visualization and a modified line chart. When time partition is applied in the TCP tree view, all related one-to-one correlations will be expanded to line chart form and the time slices will be highlighted in the view (bottom of Figure 8). Initially, this view is collapsed and is expanded when users click the expand button for more details of the correlation time-series.

The control panel provides a interface for filtering variables and adjusting the parameter of time partitioning. Once users select a node in the TCP tree view, the control panel will show the variables in correlation partitioning sub-panel and partitioning parameters in time partitioning sub-panel (top left of Figure 8). In the correlation partitioning sub-panel, users are enabled to filter the variables by checking or unchecking the variable names. In the time partitioning sub-panel, users are enabled to adjust the parameters in the time partitioning procedure, including the merge interval, upper bound of the threshold, and the lower bound of the threshold. The control panel also provide a map to show the position of the air quality monitoring sensors.

7.2 User Interactions and Explorations

Tree Traversal Supported interactions in the TCP tree include navigation, specification, and expanding. The user is expected to first traverse the variable tree nodes, and then navigate time tree nodes by visually studying and comparing the temporally aggregated correlation time-series. Therefore, interesting one-to-many and one-to-one correlation patterns might be revealed. Once a certain time node is specified, the related correlation time-series will be highlighted and moved to the top of the correlation time-series view.

Hierarchical Analysis of Aggregated Correlations Starting from analyzing many-to-many correlations depicted in the root node, nodes indicating one-to-many and one-to-one correlations can be gradually obtained by continual partition operations whose parameters determine the examine levels of correlations. The temporally aggregated correlations attached in time tree can be either one-to-many or one-to-one correlations that respectively summarizes the correlations among variable sets and variable pairs in a time interval. In addition, comparing the aggregated correlations of a sequence of time intervals is made easy because they are visually aligned around the central node.

Temporal Analysis of Correlation Time-series On one hand, the correlation time-series map is supposed to respond to interactions in the TCP tree view, such as node selection, in support of correlation analysis between various variables; on the other hand, the correlation time-series map provides guidance for correlation analysis by indicating not only the temporal trend, but also the temporal variations in terms of symmetrical and asymmetrical correlations. Users are allowed to traverse along the time line and investigate the aggregated correlations around a time point with the help of the correlation map.

8 CASE STUDY

We applied our approach to a real-world Air Quality dataset. This dataset was collected in 7 observation stations of a modern city (8 million citizens) in 3 years (2009-2011). Each station contains several sensors, recording 9 variables every hour: 5 air pollutants (CO, NO2, PM10, SO2, NOX) and 4 weather variables (Speed of Wind (WS), Temperature (TEMP), Humidity (HUMD), Pressure (PRESS)). A pre-processing stage is adopted in our research to filter anomalous cases such as silent samples, faults and outliers. After that, a TCP tree of either air quality variables or sensors are generated. All the experiments were conducted on a PC with 3.2GHz dual core, 8G memory.

8.1 Parameters

The integrated system requires tuning a sequence of user-adjustable parameters, which can be classified into two categories.

In terms of the variable tree, the variable hierarchy can be determined by users before a correlation partitioning is applied to a tree node by checking or unchecking variables in the parameter panel, as shown in Figure 8 C.

For time tree, the time-based partition of a correlation time-series is performed in an adaptive fashion. Four parameters of time partitioning are user-adjustable, including the two correlation thresholds, time interval and partitioning basis (see in Figure 8 D).

8.2 Case 1: Correlations among variables

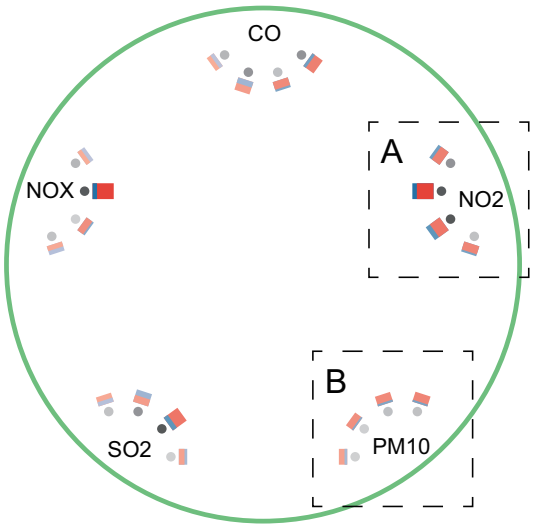

Fig. 9. Correlation summarizations among five pollutants shown by a many-to-many node. A. The aggregated asymmetrical and symmetrical one-to-many correlations of NO2. B. The aggregated asymmetrical and symmetrical one-to-many correlations of PM10.

Using the TCP tree, the analyst started by checking the aggregated correlations among variables. To analyze the correlations between pollutants, he hided the weather variable nodes. By inspecting both asymmetrical and symmetrical correlations with the tree view, he quickly found that NO2 dominantly influences NOX and SO2 while the other correlations are relatively weak (Figure 9 A). This is reasonable because NO2 is the dominant part of NOX and is released mainly by combustions such as vehicles or power plants. Further investigation into the NO2 variable suggests that there exists interesting correlations depicting an influence chain from NO2 to CO, and finally to PM10. Interestingly, correlations associated with NO2 and PM10 seem to be different (Figure 9). In contrast to NO2, which affects other variables, PM10 is affected by other variables. After examining the one-to-many correlations, the analyst discovered that NO2 is a strong influence factor on other pollutants while PM10 barely receives significant influences from others because of the weak linking edges. According to the analysis process, the analyst ultimately drew the conclusion that CO and NO2 play a predominant role in contributing to the release of PM10. This conclusion makes sense because PM10 is mainly caused by coal-based combustion.

To further investigate how weather variables influence the pollutants, the analyst studied all weather variables. Almost all the strongest asymmetrical correlations come from the HUMD (Figure 10(a)). The one-to-many aggregated correlations indicate that though the mutual information between each variable pair is strong, it is hard to determine the influence direction because the difference of the asymmetrical correlation is too small. Thus he selected the HUMD to check the one-to-many correlation time-series. The adaptive partition with a threshold yields several interesting findings which show different patterns (Fig. 10(b)). He studied the consecutive time spans and realized the

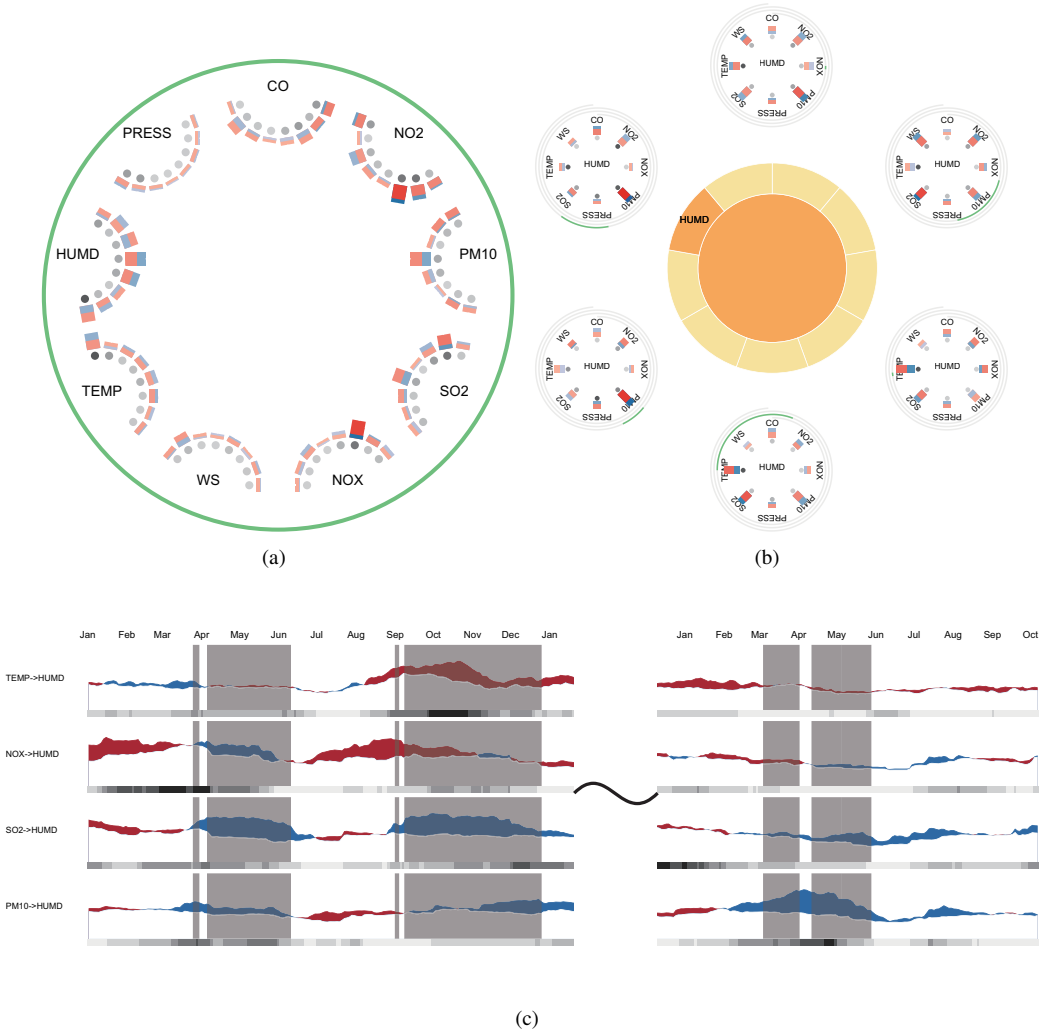

Fig. 10. Navigating the time tree node relating to HUMD. (a) The aggregated many-to-many correlations of all the pollutants and weather variables. (b) The time partitioning result of the aggregated one-to-many correlations of HUMD. Several consecutive partitions are exploited. (c) The correlation time-series that correspond to consecutive partitions in (b).

the unsteady mutual influence from HUMD to the other variables, among which TEMP always has a high symmetrical correlation. He checked the details of the correlations and corresponding time slices in the detailed view (Fig. 10(c)). According to the time tree node, most pollutants are influenced by HUMD as the color of the corresponding rectangles are relatively deeper.

The analyst continued his exploration by selecting other variables. Surprisingly, the one-to-many correlation time-series of all pollutants appear to have a similar trend from September to November every year. For example, according to the details of the correlations among CO and other variables, almost all correlations increase among these months in 2009 and 2010, especially for NO2 and PM10 (Fig. 11(a)). BY time partitioning, two time slices are obtained (see in Fig. 11(b)). From the

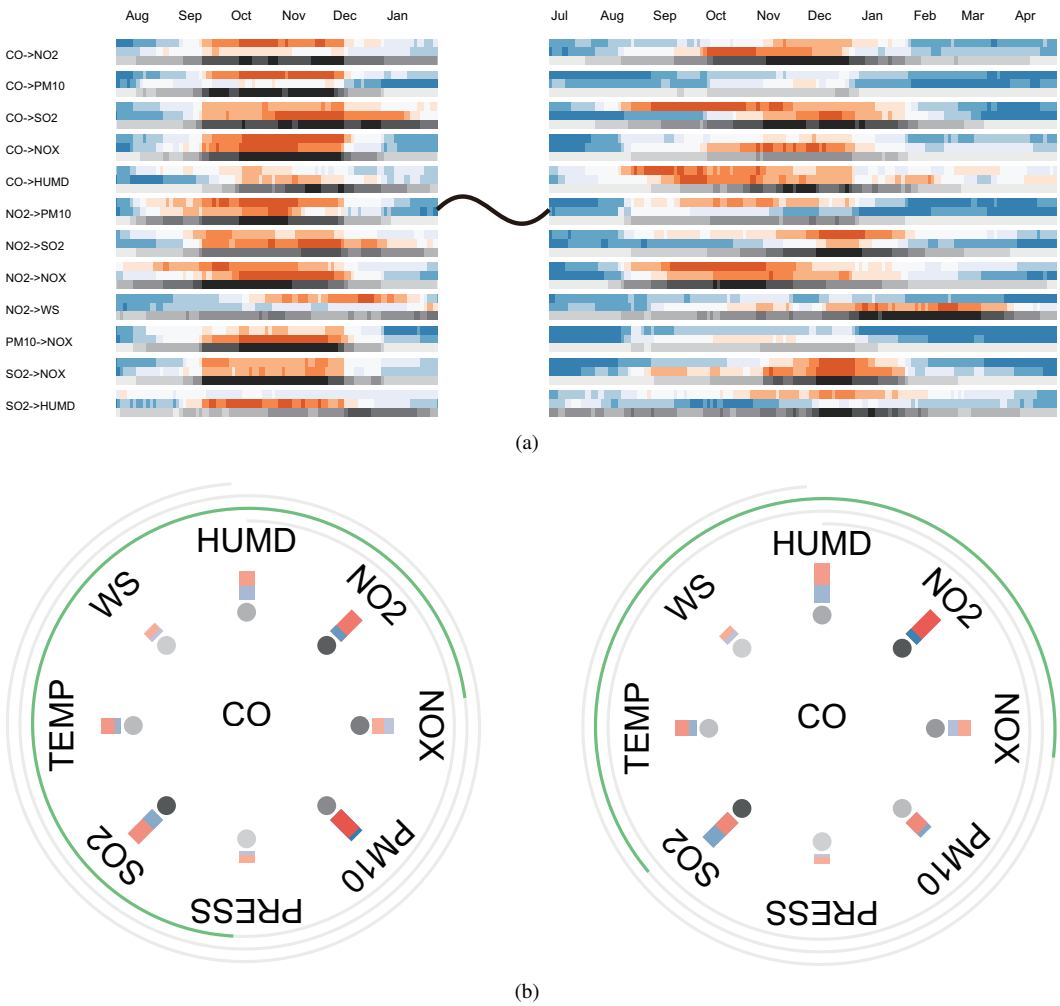

Fig. 11. (a) The correlation time-series for CO. (b) The partitions and their correlation maps. A clear periodical pattern can be seen: CO and other variables have strong correlations from September to December.

aggregated correlations, it is clear that CO significantly influences PM10 and has strong correlation with SO2 (for large $I(CO;SO2)$ and small $TD(CO,SO2)$), but it is influenced by NO2.

One possible reason might be the coal-based centralized heating service provided by the government in the city, which starts at every September, and is turned on or off depending on the actual temperature. If the service is on, it consumes a vast amount of carbon energy, and causes a dramatic increase of pollutants. The truth that the centralized heating service is not stable in these months may be related to the variations of correlations among air pollutants.

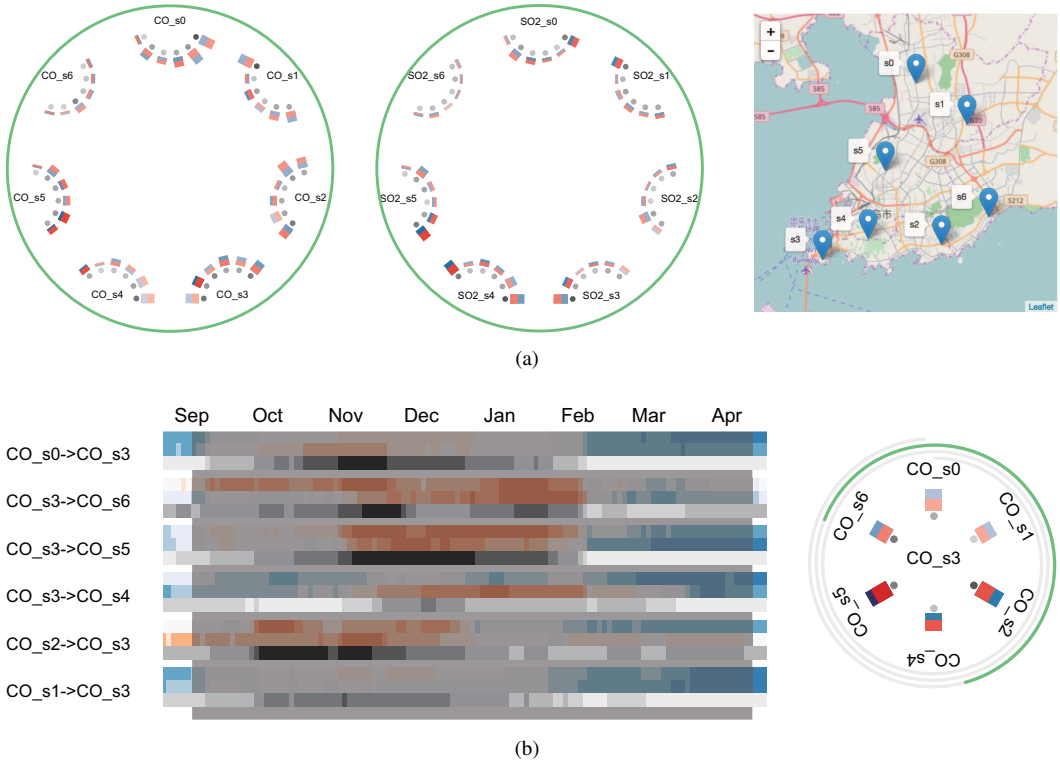

Fig. 12. (a) Aggregated correlations of s6 (sensor 6) under variable SO2 and CO and the position of sensors. (b) The partition based on s3 highlights time intervals in which s4 influenced s3, while s3 influenced other sensors.

8.3 Case 2: The correlations among observation stations

The analyst decided to navigate the second level of the variable tree to discover the correlations among observation stations. He expanded the tree node associated with one variable (e.g., CO). The tree view shows the correlations among observation stations. Two dominant influences (from s3 to s5 and from s4 to s5) attracted his attention (Fig.12(a)). By referring to the map, he found that these three stations are quite near (right of Fig.12(a)). When further studying other correlations, the analyst discovered the correlation pattern among observation stations based on CO. It seems that s3 always influenced the others. The streamgraph verified this observation except for the months from October 2010 to February 2011. In these months, s4 influenced s3 significantly (Fig.12(b)). To find more evidences, the analyst studied SO2. A similar pattern appeared (Fig. 12(a)). He concluded that s3, s4 and s5 are likely to be three interrelated pollution areas. The pollutants CO and SO2 probably spread from s3 and s4 to s5 while the influence between s3 and s4 changed in a specific time period.

Furthermore, the analyst found that s6 may be a special station that is seldom influenced by others. He then checked the sensor positions in the map and found s6 to be quite near a mountain in the southeast. The unique geographical location could be the reason of its low dependence.

9 CONCLUSION

This paper presents a novel data structure called TCP tree that captures both the variable hierarchy and the temporal variation of correlations hidden in the air quality data. The case study on a real-life dataset verifies that such a hierarchical structure can help exploit the sparsity of a large-scale air quality time series in an information-theroetic way.

As future work is concerned, we believe that the proposed hierarchical structure can be extended for characterizing dynamic network structures and applied in analyzing senor network data. A hybrid network-tree structure can be suitable for this scenario. In addition, we plan to extend the proposed method to other time-varying dataset, e.g., time-varying volume or flow dataset.

REFERENCES

[1] A. Biswas, S. Dutta, H.-W. Shen, and J. Woodring. An information-aware framework for exploring multivariate data sets. *Visualization and Computer Graphics, IEEE Transactions on*, 19(12):2683–2692, 2013.

[2] A. L. Bondi and A. Plaia. Weather variables and air pollution via hierarchical linear models. *Statistics and Environment*, pages 237–240, 2005.

[3] U. D. Bordoloi and H. W. Shen. View selection for volume rendering. In *IEEE Visualization*, page 62, 2005.

[4] S. Bruckner and T. Möller. Isosurface similarity maps. In *Computer Graphics Forum*, volume 29, pages 773–782. Wiley Online Library, 2010.

[5] M. Chen and H. Jaenicke. An information-theoretic framework for visualization. *Visualization and Computer Graphics, IEEE Transactions on*, 16(6):1206–1215, 2010.

[6] T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012.

[7] T. N. Dang, A. Anand, and L. Wilkinson. Timeseer: Scagnostics for high-dimensional time series. *Visualization and Computer Graphics, IEEE Transactions on*, 19(3):470–483, 2013.

[8] M. Demuzere, R. M. Trigo, V. G. D. Arellano, and N. P. M. V. Lipzig. The impact of weather and atmospheric circulation on o3 and pm10 levels at a rural mid-latitude site. *Atmospheric Chemistry & Physics*, 9(2009):2695–2714, 2009.

[9] M. Feixas, E. Del Acebo, P. Bekaert, and M. Sbert. An information theory framework for the analysis of scene complexity. In *Computer Graphics Forum*, volume 18, pages 95–106. Wiley Online Library, 1999.

[10] Y. Gu and C. Wang. Transgraph: Hierarchical exploration of transition relationships in time-varying volumetric data. *Visualization and Computer Graphics, IEEE Transactions on*, 17(12):2015–2024, 2011.

[11] S. Hadlak, H. Schumann, C. H. Cap, and T. Wollenberg. Supporting the visual analysis of dynamic networks by clustering associated temporal attributes. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2267–2276, 2013.

[12] M. Haidacher, S. Bruckner, A. Kanitsar, and M. E. Gröller. Information-based transfer functions for multimodal visualization. In *Proceedings of the First Eurographics conference on Visual Computing for Biomedicine*, pages 101–108. Eurographics Association, 2008.

[13] S. Havre, B. Hetzler, and L. Nowell. Themeriver: Visualizing theme changes over time. In *Information Visualization, 2000. InfoVis 2000. IEEE Symposium on*, pages 115–123. IEEE, 2000.

[14] K. Ito, G. D. Thurston, A. Ndas, and M. Lippmann. Monitor-to-monitor temporal correlation of air pollution and weather variables in the north-central u.s. *Journal of Exposure Analysis & Environmental Epidemiology*, 11(1):21–32, 2001.

[15] G. Ji and H.-W. Shen. Dynamic view selection for time-varying volumes. *Visualization and Computer Graphics, IEEE Transactions on*, 12(5):1109–1116, 2006.

[16] T.-Y. Lee and H.-W. Shen. Visualization and exploration of temporal trend relationships in multivariate time-varying data. *Visualization and Computer Graphics, IEEE Transactions on*, 15(6):1359–1366, 2009.

[17] P. Levis, N. Lee, M. Welsh, and D. Culler. Tossim: Accurate and scalable simulation of entire tinyos applications. In *Proceedings of the 1st international conference on Embedded networked sensor systems*, November 5–7 2003.

[18] J. Lin, E. Keogh, S. Lonardi, J. P. Lankford, and D. M. Nystrom. Viztree: a tool for visually mining and monitoring massive time series databases. In *Proceedings of the Thirtieth international conference on Very large data bases-Volume 30*, pages 1269–1272. VLDB Endowment, 2004.

[19] M. Ogawa and K.-L. Ma. Software evolution storylines. In *Proceedings of the 5th international symposium on Software visualization*, pages 35–42. ACM, 2010.

[20] C. Plaisant, B. Milash, A. Rose, S. Widoff, and B. Shneiderman. Lifelines: visualizing personal histories. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 221–227. ACM, 1996.

[21] H. Qu, W. Y. Chan, A. Xu, K. L. Chung, K. H. Lau, and P. Guo. Visual analysis of the air pollution problem in hong kong. *IEEE Transactions on Visualization & Computer Graphics*, 13(6):1408–1415, 2007.

[22] P. S. Quinan and M. Meyer. Visually comparing weather features in forecasts. *IEEE Transactions on Visualization & Computer Graphics*, 22(1):389–398, 2016.

[23] W. Ribarsky, Z. Wartell, and W. Dou. Event structuring as a general approach to building knowledge in time-based collections. In *Expanding the Frontiers of Visual Analytics and Visualization*, pages 149–162. Springer London, 2012.
[24] J. Rigau, M. Feixas, and M. Sbert. Shape complexity based on mutual information. In *Shape Modeling and Applications, 2005 International Conference*, pages 355–360. IEEE, 2005.
[25] T. Schreiber. Measuring information transfer. *Physical review letters*, 85(2):461, 2000.
[26] R. Sharovsky, L. César, and J. Ramires. Temperature, air pollution, and mortality from myocardial infarction in sao paulo, brazil. *Brazilian Journal of Medical and Biological Research*, 37(11):1651–1657, 2004.
[27] H.-W. Shen, L.-J. Chiang, and K.-L. Ma. A fast volume rendering algorithm for time-varying fields using a time-space partitioning (tsp) tree. In *Proceedings of the conference on Visualization'99: celebrating ten years*, pages 371–377. IEEE Computer Society Press, 1999.
[28] L. Shi, Q. Liao, Y. He, R. Li, A. Striegel, and Z. Su. Save: Sensor anomaly visualization engine. In *IEEE Conference on Visual Analytics Science and Technology*, pages 201–210, 2011.
[29] S. Takahashi, I. Fujishiro, Y. Takeshima, and T. Nishita. A feature-driven approach to locating optimal viewpoints for volume visualization. In *Visualization, 2005. VIS 05. IEEE*, pages 495–502. IEEE, 2005.
[30] C. Tominski, J. Abello, and H. Schumann. Axes-based visualizations with radial layouts. In *Proceedings of the 2004 ACM symposium on Applied computing*, pages 1242–1247. ACM, 2004.
[31] M. Turon. Mote-view: A sensor network monitoring and management tool. In *The 2nd IEEE workshop on Embedded Networked Sensors*, pages 11–17, 2005.
[32] J. J. Van Wijk and E. R. Van Selow. Cluster and calendar based visualization of time series data. In *Information Visualization, 1999.(Info Vis' 99) Proceedings. 1999 IEEE Symposium on*, pages 4–9. IEEE, 1999.
[33] P. Völgyesi, A. Nádas, X. Koutsoukos, and Ákos Lédecz. Air quality monitoring with sensormap. In *International Conference on Information Processing in Sensor Networks*, pages 529 – 530, 2008.
[34] G. Ver Steeg and A. Galstyan. Information transfer in social media. In *Proceedings of the 21st international conference on World Wide Web*, pages 509–518. ACM, 2012.
[35] R. Vicente, M. Wibral, M. Lindner, and G. Pipa. Transfer entropyła model-free measure of effective connectivity for the neurosciences. *Journal of computational neuroscience*, 30(1):45–67, 2011.
[36] I. Viola, M. Feixas, M. Sbert, and M. E. Groller. Importance-driven focus of attention. *Visualization and Computer Graphics, IEEE Transactions on*, 12(5):933–940, 2006.
[37] C. Wang and H.-W. Shen. Lod map-a visual interface for navigating multiresolution volume visualization. *Visualization and Computer Graphics, IEEE Transactions on*, 12(5):1029–1036, 2006.
[38] C. Wang and H.-W. Shen. Information theory in scientific visualization. *Entropy*, 13(1):254–273, 2011.
[39] C. Wang, H. Yu, R. W. Grout, K.-L. Ma, and J. H. Chen. Analyzing information transfer in time-varying multivariate data. In *Pacific Visualization Symposium (PacificVis), 2011 IEEE*, pages 99–106. IEEE, 2011.
[40] C. Wang, H. Yu, and K.-L. Ma. Importance-driven time-varying data visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1547–1554, 2008.
[41] J. Wang, R. Sisneros, and J. Huang. Interactive selection of multivariate features in large spatiotemporal data. In *IEEE Pacific Visualization Symposium*, pages 145–152, 2013.
[42] M. Weber, M. Alexa, and W. M. 0004. Visualizing time-series on spirals. In *IEEE Symposium on Information Visualization*, pages 7–14, 2001.
[43] L. Xu, T.-Y. Lee, and H.-W. Shen. An information-theoretic framework for flow visualization. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1216–1224, 2010.
[44] L. Yu, A. Lu, W. Ribarsky, and W. Chen. Automatic animation for time-varying data visualization. In *Computer Graphics Forum*, volume 29, pages 2271–2280. Wiley Online Library, 2010.
[45] Y. Zheng, F. Liu, and H.-P. Hsieh. U-air: when urban air quality inference meets big data. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1436–1444. ACM, 2013.
[46] Y. Zheng, F. Liu, and s.-P. HsiehH. U-air: when urban air quality inference meets big data. *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1436–1444, 2013.