

Interactive Visual Optimization and Analysis for RFID Benchmarking

Yingcai Wu, Ka-Kei Chung, Huamin Qu, *Member, IEEE*,
Xiaoru Yuan, *Member, IEEE*, and S.C. Cheung, *Senior Member, IEEE*

Abstract—Radio frequency identification (RFID) is a powerful automatic remote identification technique that has wide applications. To facilitate RFID deployment, an RFID benchmarking instrument called α Gate has been invented to identify the strengths and weaknesses of different RFID technologies in various environments. However, the data acquired by α Gate are usually complex time-varying multidimensional 3D volumetric data, which are extremely challenging for engineers to analyze. In this paper, we introduce a set of visualization techniques, namely, parallel coordinate plots, orientation plots, a visual history mechanism, and a 3D spatial viewer, to help RFID engineers analyze benchmark data visually and intuitively. With the techniques, we further introduce two workflow procedures (a visual optimization procedure for finding the optimum reader antenna configuration and a visual analysis procedure for comparing the performance and identifying the flaws of RFID devices) for the RFID benchmarking, with focus on the performance analysis of the α Gate system. The usefulness and usability of the system are demonstrated in the user evaluation.

Index Terms—RFID Visualization, Visual analytics, Visual Optimization.

1 INTRODUCTION

Radio frequency identification (RFID) is an automatic contactless identification method based on radio frequency (RF) transmission and acquisition. It has two major components – readers and a set of integrated circuits called RFID tags. Through RF signals, the readers can control the tags, transmit data from the system to the tags, and receive data from the tags to the system via reader antennae. RFID tags, on the other hand, can store data and modify it according to RF signals from the readers. RFID has been widely used in many applications such as logistics and supply chain management, industrial automation, e-Passport, ticketing, and medical applications.

Although RFID technology has received much attention in the industry in recent years, it is challenging to deploy RFID systems successfully in practice owing to various environmental, material, and device constraints. For example, missing read is a major problem in RFID deployment; this occurs when an RFID reader fails to identify the presence of all the tags in its proximity. It was reported that a small variation of the reader antenna's position and orientation can result in missing read [14], which in real life deployment can reach 30 to 40% [12]. Hence, finding the optimized position and orientation of the antenna is a critical task that is highly environment dependent. In daily RFID deployment, three to four antennae are employed because one antenna usually cannot cover all tags. This, however, makes it very difficult for engineers to find the optimum spatial configuration for the antennae, as there are numerous different combinations of antenna positions and orientations and engineers cannot examine them one by one. Furthermore, RFID devices produced by different manufacturers usually have different performance characteristics in different applications. Some RFID components may even work abnormally in certain environments. Therefore, system engineers have to perform tedious field tests to configure RFID devices for different applications.

To alleviate the problems, RFID benchmarking must be conducted

to evaluate the quality of RFID products and obtain a comparative measurement of RFID components in different environments. With the consequent benchmark reports, end users can quickly identify the strengths and weaknesses of RFID products in different scenarios, such that they can flexibly configure RFID systems and effectively deploy the systems based on their specific requirements. Towards this goal, we have developed an RFID benchmarking instrument called α Gate that is capable of automatic data acquisition, calibration, and measurement to address different constraints. Nevertheless, because of the complex features of the acquired benchmark data (Section 3.2), methods for effectively analyzing the data and finding the hidden patterns are still under intense investigation (Section 3.3).

In this paper, we focus on the real problems of the effective analysis of RFID benchmark data and extend several visualization techniques. We have improved the traditional Parallel Coordinate Plots (PCPs) to intuitively represent spatial information of the data, which otherwise is lost, and to enable the simultaneous visualization of all benchmark data. Orientation plots are introduced to visually encode the orientations of reader antennae and the signal strength of tags in a 2D image, allowing users to visualize the data efficiently without occlusion. A tree-style history mechanism and a 3D spatial viewer are also provided to coordinate with the PCPs and orientation plots in order to aid the visual analysis. With the visualization techniques, we further introduce two independent workflow procedures: a visual optimization procedure for finding the optimum spatial configuration of reader antennae and a visual analysis procedure for comparing the performance and identifying the flaws of RFID devices. Our system can help users identify the strengths and weaknesses of different RFID components and configurations in different applications. This would foster the wider adoption of RFID technologies in the future.

2 RELATED WORK

This project involves interdisciplinary research, which includes visual analytics, visualization, and RFID research.

RFID benchmarks are important for the development and deployment of RFID systems. Various industrial companies have conducted independent benchmarks and made commercial RFID benchmark reports [1, 18] measuring the performance of RFID devices. However, the testing methodologies have not been made public so far. In contrast, Cheung et al. [4] developed a combinatorial approach to reliably benchmark RFID devices and published their solution. While existing RFID benchmarking methods primarily focus on the systematic and reliable testing of RFID devices, there is less effort made on the analysis and visualization of the data. A commercial product, ep-

- Yingcai Wu, Ka-Kei Chung, Huamin Qu, and S.C. Cheung are with the Department of Computer Science and Engineering at the Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong. Email: {wuyc|kkchung|huamin|scc}@cse.ust.hk.
- Xiaoru Yuan is with the Key Laboratory of Machine Perception (Ministry of Education) and School of EECS, Peking University, Beijing, P.R. China. Email: xiaoru.yuan@pku.edu.cn.

Manuscript received 31 March 2009; accepted 27 July 2009; posted online 11 October 2009; mailed on 5 October 2009.

For information on obtaining reprints of this article, please send email to: tvcg@computer.org .




Fig. 1. (a): α Gate: an RFID measuring, testing, and calibration instrument; (b): RFID volumetric data sets; (c): An RFID system deployment in a typical warehouse scenario.

cHotSpot [9], provides a simple visualization solution for displaying the tag readability by mapping out the real-world RF-field around a product. However, epcHotSpot is used mainly by ordinary users to determine where to place the tags; thus it cannot help in analyzing a large amount of benchmark data.

History mechanisms record history information of data exploration processes and allow users to backtrack to any previous step. A novel user interface called image graphs [17] was developed to visually represent the data exploration process. Jankun-Kelly and Ma [10] presented a spreadsheet-like interface for visualization exploration and encapsulation. A volume visualization interface based on PCPs was introduced to provide users with an intuitive history mechanism [20]. Jankun-Kelly et al. [11] used a P-Set model to encapsulate, share, and evaluate the visualization process. Kreuzler et al. [16] introduced a tree-style history mechanism for visual data mining. Our system also employs a history tree for exploring complex RFID benchmark data, but it is adapted to work closely with PCPs and orientation plots.

Visual analytics has emerged as a new research field and a hot research topic [13, 22] in recent years. Visual analytics can help users gain insights into massive, heterogeneous, and dynamic volumes of information by incorporating human judgment into the analytical reasoning process with interactive visual interfaces [13]. For the advantages and major research challenges of visual analytics, interested readers can refer to an excellent book [19]. Nevertheless, to the best of our knowledge, no comprehensive visualization system for RFID benchmark analysis exists.

Parallel Coordinate Plots (PCPs) are widely used visualization tools for visualizing multidimensional data [8]. Due to PCPs' usefulness and prevalence, many improvements such as user interaction schemes [21], clustering and filtering approaches [5, 23], and visual enhancement methods [15] have been made to facilitate multivariate data exploration. PCPs are often used together with other visualization techniques such as volume rendering to explore complex data sets [3]. For instance, Akiba and Ma [2] developed a visualization interface to visually link views of PCPs, time histograms, and volume rendering to simultaneously explore the data in temporal, spatial, and variable domains. Our work makes use of linked scientific and information visualization views to help users visually analyze RFID benchmark data. However, the PCPs used in our system are adapted to encode spatial information intuitively in order to aid in the detection of spatial patterns; this differentiates our method from other 3D PCPs [6, 7].

3 RFID BENCHMARKING

In this section, the RFID benchmark system, the acquired benchmark data, and the existing data analysis approaches are introduced.

3.1 RFID Benchmark System

A successful RFID system deployment (or a good RFID system design) usually means that all tags can be covered by the antennae with a high read rate and without missing read in the working environment. However, it is challenging to design such a system because the performance of RFID devices can be constrained by a variety of environmental, material, and device factors. To address the technical challenges, a novel RFID benchmarking instrument called α Gate has been developed. The instrument, well received by the RFID industry, won an

industrial award¹ in 2008.

In practice, three to four reader antennae (or simply antennae) are usually present. This instrument, however, has three antennae on three fixed axes (one horizontal and two vertical), respectively. A number of rectangular paper boxes containing merchandise are placed on the α Gate's rotatable base to constitute a volume based on a predefined size $n_x \times n_y \times n_z$. Tags are attached to the surfaces of the boxes (see Fig. 1(a)). The instrument can aid in the identification of the best position, orientation and proximity among tags and antennae to ensure the optimum read-rate of a tag population in their application field. In addition, it can be used to identify the best mix and match among different tags and RFID antennae for specific applications. It also facilitates troubleshooting and abnormality identification or isolation during the development and deployment of RFID devices/applications.

The instrument has found wide applications in warehouse and stock flow management. Figure 1(c) gives a typical warehouse scenario where α Gate can be deployed. RFID-tagged pallets on the right hand side of Fig. 1(c) are to be delivered to different destinations by trucks parked at the entrance of each loading bay, which is equipped with four RFID antennae highlighted in red rectangles. Through these antennae the warehouse management system monitors if each pallet is loaded onto the right truck. Missing detection of RFID tags create expensive overheads in order to manually track down missing stock items. A challenge to reliable tag detection is the multipath radio waves reflection from the building structures and the metal grounding plate [14]. As shown in Fig. 1(c), four engineers are manually configuring the position of an RFID antenna in order to improve tag detection, which is time consuming and tedious. Our proposed system can ease the task through effective visualization of each configuration's performance.

3.2 Data Acquisition

The α Gate system usually takes around 13 hours to complete an RFID benchmark. In benchmarking, the α Gate's antennae are moving and rotating along each axis in turn, and at each step only one antenna is activated and can receive radio frequency signals from the testing volume of tags and determine the strength of the signals. Each RFID benchmark contains thousands of steps and can generate a large amount of data which possess some special features. The acquired data are volumetric data sets at each step (see Fig. 1(b)). The recorded RFID data are intrinsically time-varying volumetric data. In addition, the data are typically multidimensional because each voxel of the volume stores the following: (1) an activated antenna ID (top, left, or right); (2) a received signal strength indication (RSSI) value; and (3) a vector of the activated antenna's spatial configuration ($AntPos, \varphi, \theta$), where $AntPos$ is the antenna position on its axis and φ and θ indicate its orientation.

3.3 Data Analysis

After the benchmarking, researchers and engineers can retrieve the data to perform various data analysis tasks.

3.3.1 Finding the Optimum Antenna Spatial Configuration

Finding the optimum spatial configuration entails identifying a minimum number of antenna spatial settings tested in the benchmarking in which a strong signal is received from all tags. The optimum configuration can facilitate effective RFID deployment because fewer antennae will be used while the satisfactory performance is delivered. The problem is the classic set cover problem, an NP-complete problem which can be approximately solved by greedy algorithms. Suppose E is a set that contains all tags, and the tags that can be covered by an antenna spatial setting recorded at time frame t is treated as a subset S_i of E . Thus, the problem is to find a minimum collection C of sets from $S = \{S_1, S_2, \dots, S_n\}$, such that C covers all the tags in E .

Algorithm 1 is the greedy algorithm currently used by the RFID engineers. Its objective function is simply defined as the number of remaining tags that S_i can cover. Although the algorithm could run efficiently, it often fails to find the optimum solution. The usefulness

¹<http://www.cse.ust.hk/News/RFIDAward2008/>

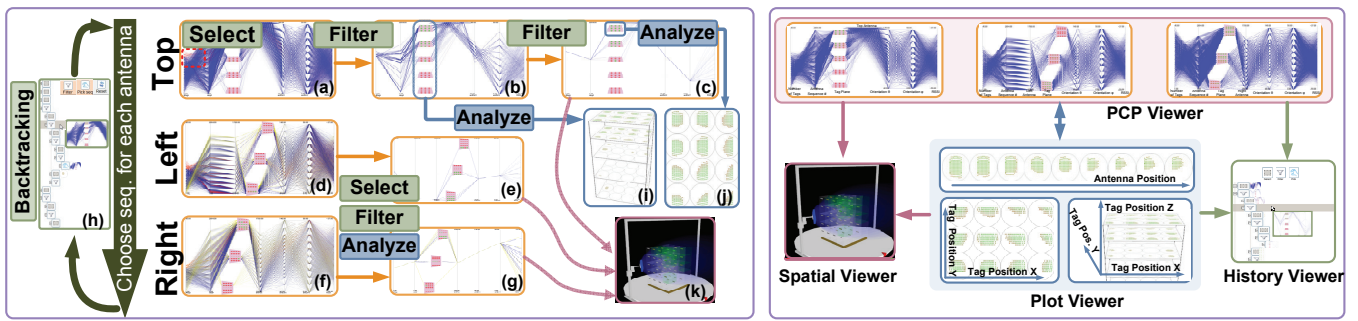


Fig. 2. Left: Visual optimization for RFID benchmarking; Right: Visual analysis for RFID benchmark data comparison and abnormality detection.

of the algorithms is also constrained by low flexibility. Although flexibility can be enhanced more or less by other appropriate objective functions, defining the functions is a demanding task. An objective function for one brand of tags is often inappropriate for another brand. In different applications, even tags of the same brand may behave quite differently and the required objective function would also be very different. Furthermore, transforming the expert knowledge of RFID deployment for a brand of tags in a typical application to an objective function is quite difficult. Due to the multipath effect [14], the spatial coverage patterns of tags are very irregular and also highly environment dependent. Currently, the irregular spatial patterns cannot be described mathematically (e.g., described as an objective function) in the RFID field.

Algorithm 1: Greedy Algorithm

```

1 Let  $C \leftarrow \phi$ 
2 while  $E$  contains elements not covered by  $C$  do
3   Find  $S_i$  containing the greatest number of uncovered tags
4   Add  $S_i$  to  $C$ 
5 end
6 Output the antennae settings associated with  $S_i$  in  $C$ 

```

Another solution is to use 3D graphics techniques (see Fig. 3) to examine the data frame by frame and find the optimum configuration manually. This solution is time consuming and may not find a good configuration, especially when the search space is huge. Thus, we introduce a visual optimization solution with a proper integration of human knowledge and visualization techniques (see Fig. 2, left). In our solution, adapted PCPs (see Fig. 2(a)-(g)) are provided to better reveal the correlations between different variables of the data. Users can interactively select an antenna setting from PCPs for each of the three antennae (the top, left, and right antennae) and visually approach the optimum configuration. A history viewer enables users to backtrack to any previous step readily (see Fig. 2(h)). Users can also inspect how the signal strength of a tag is related to the antenna orientations intuitively by orientation plots (see Fig. 2 (i) and (j)), facilitating the selection of the antenna orientation on the PCPs. A 3D spatial viewer (see Fig. 2 (k)) will highlight the antenna and the covered tags when users select an antenna setting from the PCPs.

3.3.2 Comparing the Performance and Detecting the Flaws

Comparing performance of different RFID devices in different environments can facilitate identification of their best combination, enabling the most cost-effective RFID deployment. Detecting abnormality of RFID devices, on the other hand, can facilitate the development of new RFID products, and ensure the reliability of RFID systems. Normally, signal strength received from the tags decreases with the tag-antenna distance and the number of boxes along the tag-antenna path. However it can also be affected by the environment, material and devices involved, resulting in an abnormality (see Section 7.1.3 for an example). An abnormality exists if, for instance, tags closest to the antenna have very low *RSSI* or tags further from the antenna have higher *RSSI* than those nearer to the antenna (see Fig. 7(b)).

RFID engineers currently employ a primitive 3D graphics technique (see Fig. 3) to view the benchmark data sets frame by frame for performance comparison and abnormality detection. However, this method has several major drawbacks. Owing to the occlusion problem, it would be difficult for users to see the tags clearly. In addition, comparing two 3D images frame by frame takes a long time as there are usually over 10,000 frames in a benchmark data set. Moreover, this method can only deal with the scalar values, such as signal strength, and cannot handle other types of recorded information such as 3D antenna orientations. Therefore, we developed a set of highly interactive visualization techniques (see Section 4) and integrated them into a visual analysis workflow with four linked views, namely, a PCP viewer, a plot viewer, a spatial viewer, and a history viewer, thus allowing the users to visually and intuitively analyze and explore sophisticated data from different aspects (see Fig. 2, right). The PCP viewer with three synchronized PCPs can reveal the overall patterns of the data. The plot viewer based on orientation plots allows for a detailed analysis of the data selected from the PCP viewer in order to find interesting patterns associated with orientations and signal strength. The spatial viewer models α Gate and visually presents the data selected from the PCP viewer or the plot viewer using 3D graphics techniques. Users can also use the history viewer to keep track of the analysis process.

4 VISUALIZATION TECHNIQUES

In this section, we introduce four visualization techniques: a 3D spatial viewer, a visual history mechanism, PCPs, and orientation plots.

4.1 3D Spatial Viewer

The 3D spatial viewer is used to visualize RFID benchmark data in 3D spatial space as shown in Fig. 3. It models the real benchmarking instrument using OpenGL to provide necessary and informative spatial context to users while using other visualization techniques. A pile of semi-transparent boxes are drawn on the instrument to form a big bounding box, and the attached tags are represented as rectangles on the surfaces of the smaller boxes. The color of each rectangle encodes the signal intensity (*RSSI*) value of the corresponding tag. In the system, a red-to-green lookup table is employed to encode the *RSSI* values of the tags. In addition, the system uses white to represent the tags that cannot be covered by the activated antenna. This can help users better differentiate them from the covered ones. The benchmark data are rendered in a frame-by-frame manner and the activated antenna is also drawn. The viewer can support the following operations:

- **Basic 3D Operations:** The viewer supports 3D zoom-in/out, rotation, and panning operations.
- **Information Highlight:** Important information such as the covered tags can be highlighted (see the green boxes in Fig. 3).
- **Animation:** The viewer can automatically generate an animation, in which the antenna moves and rotates on its axis and the tag color also changes based on its *RSSI*.

4.2 Tree-style History Viewer

Exploring highly complex and multidimensional data by visualization is usually an iterative and cyclic process. In the process, users fre-




Fig. 3. 3D spatial viewers: The covering range of the antenna is highlighted in (a), and not highlighted in (b).

quently modify system parameters and switch between different views to explore data from different perspectives. Thus, our system employs a tree-style history viewer similar to [16] for visualizing the history (provenance) information to facilitate the iterative and cyclic exploration process (see Fig. 7(a)). The system automatically collects intermediate results and provenance information in the process. The viewer maintains a tree data structure for the collected provenance information. The tree nodes denote user interactions, while the hierarchical relation between the tree nodes refers to the sequence order of the operations, as shown in the top right corner of Fig. 7(a). The history viewer supports various user interactions as follows:

- **Delete:** Users can delete some branches of the tree.
- **Save/Load:** The tree can be saved or loaded such that users can save their current work and restart the work later.
- **Expand/Collapse:** Just like Windows Explorer's tree view, the viewer supports the expand/collapse operations on each node.
- **Redo/Undo:** The "redo" and "undo" operations are supported.
- **Backtracking:** Users can return to a previous step by performing the backtracking operation on the corresponding tree node.

4.3 Parallel Coordinate Plots

In our system, we employ PCPs to explore RFID benchmark data in multidimensional data space, as PCPs are among the most effective visualization tools for users to visually analyze multidimensional data [8]. PCPs use parallel axes to represent dimensions in 2D image space and the dimensions used in the system include:

- Sequence ID of each antenna spatial configuration,
- Antenna's position on an axis of the instrument (*AntPos*),
- Antenna's orientation represented by two angles (φ , θ), and
- Tag's 3D positions (x, y, z) and its RSSI value (*RSSI*).

Data values are visually presented by polylines linking the axes.

4.3.1 Representation of Spatial Information

The PCPs used in our system are adapted to intuitively represent spatial information because pieces of information such as tag and antenna positions are especially important for analyzing the data. Without proper spatial information, the underlying spatial patterns would not be revealed properly. Most PCPs employed for analyzing spatial non-abstract data just add more parallel axes to represent 3D spatial information, as shown in Fig. 6(c) where Tag *X*, Tag *Y*, and Tag *Z* are drawn to encode the spatial information of tags. However, this approach does not work well in visualizing data that have strong spatial patterns. For example, the lines connecting Tag *X*, Tag *Y*, and Tag *Z* are too regular to reveal the spatial patterns in Fig. 6(c). Some PCPs are also coordinated with other views to provide spatial information. While this method helps users perform some filtering operations based on the provided spatial views, it usually cannot facilitate the revealing of the underlying spatial patterns.

To address this issue, we extended the PCPs for RFID data to preserve the spatial context (see Fig. 6(a) and 6(b)) such that spatial patterns could be readily detected. In addition, with the extended PCPs,

users do not need to mentally link the PCPs to other spatial views to obtain the spatial context, thus relieving users of the burden and increasing the efficiency of the visualization. Two new styles of PCPs are developed to facilitate detection of the underlying spatial patterns based on the types of antennae, inasmuch as the patterns in RFID data are mainly influenced by the moving patterns of antennae moving horizontally (the top antenna) or vertically (the left and right antennae).

- **PCPs with horizontal planes:** The PCPs have parallel horizontal planes to represent spatial information of RFID tags for the RFID data acquired by the top antenna (see Fig. 6(a)). Furthermore, the axis indicating the top antenna position is placed horizontally to model how the antenna moves in 3D space.
- **PCPs with vertical planes:** Parallel vertical planes are used to encode spatial information of RFID tags for the RFID data acquired by the left or right antenna (see Fig. 6(b)).

Each tag position originally represented by a polyline linking the three axes, e.g., Tag *X*, Tag *Y*, and Tag *Z* in Fig. 6(c), is now represented by a circular node on a plane according to its 3D spatial position. The node color encodes the number of polylines connected to the tag. When an antenna spatial setting (*AntPos*, φ , θ) is selected, each node color will encode the tag's *RSSI* value under the antenna configuration. To reduce visual clutter, each polyline in the PCPs does not directly connect to its associated tag on the plane, but instead connects to the edges of the plane in terms of the tag's position. The polylines connect to the same position on the edge of a plane if their associated tags are in the same row on the plane.

4.3.2 Interactive Operations

PCPs support highly interactive operations in the following ways:

- **Brushing:** Users are allowed to perform the brushing operation on PCPs to select a subset of the data.
- **Axis Switching:** Users can configure the PCPs to determine the order of the axis based on their requirements.
- **Linking to Other Views:** The PCPs can be coordinated automatically or manually with other views by brushing.
- **Drag & Drop:** Users can manually drag and drop the tag planes in the PCPs to reduce visual cluster. The position adjustment can be finished quickly as there are only several planes in PCPs.

4.4 Orientation Plots

Every reader antenna rotates systematically and has 200 different orientations in total at each position (*AntPos*) on its axis. Each orientation is determined by two angles (φ , θ) in the spherical coordinate system (see Fig. 4(a)). The yellow surface area on the sphere represents all possible antenna orientations. The orientation information is quite useful for data analysis tasks such as the flaw detection. However, orientations cannot be revealed properly by existing methods. An orientation plot is hereby proposed to visually and intuitively represent all orientation information, such that interesting patterns related to antenna orientations are more likely to be observed.

4.4.1 Representation of Orientation Plots

Given a tag and an antenna fixed at a position, an orientation plot can represent all antenna orientations and tag *RSSI* values simultaneously on a circular 2D plane. The plane is created by flattening the yellow surface area on the sphere (see Fig. 4(a)). Assume that the center of the 2D plane is located at the origin of the 2D Cartesian coordinate (*X-Y*) system. For an orientation (φ , θ), its position in the plane is determined as follows: $x = \rho \sin \varphi \cos \theta$ and $y = \rho \cos \varphi$, where ρ is the radius of the sphere. The color at the position encodes the associated *RSSI* value so that its relation with the orientation is clearly displayed. The flattening process is illustrated in Fig. 4(a). The orientation plots, called **absolute orientation plots**, are helpful in some scenarios of visual analysis, e.g., in the performance comparison of two tags at the same tag position (see Fig. 5(b)). However, the plots cannot help users determine whether the *RSSI* distribution is related to the relative angles between an antenna and a tag or not. Thus, we have developed another type of orientation plots called **relative orientation plots** to address the concern. Compared with absolute plots,

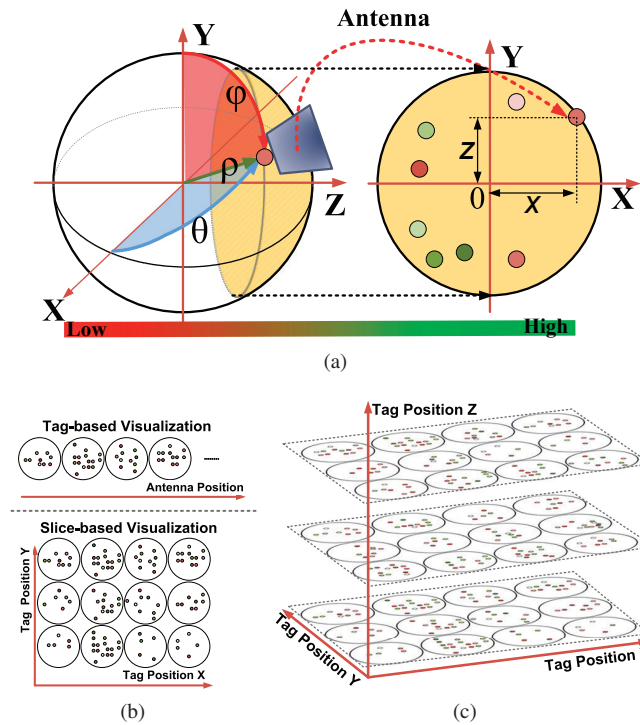


Fig. 4. Orientation plots: (a) Encoding scheme; (b) Top: tag-based and bottom: slice-based visualization of orientation plots; (c) Slice-based visualization of orientation plots in 3D spatial space.

relative plots use the pixel positions to encode relative angles rather than absolute angles. Assume that the direction from an antenna to a tag is denoted as (φ_0, θ_0) in the spherical coordinate system. For an orientation (φ, θ) , the relative orientation can then be calculated as $(\varphi' = \varphi - \varphi_0, \theta' = \theta - \theta_0)$, and its position in the 2D plane is determined as follows: $x = \rho \sin \varphi' \cos \theta'$ and $y = \rho \cos \varphi'$. Figure 5(c) is an example where the relative plots are useful in comparing two tags at different tag positions. A comparison of absolute and relative plots is shown in Fig. 5(a). In the absolute plot the best signal (in green) is found near the edge. In contrast, in the relative plots, the best signal is near the center because the antenna is facing the tag at that position.

4.4.2 Visualization of Orientation Plots

In analyzing the data, users often need to compare the performance of different tags when the activated antenna is fixed, or to compare the performance of the same tag when the antenna is moving on the axis. Some interesting patterns would be revealed only when multiple orientation plots are displayed side by side. Hence, showing multiple orientation plots simultaneously for data comparison and pattern identification is desirable. Our system provides two different methods, **tag-based** and **slice-based** methods, to allow users to compare tag performance efficiently and identify interesting patterns intuitively.

- **Tag-based Visualization** This method arranges all plots of a certain tag side by side (see Fig. 4(b), top) along the moving direction of an antenna on its axis. This can help users analyze how the tag's RSSI distribution changes when the antenna is moving along the axis in order to find interesting patterns.
- **Slice-based Visualization** This method simultaneously visualizes the orientation plots of the whole tag volume (see Fig. 4(c)) or a slice of the volume (see Fig. 4(b), bottom) when the activated antenna is fixed at a certain position. Compared with the tag-based method, this method shows the orientation plots of multiple tags in which the antenna position is fixed.

4.4.3 Interactive Operations

The orientation plots support the following interactive operations:




Fig. 5. (a) Absolute and relative plots for the same tag; (b) Comparison of two tags at the same tag position by absolute plots; (c) Comparison of two tags at different tag positions by relative plots.

- **Slicing:** Users can perform slicing in a top-down, left-right, or front-back manner to change the slice of the tags for a slice-based visualization of the plots.
- **Brushing:** Users can choose a set of pixels on a plot by brushing; the corresponding data records will then be highlighted in other viewers (PCPs or the spatial viewer).

5 VISUAL OPTIMIZATION FOR RFID BENCHMARKING

In this section, we introduce a new visual optimization procedure for users to interactively estimate the optimum spatial configuration using the proposed visualization techniques. While the approach is highly flexible and intuitive and can take advantage of the users' experience, background knowledge, and intuition, its usefulness would be constrained by its intrinsic low efficiency and aimlessness, especially for visualizing large benchmark data. The greedy algorithm, on the other hand, always makes the locally optimal choice that appears best at the moment and is often used to estimate NP-complete problems efficiently. However, it has major drawbacks such as low flexibility in meeting different user requirements as described in Section 3.3.1. Thus, the greedy algorithm is the perfect complement to the visualization techniques. Based on the observation, we integrate the visualization techniques and greedy algorithm into a coherent visual optimization procedure leading to an efficient, flexible, and intuitive solution for estimating the optimum spatial configuration. The visual optimization procedure is outlined as follows:

1. Start from an antenna (top, left, or right) and an empty set C .
2. Display the PCP of the data for the selected antenna where all polylines connecting to the tags in C have been removed.
3. Select one antenna spatial setting ($AntPos, \varphi, \theta$) from the PCP that can cover most uncovered tags (i.e., in a greedy manner).
4. Add the tags covered by the selected antenna to C .
5. Choose another antenna and repeat from Step 2 to 4 until all three antennae have been selected.
6. Return to a previous step via the history viewer if there are still some uncovered tags.

This procedure allows users to make decisions based on their analysis of the information displayed. Users can get instant feedback of what their operations lead to. The antennae operate in a time-varying multiplex mode in which only one antenna is activated at anytime in the benchmarking, thus it is reasonable to select the antennae one by one. In the above procedure, users mainly manipulate the PCPs for visual optimization to find the optimum configuration. Thus, we further extend the PCPs to facilitate visual optimization as follows:

- We add one more axis to the left side of the PCPs to represent the number of tags covered by an antenna setting and help users select an appropriate antenna. By brushing on the axis, users can filter out most of the polylines representing the antenna settings that cover fewer tags in the PCPs, reducing the visual clutter.
- We retain and highlight the polylines that indicate the relatively weak tags (i.e., with low $RSSI$ values) covered by the previously selected antennae rather than filter them out in Step 2 in the visual optimization process. Users can select a new antenna that can cover the weak tags with strong signals, ensuring that the selected antennae can complement one another.

The whole workflow procedure is shown in Fig. 2. Following the basic greedy principle, users select antenna settings that can cover more tags (highlighted by the red polylines in Fig. 2(a)) by brush-

ing on the leftmost axis to narrow down the search space. The users continue to search for a good antenna setting by taking the number of covered tags, tag positions, antenna orientations, and *RSSI* values into account by brushing on different axes to get instant feedback and to finally obtain a potentially good setting (see Fig. 2(c), 2(e), and 2(g)). The spatial viewer is then automatically updated to show the effect of the antenna setting in 3D space (see Fig. 2(k)). The weakly covered tags are highlighted in PCPs (see Fig. 2(d) and 2(f)) to help users find an antenna configuration in which the antennae can complement one another. If users are unsatisfied with the result, they can use the history viewer to backtrack to any previous step easily (see Fig. 2(h)). Orientation plots are also provided to facilitate the search, especially when users fail to find a good spatial configuration after several iterations. The plots can help users figure out, for example, which antenna orientations perform poorly in covering a tag. Users can select a tag, a slice of tags, or even multiples of slices of tags directly from the PCPs. The visualizations of the plots can be provided to the users instantly (see Fig. 2(i) and 2(j)).

6 VISUAL ANALYSIS FOR DATA COMPARISON AND ABNORMALITY DETECTION

In this section, we introduce a visual analysis procedure for comparing performance and detecting possible flaws of RFID devices. With the procedure, users can intuitively compare and examine the data from different perspectives at different levels of detail in a highly interactive manner. The general workflow procedure, which includes four well coordinated views, is shown in Fig. 2.

- **PCP Viewer** The PCP viewer has three synchronized PCPs placed side by side for the top, left, and right antennae, respectively. The viewer presents the benchmark data in an informative manner and provides an overview of the data to users; thus, it can be used as the starting point for the visual analysis. Users can interactively explore the data by brushing PCPs to determine which parts of the data should be examined in detail.
- **Plot Viewer** The plot viewer is based on orientation plots. It allows users to visually analyze the data selected by the PCP viewer thoroughly in order to detect interesting patterns associated with the orientations and signal strength. If the users finish the analysis for the selected data with the viewer, they can return to the PCP viewer to select other data for further analysis.
- **Spatial Viewer** The spatial viewer presents the benchmark data selected from the PCP viewer or the plot viewer to users with the 3D graphics model. It facilitates the detailed analysis of the benchmark data in 3D space in a frame-by-frame manner and convey some found patterns to users intuitively.
- **History Viewer** The history viewer displays the historical information including previous operations and the rendered images from the PCP viewer and the plot viewer hierarchically. It is primarily used for detecting flaws and abnormalities in the cyclic exploratory process to keep users informed on what had already been explored and hence avoid redundant operations.

The above workflow works particularly well for pattern detection. To compare the performance of RFID tags, however, it needs a few minor adjustments. Given two benchmark data sets for comparison, each viewer will have two side-by-side views that are synchronized through the brushing and filtering operations for the two data sets, such that users can visually compare the performance quickly.

7 EXPERIMENTS AND DISCUSSION

The system was developed in C++ based on the Visualization Toolkit (VTK). We tested the system on a Dell XPS 730 PC equipped with Intel Core Quad CPU Q9550, 4 GB DDR3 memory, and dual NVIDIA GeForce 9800GT (SLI) with 1 GB memory. Interactive performance is achieved after the data are loaded into the memory in all experiments.

7.1 Experiments

To validate the usability and usefulness of our visualization solutions, we conducted an informal task performance test and invited three RFID engineers to participate in the test. They all have considerable

experience in analyzing RFID benchmark data using the 3D spatial viewer and greedy algorithm. In the test, the engineers were asked to perform three tasks: estimating the optimum antenna configuration, comparing tag performance, and detecting abnormalities using different approaches. Three RFID benchmark data sets of the tags produced by different companies were used in the test. All data sets are time-varying volumetric data sets each of which has 8415 time steps in total with the same volume dimension ($3 \times 4 \times 5$). Prior to the test, a 10-minute introduction to our system was made and a 10-minute dry-run trial was provided for each engineer.

7.1.1 Estimating the Optimum Antenna Configuration

In the first task, we asked the participants to estimate the optimum configuration of the antennae for the three data sets. Each participant used three methods, namely, the 3D spatial viewer, the greedy algorithm, and our visual optimization solution for the three data sets in random order, respectively. Every participant applied each method to a different data set. The one-to-one correspondence of the methods and the data sets was randomized to reduce possible bias. The experiment was carried out to demonstrate the intuitiveness, efficiency, and flexibility of the visual optimization method.

In the experiment, the greedy algorithm failed to find the optimum configuration for two out of the three data sets. For instance, for one failed case, the algorithm found a top antenna setting for covering 26 tags (60 tags in total), and a left antenna setting for covering 13 tags out of the remaining 34 tags. For the rest, at most 18 tags can be covered by any right antenna setting, leading to the missing read of three tags. The participants had to count on their own experience in manually searching for the desired results in the failed cases. The participants spent more than 15 minutes on the manual search and finally gave up.

Three synchronized spatial views, for the top, left, and right antennae respectively, were provided by the spatial viewer for the task. With the spatial views, two participants finally found a spatial configuration of antennae that can cover all tags in 15 minutes on average, while one participant failed. The participants confirmed that the major difficulty in using the spatial viewer for this task lies in the frame-by-frame viewing mechanism, leading to a time-consuming and tedious optimization process. They also complained about the severe occlusion for obtaining undesired configurations that merely cover partial tags.

All participants successfully completed the tasks with the visual optimization method; the time needed was three minutes on average. They all agreed that the method is much more efficient and flexible than the other methods. The efficiency was achieved by integrating the greedy principle and presenting all information into a coherent and informative visualization rather than by the frame-by-frame visualization, while the increased flexibility was obtained through the highly interactive user interface with multiple linked views. For instance, for the failed case of the greedy algorithm we mentioned earlier, a participant first followed the greedy principle using the PCPs. After selecting the spatial settings of the top, left, and right antennae, respectively, with our method, it turned out that merely selecting the most remaining tags that could be covered by the current antenna in each iteration could not deliver a good result. The participant then used PCPs and the spatial viewer to check whether there were any spatial patterns of the uncovered tags and found that the tags can be mostly covered by some antenna settings. Thus, he backtracked to a previous step using the history viewer and selected the antenna setting for covering those tags. The process was repeated until he found a satisfactory solution. During the process, the participant frequently employed the orientation plots to visually examine the uncovered tags so that he could keep track of the overall antenna orientation patterns to facilitate the visual optimization process. With our visualization method, an appropriate antenna spatial configuration was obtained in minutes.

They also felt that the visual optimization is intuitive and easy to use. Two participants commented that the history tree was very helpful in the cyclic exploratory process (see Fig. 7(a)) because they could easily backtrack to a previous step through the viewer without remembering the parameter settings. They confirmed that the spatial informa-

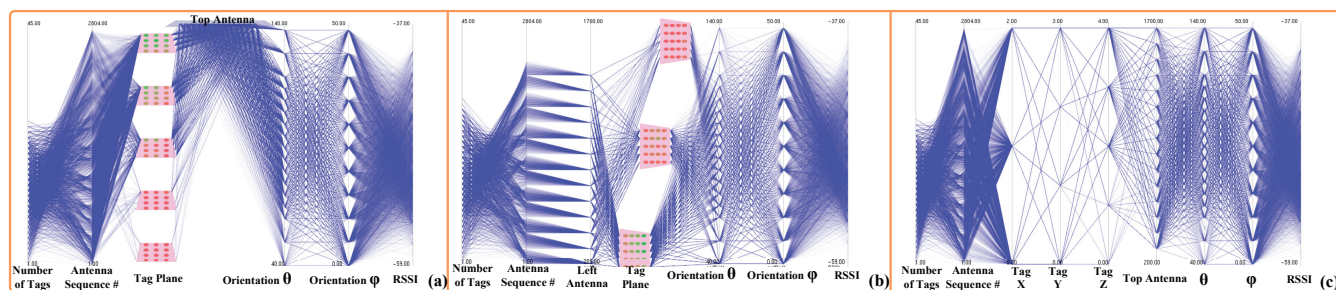


Fig. 6. (a): A PCP for the top antenna; (b): A PCP for the left antennae, and a similar PCP can also be constructed for the right antenna. (c): A traditional PCP for the top antenna.

tion encoded in our new PCPs (see Fig. 6(a) and 6(b)) was much more intuitive to perceive than that of the traditional PCPs (see Fig. 6(c)). In Fig. 6(a), we can readily and intuitively observe some spatial patterns. For example, some tag positions in the bottom slice cannot be covered by the antenna. However, in the traditional PCPs, the spatial patterns could not be easily perceived (see Fig. 6(c)). Two participants frequently used orientation plots when they found that several tags could not be well covered by their selected configuration. Based on the plots, they analyzed the *RSSI* distributions among different antenna orientations and finally identified a good antenna configuration.

7.1.2 Comparing Performance

In the second task, we asked the participants to compare the performance of different tags. Each pair of the data sets was compared by two participants, each of whom used a different method – the 3D spatial viewer or our visual analysis solution. In addition, the experiment was carefully designed to ensure that each participant tried both methods in random order. This experiment was conducted to show the limitations of the existing method and to validate the intuitiveness and effectiveness of our approach.

Two spatial views were provided by the spatial viewer for the participants to compare a pair of the benchmark data sets side by side. The average time with the spatial viewers was around 20 minutes for each comparison. They all felt that it took them a long time to complete the task. To compare the tag performance effectively, they had to perform the rotate and zoom in/out operations frequently to view each tag performance clearly and reduce occlusion. In addition, they felt that it was quite inconvenient for them to perform the comparison frame by frame. They wanted a simultaneous comparison of all data rather than a step-by-step comparison. Moreover, they commented that when they wanted to compare the tag performance under different antenna orientations, the approach would not work.

The average time needed to perform the comparison with our visual analysis solution was about 15 minutes. The participants all agreed that the PCP viewer is a good starting point to gain a general spatial pattern of the *RSSI* distribution by interactive brushing operations, such that they could choose an interesting subset of the data for further analysis. The orientation plots allow users to further analyze data selected from the PCPs. Both absolute and relative plots were received quite well by all participants. The absolute plots were used to compare the performance of tags at the same position, while the relative plots were often used to compare the performance of tags at different positions because the tag positions were also considered in the plots (see Fig. 5(a)). The tag-based and slice-based visualizations were both considered very intuitive and were used frequently by all participants for the comparison. The tag-based method (see Fig. 5(b) and 5(c)) helped the participants focus on comparing the performance of tags under all antenna configurations at the same time, while the slice-based method helped them simultaneously compare the performance of tags on one or multiple slices under one antenna configuration (see Fig. 7(b)).

After all the participants have finished the task, for a pair of data sets, we found that two participants made conflicting conclusions on

the corresponding tag performance. After investigating the two data sets and discussing with one another, the participants agreed that the conclusion made by the visual optimization was correct while the other one made based on the spatial viewer was wrong.

7.1.3 Detecting Abnormalities

In the third task, the participants were asked to freely explore the data sets using either the spatial viewer or our comprehensive visual analysis method to find abnormalities in 20 minutes. The one-to-one correspondence of the methods and the data sets was randomized to reduce possible bias. This task was conducted to demonstrate the advantages of our visual analysis method in detecting unexpected patterns.

No abnormality was observed in all data sets by the participants using only the spatial viewer. Two participants suggested that 3D occlusion was a major obstacle preventing them from detecting any unexpected pattern, while one participant argued that the tedious step-by-step display mechanism was the main reason. They all agreed that 3D orientations of the antennae were difficult to perceive in the viewer (see Fig. 3), such that the patterns related to the orientations would be hardly detected. On the other hand, with our visual analysis method, two out of three participants were able to detect several interesting unexpected patterns in the data sets never discovered before. One of the patterns is shown in Fig. 7(b) where the 3D slice-based visualization is used to display the orientation plots of all tags when the top antenna was fixed at a position. Generally speaking, the top antenna should receive no signal from the tags at the bottom slice. However, in the case of Fig. 7(b), the antenna received strong signals from a tag at the bottom slice. These strange patterns confused the participants. After investigating the data and inquiring with the α Gate operator, the RFID engineers concluded that the strange patterns were possibly caused by natural environmental noise such as reflectance, refraction, diffraction, and absorption of radio frequency signals coming from the neighboring walls or even the human body. The exact reason is still under intensive study.

All participants agreed that the visual solution is much more robust than the spatial viewer for the task. The supported levels-of-detail visualization methodology was well received by the participants. They could start from the PCPs to select some suspected patterns by brushing and then analyze the selected data using orientation plots and the spatial viewer, thus enabling efficient pattern detection. Two participants felt that the history viewer was extremely helpful in pattern detection, as the viewer kept them informed of what they had already explored to reduce exploration redundancy, while one participant considered it less useful because he could remember the exploration process by himself quite well. They especially liked the 3D slice-based visualization method (see Fig. 7(b)) because this method allowed them to visualize the orientation plots for all tags when the activated antenna is fixed at a certain position. The orientation plots improved their understanding of how the antenna orientation influences the signal strength, allowing them to find the orientation patterns quickly (see Fig. 7(b)).

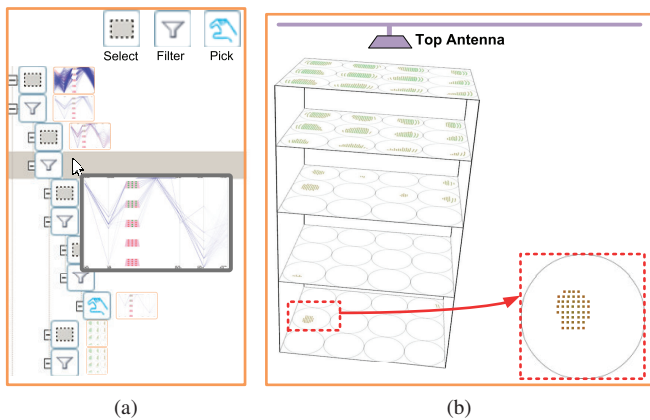


Fig. 7. (a) The history viewer; (b) The 3D slice-based visualization of plots where a detected unexpected pattern is shown.

7.2 Limitations and Discussion

The experiments confirm the applicability of our visualization techniques, inasmuch as the experiments were conducted in real applications. Through the task performance evaluation test, some limitations of our visualization techniques were also identified. Users may be still overwhelmed by visual clutter in the PCPs, preventing them from seeing the overall patterns. This problem can be greatly alleviated with the provided interactive operations such as the brushing and drag-and-drop operations. In the future, we plan to employ some advanced visual clutter reduction techniques to further reduce visual clutter in PCPs. The visual optimization method works well in our experiments for all the tested data. However, users may fail to find a good antenna configuration, especially when there is no antenna configuration in the data that can cover all tags. For that case, the engineers suggested that the uncovered tags should be localized, such that adding one more antenna for the tags can deliver a good configuration. Thus, we want to improve our visual optimization method to provide visual cues to users for finding the localized uncovered tags. A formal user study will be conducted and more engineers will be invited to participate in the study. Furthermore, in the future, we plan to conduct more sophisticated benchmarking tasks using our system.

8 CONCLUSIONS

In this paper, a set of visualization techniques with coordinated views were introduced for visualizing special RFID benchmark data. Based on the proposed techniques, we developed two solutions, visual optimization and visual analysis, to the real problems in RFID benchmarking. Our visualization solutions have the following advantages:

- ◆ **Robust:** They are much more robust than other methods because the multiple coordinated views enable users to visualize the data from different perspectives at different levels of detail.
- ◆ **Intuitive:** Through highly interactive visualization techniques, the solutions are much more intuitive than other methods.
- ◆ **Efficient:** Compared with the 3D graphics methods, the visual optimization solution achieves greater efficiency by integrating the greedy principle into the visual exploratory process. It will also be more efficient than the greedy algorithm when the greedy algorithm fails to find the optimum result and a great deal of effort is needed for manual search for the result.
- ◆ **Flexible:** The visual optimization allows maximum flexibility by the proper integration of human background knowledge, intuition, and experience in contrast to the greedy algorithm requiring different objective functions that are challenging for users to define.

The visualization solutions for RFID benchmarking will have great value and a broad impact considering the wide applications and bright future of RFID technologies. The success of our system will definitely facilitate the effective deployment of RFID technologies and promote the fast development of the RFID industry.

ACKNOWLEDGMENTS

This work was supported by HK RGC CERG 618706, HK RGC RPC07/08.EG24, SSRI08/09.EG01, ITP/022/08LP, Beijing Municipal NSF (No. 4092021), and RFDP 200800011004. We thank Wing-Yi Chan, Wai-Kei Tsang, Ping-Luen Ng, and Lap-Hin Leung for developing the initial 3D spatial viewer and help with the user evaluation. We thank the anonymous reviewers for their valuable comments.

REFERENCES

- [1] AberdeenGroup. The RFID benchmark report: Finding the technology's tipping point. <http://www.aberdeen.com>.
- [2] H. Akiba and K.-L. Ma. A tri-space visualization interface for analyzing time-varying multivariate volume data. In *EuroVis*, pages 115–122, 2007.
- [3] J. Blaas, C. P. Botha, and F. H. Post. Extensions of parallel coordinates for interactive exploration of large multi-timepoint data sets. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1436–1451, 2008.
- [4] S. Cheung, W. Chan, P. M. Lee, L. M. Ni, and P. Ng. A combinatorial methodology for RFID benchmarking. In *RFID Academic Convocation*, 2006.
- [5] G. Ellis and A. Dix. A taxonomy of clutter reduction for information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1216–1223, 2007.
- [6] E. Fanea, S. Carpendale, and T. Isenberg. An interactive 3D integration of parallel coordinates and star glyphs. In *IEEE Symposium on Information Visualization*, pages 149–156, 2005.
- [7] C. Forsell and J. Johansson. Task-based evaluation of multi-relational 3D and standard 2D parallel coordinates. In *Proceedings of SPIE-IS&T Electronic Imaging, SPIE Vol.6459, 64950C-1-12*, 2007.
- [8] A. Inselberg and B. Dimsdale. Parallel coordinates: a tool for visualizing multi-dimensional geometry. In *IEEE Visualization*, pages 361–378, 1990.
- [9] Integral RFID. epcHOTSPOT. <http://www.epchotspot.com/>.
- [10] T. Jankun-Kelly and K.-L. Ma. Visualization exploration and encapsulation via a spreadsheet-like interface. *IEEE Transactions on Visualization and Computer Graphics*, 7(3):275–287, 2001.
- [11] T. Jankun-Kelly, K.-L. Ma, and M. Gertz. A model and framework for visualization exploration. *IEEE Transactions on Visualization and Computer Graphics*, 13(2):357–369, 2007.
- [12] S. R. Jeffery, M. Garofalakis, and M. J. Franklin. Adaptive cleaning for RFID data streams. In *International Conference on Very Large Data Bases*, pages 163–174, 2006.
- [13] D. A. Keim, F. Mansmann, J. Schneidewind, J. Thomas, and H. Ziegler. Visual analytics: Scope and challenges. In *Visual Data Mining: Theory, Techniques and Tools for Visual Analytics*. Springer Berlin / Heidelberg, 2008.
- [14] D. Kim, M. A. Ingram, and W. W. Smith. Measurements of small-scale fading and path loss for long range rf tags. *IEEE Transactions on Antennas and Propagation*, 51(8):1740–1749, 2003.
- [15] R. Kosara, F. Bendix, and H. Hauser. Parallel sets: Interactive exploration and visual analysis of categorical data. *IEEE Transactions on Visualization and Computer Graphics*, 12(4):558–568, 2006.
- [16] M. Kreuseler, T. Nocke, and H. Schumann. A history mechanism for visual data mining. In *IEEE Symposium on Information Visualization*, pages 49–56, 2004.
- [17] K.-L. Ma. Image graphs - a novel interface for visual data exploration. In *IEEE Visualization*, pages 81–88, 1999.
- [18] ODIN Technologies. RFID benchmark reports. <http://www.odintechnologies.com/>.
- [19] J. J. Thomas and K. A. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE CS Press, 2005.
- [20] M. Tory, S. Potts, and T. Möller. A parallel coordinates style interface for exploratory volume visualization. *IEEE Transactions on Visualization and Computer Graphics*, 11(1):71–80, 2005.
- [21] P. C. Wong and R. D. Bergeron. Multiresolution multidimensional wavelet brushing. In *IEEE Visualization*, pages 141–148, 1996.
- [22] P. C. Wong and J. Thomas. Visual analytics. *IEEE Computer Graphics and Applications*, 24(5):20–21, 2004.
- [23] H. Zhou, W. Cui, H. Qu, Y. Wu, X. Yuan, and W. Zhuo. Splatting the lines in parallel coordinates. *Computer Graphics Forum*, 28(3):759–766, 2009.